

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*

August 5, 2025

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 Daniel Broxterman, Pierre-Philippe Combes, Yige Duan, Gilles Duranton, Cecile Gaubert, Guangbin Hong, Tomoya Mori, Elio Nimier-David, Yanos Zylberberg, and seminar participants at Brock University, CREI, Liaoning University, Peking University, Singapore Management University, University of Calgary, Wuhan University, the 1st Summer Meeting in Urban Economics at Peking University, the 2023 Asia Meeting of the Econometric Society in China, the 2023 Asian Meeting of the Econometric Society in East and South-east Asia, the 2023 European Meetings of the Urban Economics Association, the 2023 North American Meeting of the Urban Economics Association, the 2024 ASSA-AREUEA Annual Meeting, the Institute of Developing Economies (IDE) Workshop on Advances in Spatial Economics, and the SMU-Jinan Conference on Urban and Regional Economics. Naqun Huang and Yanmin Yang acknowledge the support from the National Natural Science Foundation of China (Grant Nos. 72003095 and 72303107). Jing Li acknowledges the support from the Singapore Ministry of Education (MOE) Academic Research Fund Tier 1 Grant C244/MSS23E006. All errors remain ours.

[†]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.

[¶]Corresponding Author. 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 [...] 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, fueled by a growing availability of effective tools of online communication, renewed concerns that the advantages traditionally associated with urban density might be persistently eroded. Whether this shift to remote communication will lead to a lasting change in the geography of innovation, or if in-person interactions in cities will ultimately prove irreplaceable, remains a widely debated question (e.g., Nathan and Overman, 2020; Barrero et al., 2021; Monte et al., 2022).

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

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 (referred to as the “density premium in innovativeness”) decreased significantly after the onset of the pandemic. Leveraging smartphone data on local mobility, we also show that a substantial part of this decline in the density premium can be attributed to the reduced frequency of face-to-face encounters. These results highlight that a key advantage of urban density lies in its capacity to create opportunities for random interactions, which facilitate the exchange and recombination of ideas across diverse fields, driving the creation of the most novel inventions. 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 with the United States Patent and Trademark Office (USPTO) between 2011 and 2021, geolocated to the level of County Sub-Divisions (CSDs). 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 (for example, between the urban core and less densely populated suburbs) 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 based on the count of new word pairs that occur in the application’s text but are not found in the full text of any earlier patents granted since 1976. Inventions with exceptionally high novelty, as captured by this measure, appear qualitatively different from the majority of inventions, as they display disproportionately high impact and breadth of diffusion.² We argue that these highly novel inventions often result from the recombination of ideas across diverse fields, which is facilitated by random interactions in dense environments. To capture local capacity for such high-

²This aligns with the well-documented observation that the distribution of technological value is highly skewed (e.g., Scherer, 1965; Schankerman and Pakes, 1986; Hall et al., 2005).

novelty inventions, we define the local degree of innovativeness as the right tail of the novelty distribution for each CSD-year observation (in our baseline specifications, the 99th and 95th percentiles).

The empirical framework is a difference-in-differences (DiD) 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 local degree of innovativeness between the years before (2011-2019) and during (2020-2021) the pandemic, 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 DiD estimates suggest that the density premium declined by approximately 20% 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 consequence of the onset of the pandemic.

Using the same empirical framework, we show that the density premium in the overall quantity of patents was not significantly impacted by the onset of COVID-19. Moreover, the change in the density premium only affected the right tail of the local novelty distribution, leaving lower percentiles such as the median and the 75th percentile unchanged. These findings suggest that the pandemic primarily disrupted dense locations' ability to foster the random and unplanned encounters underlying the most novel ideas, while providing, in the form of online communication tools, an effective substitute for interactions underlying less novel inventions.

Having established the main fact, we provide direct evidence for our candidate mechanism. Using smartphone data on local mobility provided by SafeGraph, we construct local proxies for the intensity of in-person interactions (such as prevalence of WFH and visits to dining venues) in the months surrounding the onset of COVID-19. We consistently




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. Fitted lines and 95% confidence intervals correspond to a CSD-year-level regression of innovativeness on the density decile, with standard errors clustered at the CSD level. The χ^2 test based on seemingly unrelated regressions delivers $\chi^2(1)$ equal to 36.38, and $\text{Prob} > \chi^2$ is less than 0.001. Dots and lines are shifted vertically so that innovativeness in the first decile is the same (and normalized to zero) for both sub-samples.

find that the drop in these proxies was larger in more densely populated locations. Controlling for these measures in our empirical model reduces the magnitude of the estimated effects by one-third to two-thirds relative to the baseline, suggesting 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. Other potential channels, such as differences across locations in demographic and sectoral composition, do not appear to play a meaningful role in explaining the main result.

Overall, our findings support the idea that the advantage of dense cities in the innovation process lies in their ability to foster the random and unplanned interactions underlying the most novel inventions. The widespread restrictions implemented during the pandemic resulted in a temporary loss of these innovation advantages. However, despite the widespread adoption of online communication tools, virtual interactions are unlikely to fully replace the richness and variety 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.

The remainder of the paper is organized as follows: Section 2 builds on the existing literature to provide a guiding framework for the empirical analysis; 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 Conceptual Framework and Relationship to Literature

Our paper brings together the extensive literature on the role of cities as drivers of productivity and innovation with recent quasi-experimental research on the effect of local interactions on idea diffusion and innovation outcomes.

Jacobs (1969) first popularized the intuition that urban density encourages informal and unplanned in-person interactions, facilitating the flow and recombination of ideas. Over the following decades, an extensive body of work has provided empirical support for this intuition. Glaeser et al. (1992) showed that urban variety is a strong predictor of local employment growth, highlighting the role of cross-industry spillovers as a key agglomeration force. Lin (2011) found that new occupations are more prevalent in cities with high industry variety and density of college graduates. With a specific focus on innovation activities, Carlino et al. (2007) showed that patenting rates are systematically higher in more densely populated areas. Later refinements of this finding have shown that novel ideas originate disproportionately from high-density cities (Packalen and Bhattacharya, 2015) and that inventions from dense locations tend to build on more atypical combinations of prior knowledge (Berkes and Gaetani, 2021).³

³This literature is also rich of theoretical contributions. Duranton and Puga (2001) proposed a theory in which local diversification (one of the defining features of dense cities) provides new firms with access to a broad range of ideas to experiment with in the early stages of a product cycle. Davis and Dingel (2019) developed a model in which local learning depends on how frequently people choose to interact. In equilibrium, the intensity of interactions is higher in larger cities, leading to a positive

One of the limitations of this literature is that it is mainly descriptive, and evidence on specific mechanisms is mostly indirect. More recent empirical work has used quasi-experimental settings to provide direct evidence of the effect of in-person interactions on idea diffusion and invention rates. For example, [Andrews \(2019\)](#) found that the onset of prohibition in the 1920s United States, which reduced opportunities for informal in-person interactions in bars and saloons, depressed patenting rates in dry counties relative to wet counties. Using smartphone location data to measure face-to-face interactions in Silicon Valley, [Atkin et al. \(2022\)](#) showed that random encounters among workers increase patent citations between their respective firms.⁴

This paper integrates these strands of literature and provides direct evidence on the determinants of cities' innovation advantages. Drawing from this literature, our empirical design builds upon a simple intuition. Dense locations have the advantage of hosting diverse knowledge and offering opportunities for random, unplanned in-person interactions. These interactions are crucial for high-novelty inventions, which require informal cross-field idea exchanges that are difficult for organizations to foresee and internalize. In contrast, low-novelty inventions rely on more predictable idea exchanges that are easier to plan and can occur online without a prohibitive loss of efficiency. As a consequence, a shock that widely suppresses in-person encounters and spurs online interactions will have a disproportionate negative effect on the capacity for novel inventions in high-density locations. Meanwhile, since the interactions underlying the majority of inventions can take place remotely, overall invention rates will largely remain unaffected across the density distribution. In the empirical analysis, we will test this hypothesis in the context of the sharp drop in local interactions at the onset of COVID-19.

One source of concern for our empirical design is that the pandemic may have affected innovation activities differentially across locations through channels that are unrelated to local interactions. For example, the pandemic had a disproportionate effect on some demographic groups, including women ([Myers et al., 2020](#); [Alon et al., 2022](#)) and young

relationship between productivity, skill premium, and city size. See [Kerr and Robert-Nicoud \(2020\)](#) for a comprehensive overview of this literature.

⁴Other empirical contributions to this literature include [Pennington \(2020\)](#), [Roche \(2020\)](#), and [Koh et al. \(2025\)](#).

workers in the early stages of their career (Lee et al., 2021; Walker et al., 2022). The pandemic also had a large effect on migration, with people relocating outside of the dense urban cores towards less dense suburbs (Ramani and Bloom, 2021; Coven et al., 2023), and generated unbalanced changes in demand for different sectors and technological areas. Our empirical approach allows us to confront these channels directly, and, in addition, to provide direct evidence on the role of in-person interactions by using local mobility data. In providing this evidence, we contribute to a growing literature that uses smartphone location data to measure local interactions and study their implications (e.g., Miyauchi et al., 2021; Couture et al., 2022; Monte et al., 2022).

A recent extension of the literature on the geography of innovation goes beyond the generic focus on invention and citation rates and explores the micro-level mechanisms linking the type and frequency of interactions to innovation outcomes. Leveraging a quasi-random spatial allocation of labs in a research institution in France, Catalini (2018) showed that colocation increases the likelihood of forming new collaborations, but separation does not negatively affect existing collaborations. In a field experiment in which opportunities for face-to-face encounters were randomly generated among participants to an academic conference, Lane et al. (2021) found that in-person interactions lead to more collaborations among scientists with moderately overlapping research interests, but fewer among scientists in the same field, for whom competitive effects dominate. Studying startups in a major co-working hub, Roche et al. (2024) found that randomly co-located firms exchange more knowledge, with the effect driven by the frequency of social interactions. In a comprehensive comparison of local and remote collaborations in scientific production, Lin et al. (2023) showed that remote teams are more likely to focus on late-stage technical tasks as opposed to conceptual tasks that require tacit knowledge, resulting in a lower likelihood of producing breakthrough discoveries.

Our empirical design allows us to shed light on some of these micro-level mechanisms. First, we investigate whether the effect of local interactions on innovativeness differs between in-person encounters at the workplace and those in informal settings, thereby determining the relative importance of each type of interaction. Second, we examine

whether local innovativeness is predominantly driven by face-to-face encounters with third parties outside the inventing team, as opposed to interactions among members of the team itself. Third, we explore whether restricting interactions has a larger negative effect on the novelty of inventions with a high degree of complexity, highlighting the importance of in-person communication for the transmission of complex and non-codified knowledge.

Finally, this paper relates 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). 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 learning (Emanuel et al., 2023), information sharing (Yang et al., 2022), and overall productivity (Emanuel and Harrington, 2024; Gibbs et al., 2023). Our paper contributes to this ongoing debate by showing that, in the innovation process, remote communication can only imperfectly substitute for the richness of in-person interactions provided by dense urban environments.

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, Manson et al., 2022), and two datasets on local interactions and mobility 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,⁵ 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.⁶ 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.⁷ 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.⁸

⁵CPC 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>.

⁶For 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.

⁷See Online Appendix B for details on the nature of CSDs and how the matching is conducted.

⁸The 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

Our main outcome of interest is based on a patent-level measure of novelty, constructed following 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.⁹

The construction of the measure occurs in two steps, which we briefly sketch here. Since we closely follow Arts et al. (2021), we refer readers 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 $t - 1$).¹⁰

Using this patent-level measure of novelty, we define two CSD-year-level indexes of innovativeness as the weighted 99th and 95th percentiles of the novelty distribution among all patent applications in the corresponding CSD-year. In other words, to capture the local capacity for high-novelty inventions, we measure 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 99th percentile is 872 and that based on the 95th percentile is 445.

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. We refer readers to the USPTO for examples.

⁹We 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.

¹⁰As 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 99th percentile of all the patent applications.

3.1.2 Novelty predicts impact and span of diffusion

Using citations data, we now 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.¹¹ 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 shocks and policies that affect their supply.

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

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

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 $\mathbf{1}(d(NWC_{ict}) = d)$ of the decile of patent i ’s novelty (count of new word pairs) among patents filed in the same year t and 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 (1), we restrict the sample to patents filed between 2011 and

¹¹Arts 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.

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.

Table A3 presents the regression results, which are visually complemented by Figure 2, illustrating the estimated coefficients and their 95% confidence intervals. Panels A, B, and C of Figure 2 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 8th, the 9th, and the 10th 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 9th to the 10th 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 8th 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 9th decile (0.3) and 10th 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 8th 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 9th decile (79 kilometers) and the 10th decile (117 kilometers).

In summary, exceptionally novel inventions are qualitatively different from the majority of inventions, displaying significantly higher technological impact and broader geo-




Figure 2: Novelty Predicts Impact and Span of Diffusion

Notes: This figure presents estimated coefficients and 95% confidence intervals from Equation (1).

graphical and technological diffusion. We argue that these inventions often emerge from the exchange and recombination of ideas across diverse fields, a process facilitated by random interactions in dense urban environments. Consequently, shocks that restrict such interactions, as we document in the remainder of the paper, are likely to significantly reduce the local capacity to generate these high-novelty inventions.

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 residents per square kilometer, as reported in Table A1.¹²

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 1st, 2019 to December 30th, 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 1st to May 9th, 2020 are constructed following a different methodology from that used from May 10th, 2020, onwards.¹³ Hence, to quantify the effect of COVID-19 on local mobility, we restrict the sample to the period between January 1st and April 30th in 2020, and use the same months in 2019 as a point of comparison.

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

¹²We also constructed alternative density measures by taking the population-weighted average of census tract density, or the commuting flow-weighted employment density experienced by residents in each CSD at their workplace. We provide detailed discussions in Online Appendix C. Overall, we find robust evidence when using those alternative measures as reported in Online Appendix Tables C1 and C2.

¹³See the release note in <https://docs.safegraph.com/docs/social-distancing-metrics>.

System (NAICS) code 7225 which stands for restaurants and other eating places, where the largest part of informal meetings are likely to take place.¹⁴ 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

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 COVID-19. In Section 5, we will then provide evidence that the drop in the frequency of in-person interactions can explain a substantial portion of this effect. Overall, this evidence supports the idea that the innovation advantage of dense locations lies in their capacity to generate opportunities for face-to-face interactions leading to novel inventions. While the pandemic resulted in a temporary loss of this advantage, the rise of remote communication is unlikely to permanently erode cities' centrality in the innovation process.

4.1 Empirical framework

The research design is a continuous DiD, 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:

$$I_{lt} = \alpha + \beta \log(Density_l) \times \mathbb{1}(After_t) + \eta_l + \theta_t + \varepsilon_{lt}, \quad (2)$$

where the dependent variable, I_{lt} , represents the degree of innovativeness, defined as either the 99th or the 95th 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 ($\mathbb{1}(After_t)$)

¹⁴Specifically, 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).

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_i , capturing time-invariant characteristics that affect local innovativeness, and year fixed effects, θ_t , capturing changes over time in the average degree of innovativeness.

Under the strong parallel trends assumption outlined by Callaway et al. (2024), the parameter β has a direct interpretation as a weighted mean of the average causal response of local innovativeness to a marginal increase in log-density following the COVID-19 shock. The strong parallel trends assumption requires that, at each point of the density distribution, the observed paths of outcomes reflect the average counterfactual path of other locations had they been at that point of the density distribution.¹⁵ Strong parallel trends is likely a stronger assumption than the standard parallel trends required in DiD applications with binary treatment. While this assumption cannot be directly tested, in Section 4.2.1 we will probe its plausibility by showing robust evidence of the absence of differential pre-trends across the density distribution.

4.2 Results

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

In columns 1 and 3, we leave out CSD fixed effects. This specification 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 imply that a 10% increase in population density increases the count of new word combinations at the 99th percentile by about 22.1, and that at the 95th 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

¹⁵This assumption also allows us to interpret our estimated coefficient as capturing how the level treatment effect of COVID-19 on local innovativeness varies with log-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
$\log(\text{pop density})$	221.49*** (9.19)	- -	72.80*** (5.38)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-40.56*** (6.63)	-41.12*** (6.64)	-17.28*** (5.05)	-16.72*** (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 (2). The dependent variable is the local degree of innovativeness defined as the 99th percentile value of the local novelty distribution (columns 1 and 2) or the 95th percentile value of the local novelty distribution (columns 3 and 4). Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

population density of 192 residents per square kilometer) 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 residents per square kilometer) increases the local degree of innovativeness by 1,474.4 and 484.6, respectively, corresponding to about 169% and 109% of its unconditional mean.¹⁶

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 4.2.2 below, we find no evidence that the

¹⁶This 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 premium in innovativeness is also visible with our text-based measure of novelty.

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 (2) by including a range of controls for the local number of patent applications. In Panels A and B of Table A4 we control for a linear and for a quadratic function of patent applications in each CSD-year observation, respectively. In Panels A-C of Table A5 we further include a set of five, ten, and twenty indicators of each observation’s rank in the distribution of patent count, to flexibly control for the potential bias. In all these specifications, the inclusion of these controls has no significant effect on the estimates of our coefficients of interest.¹⁷

Note that local innovativeness in Equation (2) is defined in levels, meaning that the coefficient β reflects absolute changes in tail novelty. We see this specification as the most appropriate in our context. Shocks to local interactions—which support idea exchange across fields and contribute to the development of exceptionally novel inventions—are expected to generate large shifts at the top of the novelty distribution. Measuring absolute changes captures these shifts more directly, while percentage specifications tend to give greater weight to small changes in the lower tail. To further support this approach, in Section 5 we present patent-level analyses where the outcome is an indicator for whether an invention displays an exceptionally high degree of novelty. These specifications, which avoid potential mechanical bias and are independent of specific outcome transformations, fully corroborate our baseline results.

Nonetheless, examining relative changes in innovativeness allows us to ask whether the decline in the density premium reflects a broad reduction in the productivity of dense locations, which would manifest as a disproportionate relative decline in tail novelty, or instead a more specific disruption to high-novelty projects, which would result in large absolute declines at the top without necessarily affecting relative measures. To explore this, Table A6 presents results from estimating Equation (2) using Poisson regressions (Panel A) and log-linear specifications (Panel B), where the coefficient β reflects percentage changes in tail novelty. In both cases, we do not find a statistically significant interaction

¹⁷In Online Appendix E, we provide further evidence that our results are not driven by other potential sources of mechanical bias, such as low-density CSDs having lower initial degrees of innovativeness.

between density and the post-COVID-19 indicator. These results are consistent with the idea that the COVID-19 shock did not broadly reduce the relative productivity of dense areas, but rather disrupted the types of interactions that are especially important for generating exceptionally novel inventions.

Another potential concern with the yearly specification of Equation (2) is that 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 A7, we estimate Equation (2) 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 affected the demand for new ideas, and this change in the demand for innovation 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 A7, we estimate Equation (2) at the CSD-CPC Section level, again at a monthly frequency.¹⁸ 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.

4.2.1 Testing for pre-trends

We now explore event-study variants of our empirical model, and show that there was no detectable trend in the density premium in innovativeness before the onset of the pandemic. The specification is as follows:

$$I_{it} = \alpha + \sum_{\tau \in T} \beta_{\tau} \log(\text{Density}_i) \times \mathbf{1}(t = \tau) + \eta_i + \theta_t + \varepsilon_{it}, \quad (3)$$

¹⁸CPC Sections capture nine broad technological categories assigned by the patent examiner to the patent application based on its contents.




Figure 3: Temporal Changes in Density Premium of Patent Innovativeness

Notes: This figure plots the estimated coefficients and the corresponding 95% confidence intervals from estimating Equation (3). The dependent variable is the local degree of patent innovativeness defined as the 99th percentile value of the local novelty distribution in Panel A or the 95th percentile value of the local novelty distribution in Panel B.

where T is the set of years between 2011 and 2021 (with 2019 omitted as the reference year). This specification is analogous to Equation (2), 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 3 plots the estimated coefficients, together with the corresponding 95% confidence intervals, when the local degree of innovativeness is measured as the 99th (upper panel) or the 95th (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 A2 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 99th 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 95th percentile value (Figure A3). 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.

Although event-study plots are important in evaluating the plausibility of the parallel trends assumption (Roth et al., 2023), Roth (2022) cautions that tests of pre-trends may be under-powered to detect possible violations of parallel trends. In our setting, this concern is attenuated by the relatively long pre-treatment period and short post-treatment period, which increase the probability of detecting a significant pre-trend (Roth, 2022, p. 312). However, we can formally assess the power of our test against economically relevant violations of parallel trends.

Following Roth (2022), in Figure A4, we augment the event-study plots by superimposing in the solid red lines the positive linear trends that would be detected in 50% (panels A and B) or 80% (panels C and D) of the samples. We do this separately for when the local degree of innovativeness is measured as the 99th (panels A and C) and the 95th (panels B and D) percentile value. We present analogous results for negative linear trends in Figure A5. In the same figures, we also display in the dashed blue lines the coefficients that we would expect to estimate, given the aforementioned linear trends, conditional on not detecting a significant pre-trend. In all these plots, both the solid and dashed lines are mostly not contained within the confidence intervals following the treatment. This suggests that the treatment effect is likely to be significant even under

conservative assumptions on the pre-trends.

In addition, following the approach proposed by [Rambachan and Roth \(2023\)](#), we assess the sensitivity of our conclusions to violations of parallel trends. Specifically, we report the “breakdown” value which captures how big the post-treatment violation of parallel trends would have to be (relative to the pre-treatment trend) to invalidate our conclusion. These sensitivity analyses are presented in Figures [A6](#) (the effect for the first post-treatment period) and [A7](#) (the average effect over the two post-treatment periods). We find that our results are highly robust to violations of parallel trends, with the effect of COVID-19 remaining significant unless the post-shock differences in trends are more than three times as large as the corresponding differences before COVID-19.

4.2.2 Impact on the density premium in patent quantity

These findings raise an immediate question: Did COVID-19 have a broader impact on the density premium in patent quantity, or was the effect confined to the most novel inventions? While restrictions to in-person interactions may have curtailed opportunities of random idea exchange, they may also have induced inventors to refocus their effort towards less novel inventions, for which idea exchange can occur online and is 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 [A8](#) shows estimates of regressions analogous to Equation (2), 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 A9 shows that, while there is no significant change in the overall patent count across the density distribution following the onset of COVID-19, the number of high-novelty inventions does drop significantly more in denser locations. Panel A (B) shows estimates of regressions analogous to Equation (2), but in which the outcome variable is the count of patents in the top 5% (bottom 95%) of the overall novelty distribution.¹⁹ 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 ones. Tables A10 and A11 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

An important related 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 invention. To address this question, we conduct a set of tests by choosing alternative percentile values from the local novelty distribution—the 90th, 75th, or the 50th percentile—and examine whether the density gradient of those values is also affected by the COVID-19 shock. Table A12 reports the corresponding results, showing that the coefficients associated with the interaction terms are small and not significant.

This finding is consistent with the notion that, while the random interactions promoted by urban density are crucial in driving the most novel ideas, they play a limited role in generating the majority of inventions. 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 inventions at the very top of the novelty distribution, which,

¹⁹High-novel applications are defined as those exceeding the 95th percentile of novelty, while less novel ones fall below this threshold, based on all applications filed during the pre-treatment period (2011-2019).

as shown in Section 3.1.2, are qualitatively different from less novel ones, in that they generate disproportionate technological impact and spillovers.

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. Finally, we discuss the heterogeneity of the results across different characteristics of the invention, shedding light on some of the micro-level mechanisms driving our results.

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 types 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 types 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 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.²⁰ 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 A13, 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 workplace.²¹

²⁰Safegraph 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>.

²¹Note 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>.

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%).²² 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 A13, Panels A and B in Table A14 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.

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 following 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

²²For 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>.

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:

$$y_{lt} = \alpha + \sum_{\tau \in T} \beta_{\tau} \log(\text{Density}_l) \times \mathbf{1}(t = \tau) + \eta_l + \theta_t + \varepsilon_{lt}, \quad (4)$$

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 1st and April 30th, with February 28th 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 4, 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 traveled distance equal to zero, with a median traveled 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. Figure A8 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.

In Figure 5, 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. Figure A9 shows results when weekend days are excluded from the




Figure 4: Temporal Changes in Density Gradient of the WFH Intensity

Notes: The empirical specification follows Equation (4). It reports the day-to-day changes in the density gradient of the share of devices with a median travel distance of zero kilometers in Panel A, a median travel distance of less than 2 km in Panel B, those that are completely at home in Panel C, and those that are full-time at work in 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.

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.²³ 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).²⁴

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 of in-person interactions as a result of COVID-19. Specifically, we estimate the following specification at a daily frequency between January and April 2020:

$$y_{lt} = \alpha + \delta_l \mathbb{1}(After_t) + \eta_l + \theta_t + \varepsilon_{lt}, \quad (5)$$

²³Monte et al. (2022) find that trips to the Central Business District (CBD) declined more in larger Core-Based Statistical Areas (CBSAs) relative to smaller ones at the onset of COVID-19, although the more striking difference was in the slower recovery of trips to the CBD in large CBSAs over the subsequent two years.

²⁴For 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 5: Temporal Changes in Density Gradient of the Intensity of Visits to Dining Venues

Notes: The empirical specification follows Equation (4). It reports the day-to-day changes in the density gradient of the logarithm of one plus the number of visits to dining venues 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.

where $\mathbb{1}(After_t)$ 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 2), in which we include as additional covariate the coefficient $\hat{\delta}_l$ estimated via Equation (5), interacted with the post-COVID-19 dummy variable, $\mathbb{1}(After_t)$:

$$I_{lt} = \alpha + \beta \log(Density_l) \times \mathbb{1}(After_t) + \gamma \hat{\delta}_l \times \mathbb{1}(After_t) + \eta_l + \theta_t + \varepsilon_{lt}. \quad (6)$$

A drop in the estimate of the coefficient β relative to the baseline would suggest that the decline of 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. Panels A and B show results with the outcome variable defined as the 99th and 95th 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, 3, and 4, 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 positive in column 5, where the channel is proxied by the tendency to work away from one's own residence. These patterns suggest that, conditional on density, locations that experienced a larger drop in the frequency of interactions also experienced a larger drop in their degree of innovativeness.

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).

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

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. 99th</i>						
$\mathbb{1}(\text{after 2020}) \times$						
$\log(\text{pop density})$	-41.12*** (6.64)	-27.37*** (7.18)	-38.32*** (6.75)	-27.54*** (7.18)	-36.16*** (6.71)	-25.71*** (7.89)
$\Delta \text{ dist} = 0$	-	-937.12*** (195.71)	-	-	-	-6429.42* (3563.72)
$\Delta \text{ dist} \leq 2\text{km}$	-	-	-631.83*** (158.04)	-	-	350.04 (348.55)
$\Delta \text{ home device}$	-	-	-	-920.70*** (195.35)	-	5470.53 (3526.46)
$\Delta \text{ work device}$	-	-	-	-	2439.35*** (468.84)	1623.58*** (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
<i>Panel B. 95th</i>						
$\mathbb{1}(\text{after 2020}) \times$						
$\log(\text{pop density})$	-16.72*** (5.09)	-7.60 (5.30)	-14.96*** (5.09)	-7.77 (5.31)	-14.02*** (5.08)	-5.76 (5.98)
$\Delta \text{ dist} = 0$	-	-625.49*** (150.81)	-	-	-	-6182.05** (2884.27)
$\Delta \text{ dist} \leq 2\text{km}$	-	-	-410.63*** (115.98)	-	-	269.81 (246.53)
$\Delta \text{ home device}$	-	-	-	-610.31*** (150.61)	-	5416.71* (2855.89)
$\Delta \text{ work device}$	-	-	-	-	1379.65*** (352.47)	709.03* (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 (6). The dependent variable is the local degree of innovativeness defined as the 99th percentile value of the local novelty distribution in Panel A and the 95th percentile value in Panel B. $\Delta \text{ dist} = 0$ represents the CSD-specific change in the share of mobile devices with median distance traveled equal zero. $\Delta \text{ dist} \leq 2\text{km}$ represents the CSD-specific change in the share of mobile devices with median distance traveled less than or equal to two kilometers. $\Delta \text{ home device}$ represents the CSD-specific change in the share of mobile devices at home. $\Delta \text{ work device}$ 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

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, Panels 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, conditional on density, the decline in these activities reduced innovativeness after the onset of COVID-19.²⁵ 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 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

²⁵Visitors and visits are highly correlated, so including them together as covariates in column 4 changes the sign of the estimates. However, the coefficient of $\log(Density_t) \times \mathbb{1}(After_t)$ is stable across columns 2 to 4.

Table 3: The Decline in Density Premium of Patent Innovativeness Explained by Changes in Visits to Dining Venues

	(1)	(2)	(3)	(4)
<i>Panel A. 99th</i>				
$\mathbb{1}(\text{after 2020}) \times$ $\log(\text{pop density})$	-42.02*** (6.89)	-27.73*** (7.93)	-28.17*** (7.94)	-28.34*** (7.95)
Δ visits	-	90.47*** (29.21)	-	278.28 (270.81)
Δ visitors	-	-	89.59*** (29.80)	-195.84 (277.36)
Adj. R-squared	0.57	0.57	0.57	0.57
Obs.	45,043	45,043	45,043	45,043
<i>Panel B. 95th</i>				
$\mathbb{1}(\text{after 2020}) \times$ $\log(\text{pop density})$	-17.43*** (5.28)	-12.58** (6.05)	-12.25** (6.07)	-12.19** (6.07)
Δ visits	-	30.69 (22.09)	-	-89.92 (196.69)
Δ visitors	-	-	33.54 (22.83)	125.78 (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 (6). The dependent variable is the local degree of innovativeness defined as the 99th percentile value of the local novelty distribution in Panel A and the 95th percentile value in Panel B. Δ visits represents the CSD-specific change in the logarithm of one plus the number of visits to dining venues. Δ visitors represents the CSD-specific change in the logarithm of one plus the number of visitors to dining venues. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

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

We extract additional information on inventors’ demographic characteristics from the PatentsView data and merge it with our baseline data to test for plausible alternative mechanisms. We briefly present these tests here and leave more details on data and analyses to Online Appendix D.

We first explore the hypothesis that our results are driven by the fact that women and young workers were more severely affected by the onset of COVID-19. This fact may play a role in driving our main result if female and young 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 2015.²⁶ We then run application-level regressions of the following form:

$$HN_{ilt} = \alpha + \beta \log(Density_l) \times \mathbb{1}(After_t) + \gamma_0 X_i + \gamma_1 X_i \times \mathbb{1}(After_t) + \delta Z_i + \eta_l + \theta_t + \varepsilon_{ilt}, \quad (7)$$

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.

Table 4 summarizes the key estimates from Equation (7), where high-novelty applica-

²⁶In a robustness analysis, we obtain accurate information on age for a subset of inventors by merging in the age information provided by Kaltenberg et al. (2021). Our findings remain robust across alternative definitions of young inventors based on their age, as shown in Table D5.

tions are defined as those with novelty above the 99th percentile of all applications filed during the pre-treatment period (2011-2019). Column 1 in Panel A presents the baseline results, which do not include indicators for female or young inventors. Columns 2 and 3 indicate that controlling for these factors does not significantly change the main coefficient of interest (β), indicating that the decline in the density premium is unlikely to be due to a more severe impact of the pandemic on the productivity of women or young inventors.²⁷

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 (7), in which the term X_i is this indicator for “mover” inventors. Column 4 in Panel A of Table 4 shows that 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.²⁸

Panels B and C of Table 4 provide robustness analyses to rule out other potential confounding mechanisms, including changes in the demand for innovation in specific technological areas, shifts in the political environment, and geographical variation in macroeconomic conditions and pandemic severity.²⁹

In Panel B, we show that our main results are unlikely to be driven by the pandemic’s impact on demand for inventions in specific technological areas. In column 1, we exclude patent applications containing one or more terms from the WFH technology dictionary developed by Bloom et al. (2021) in their titles or abstracts. In columns 2 and 3, we

²⁷For a comprehensive view of the results, please refer to Tables D2 - D5, which present additional findings, including coefficients on X_i and $X_i \times \mathbb{1}(After_t)$, as well as alternative high-novelty indicators based on the 95th percentile and various measures for young inventors.

²⁸Table D6 displays the complete results of this analysis.

²⁹Online Appendix E provides more details on these analyses.

exclude applications with keywords related to the pandemic (column 2) and in the top 10 CPC subclasses associated with COVID-19 diagnostic, as identified by [Toole et al. \(2023\)](#) (column 3). In column 4, we omit applications classified under Y02 (technologies or applications for mitigation or adaptation against climate change). Across these specifications, the coefficients for the interaction terms remain virtually unchanged, confirming that the observed decline in the density premium is not driven by changes in demand for innovation related to WFH technologies, COVID-19 diagnostics, or climate change.

In Panel C, we test for remaining potential channels by incorporating different sets of fixed effects. In column 1, we include state-year fixed effects to account for variation in the political environment and potential confounders at the state-year level, such as local economic conditions and pandemic severity. In columns 2-4, we introduce technology class, technology class-by-year, and technology class-by-year-month fixed effects, respectively, to control for different trends in innovation across technological areas. The results show that the decline in the density premium is large and significant across these specifications, corroborating the robustness of our findings.

Lastly, we address the potential influence of small-density CSDs with minimal or zero pre-treatment patent innovativeness, which might mechanically bias our results. To address this, we restrict the sample to CSDs with positive innovativeness during the pre-treatment period (Table [E6](#)) and exclude those in the bottom decile of population density (Table [E8](#)). Results from these robustness checks confirm that our findings are not driven by small or sparsely populated CSDs.

5.5 Heterogeneous effects: Type of collaboration, team formation, and technological complexity

As a last step, we use the application-level specification in Equation (7) to explore potential heterogeneity in the effects, shedding further light on the micro-level mechanisms behind the results. Results are summarized in Figure [6](#), and details on these analyses are provided in Online Appendix [F](#).

First, we investigate whether local innovativeness is primarily driven by face-to-face

Table 4: The Impact of COVID-19 on Density Premium (Accounting for Inventor Heterogeneity and Addressing Confounding Mechanisms)

	(1)	(2)	(3)	(4)
<i>Panel A. Control for Inventor Type</i>	Baseline	Female	Young	Mover
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.0011*** (0.0003)	-0.0010** (0.0004)	-0.0011*** (0.0003)	-0.0010*** (0.0004)
Adj. R-squared	0.020	0.021	0.020	0.020
Obs.	1,739,122	1,635,107	1,739,122	1,712,770
<i>Panel B. Exclude Patents</i>	WFH	COVID Keywords	COVID Top 10 CPC Subclasses	Climate Change
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.0011*** (0.0003)	-0.0011*** (0.0003)	-0.0008*** (0.0002)	-0.0010*** (0.0003)
Adj. R-squared	0.020	0.020	0.023	0.020
Obs.	1,737,491	1,738,504	1,458,473	1,586,124
<i>Panel C. Control for Different FE</i>	State-Year	Class	Class-Year	Class-Month
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.0009*** (0.0003)	-0.0006** (0.0003)	-0.0006** (0.0003)	-0.0006** (0.0003)
Adj. R-squared	0.020	0.052	0.052	0.049
Obs.	1,739,122	1,739,064	1,739,064	1,738,426

Notes: All regressions follow Equation (D1). The dependent variable is a dummy variable indicating whether a patent belongs to the top 1% of the novelty distribution. In Panel A, column 1 shows the baseline results without additional controls (column 1 of Table D2), column 2 controls for the presence of at least one female inventor and its interaction with the post-COVID-19 filing variable (column 2 of Table D2), column 3 includes controls for the presence of at least one young inventor and its interaction with the same filing variable (column 3 of Table D2), column 4 accounts for the presence of “mover” inventors—those who changed CSDs in 2020 or later—and their interaction with the post-COVID-19 variable (column 2 of Table D6). In Panel B, column 1 excludes patent applications related to Working from Home (WFH) technologies (column 1 of Table E1), column 2 excludes patent applications that contain at least one word related to COVID-19 diagnostics in the title or abstract (column 2 of Table E1), column 3 excludes the top 10 CPC subclasses that include COVID-19 diagnostic patents (column 1 of Table E2), column 4 excludes applications categorized under Y02, which pertains to climate change mitigation technologies (column 1 of Table E4). In Panel C, columns 1 to 4 incorporate state-by-year fixed effects, technology class fixed effects, technology class-by-year fixed effects, and technology class-by-year-month fixed effects, respectively (column 1 of Table E3, columns 1 to 3 of Panel A in Table E5). All regressions control for CSD fixed effects, year fixed effects, and team size. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

interactions with third parties outside the inventing team or by exchanges among team members. Patent applications with multiple inventors are categorized as either (1) close-distance collaborations, where at least two inventors are from the same CSD, or (2) far-distance collaborations, where all inventors are from different CSDs. While both types may benefit from external knowledge exchange, close-distance collaborations allow for more frequent in-person interactions within the team. A stronger effect for close-distance inventions would indicate that the decline in novelty is driven by the loss of in-team idea exchange, whereas no difference would suggest that external interactions play a more central role.

Panel A of Figure 6 plots the coefficients and confidence intervals for key interaction terms from the application-level specification. The results confirm a decline in the density premium (left coefficient). However, they also show that there is no difference between close-distance and remote collaborations in the impact of COVID-19 on the likelihood of high-novelty inventions (center coefficient), and this difference does not systematically vary with population density (right coefficient). This result suggests that in-person interactions with individuals external to the inventing team are a key determinant of novelty. Details of this analysis can be found in Online Appendix F.1.

Second, we examine how the COVID-19 shock influenced the stability of existing teams and the formation of new collaborations. To assess team stability, we construct an unbalanced panel at the team-by-year level and introduce a binary dependent variable, “work together,” which equals one if team k collaborates in year t , and zero otherwise. Panel B of Figure 6 shows that, after the onset of the pandemic, close-distance teams—that are more likely to rely on in-person interactions—experience a lower probability of continuing working together than far-distance teams.

To analyze new team formation, we create an outcome variable capturing the number of newly formed close-distance teams (teams with at least two inventors from the same CSD that submit their first patent application in a given year). Conducting this analysis at the CSD-by-year level, we find in Panel C of Figure 6 that while high-density areas consistently display a higher number of new teams (left coefficient), this difference de-

creases after the COVID-19 shock (right coefficient). More details on this analysis are described in Online Appendix [F.2](#).

Third, we explore whether the magnitude of the effect differs by the technological complexity of the invention. To the extent that creating complex ideas requires exchange of tacit knowledge, we should expect novelty in more complex technological areas to be more responsive to the availability of in-person interactions. To test for this channel, we use the measure of complexity developed by [Broekel \(2019\)](#) to split applications into high- and low-complexity. Using the application-level specification, Panel D of Figure [6](#) reveals that patents exhibiting higher technological complexity are more likely to rank in the top 1% in terms of novelty (left coefficient). However, the drop in the density premium following the onset of COVID-19 is larger for high-complexity inventions (right coefficient), highlighting the importance of in-person communication for the exchange of non-codifiable knowledge. Online Appendix [F.3](#) provides details on this test.

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 the onset of the pandemic 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 have implications for the spatial organization of the workforce within




Figure 6: Additional Analyses on Micro-level Mechanisms

Notes: The figures show estimated coefficients and 95% confidence intervals. Panel A is from Equation (D1), with the dependent variable indicating if a patent is in the top 1% of the novelty distribution (column 1 of Table F1). “Close” refers to close-distance collaborations with at least two inventors from the same CSD. Panel B is from team-year data and regressions with the binary dependent variable “work together,” equal to one if team k collaborates in year t (column 1 of Table F2). Panel C is from Equation (2), with the dependent variable being the number of new close-distance teams (column 1 of Table F3). New close-distance teams include those with at least two inventors from the same CSD submitting their first patent. Panel D is from Equation (D1), with the dependent variable indicating if a patent is in the top 1% of the novelty distribution (column 1 of Table F4). “Complexity” equals one if the patent’s CPC class ranks in the top 50% for complexity and zero otherwise.

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. Our results suggest that the communication constraints inherent in remote work environments could reduce innovation potential, 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.
- Alon, T., S. Coskun, M. Doepke, D. Koll, and M. Tertilt (2022). From Mancession to Shecession: Women’s Employment in Regular and Pandemic Recessions. *NBER Macroeconomics Annual* 36(1), 83–151.
- 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*.
- Barrero, J. M., N. Bloom, and S. J. Davis (2021). Why Working from Home Will Stick. *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., S. J. Davis, and Y. Zhestkova (2021). Covid-19 Shifted Patent Applications Toward Technologies that Support Working from Home. *AER Papers and Proceedings* 111, 263–266.
- 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.
- Broekel, T. (2019). Using Structural Diversity to Measure the Complexity of Technologies. *PLOS One* 14(5), e0216856.
- 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.
- Callaway, B., A. Goodman-Bacon, and P. H. Sant’Anna (2024). Difference-in-Differences with a Continuous Treatment. *Working Paper*.
- 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.
- Couture, V., J. I. Dingel, A. Green, J. Handbury, and K. R. Williams (2022). JUE Insight: Measuring Movement and Social Contact with Smartphone Data: A Real-time Application to COVID-19. *Journal of Urban Economics* 127, 103328.
- Coven, J., A. Gupta, and I. Yao (2023). JUE Insight: Urban Flight Seeded the COVID-19 Pandemic Across the United States. *Journal of Urban Economics* 133, 103489.
- 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. and E. Harrington (2024). Working Remotely? Selection, Treatment, and the Market for Remote Work. *American Economic Journal: Applied Economics* 16(4), 528–559.
- Emanuel, N., E. Harrington, and A. Pallais (2023). The Power of Proximity to Coworkers: Training for Tomorrow or Productivity Today? *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.
- Kaltenberg, M., A. B. Jaffe, and M. Lachman (2021). The Age of Invention: Matching Inventor Ages to Patents Based on Web-Scraped Sources. *Working Paper*.

- 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 (2025). Subway, collaborative matching, and innovation. *Review of Economics and Statistics* 107(2), 476–493.
- 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.
- Lane, J. N., I. Ganguli, P. Gaule, E. Guinan, and K. R. Lakhani (2021). Engineering Serendipity: When Does Knowledge Sharing Lead to Knowledge Production? *Strategic Management Journal* 42(6), 1215–1244.
- Lee, S. Y. T., M. Park, and Y. Shin (2021). Hit Harder, Recover Slower? Unequal Employment Effects of the COVID-19 Shock. *Working Paper*.
- Lin, J. (2011). Technological Adaptation, Cities, and New Work. *Review of Economics and Statistics* 93(2), 554–574.
- Lin, Y., C. B. Frey, and L. Wu (2023). Remote Collaboration Fuses Fewer Breakthrough Ideas. *Nature* 623(7989), 987–991.
- Manson, S., J. Schroeder, D. Van Riper, T. Kugler, and S. Ruggles (2022). IPUMS National Historical Geographic Information System: Version 17.0 [dataset]. Minneapolis, MN: IPUMS.
- Miyauchi, Y., K. Nakajima, and S. J. Redding (2021). The Economics of Spatial Mobility: Theory and Evidence Using Smartphone Data. *Working Paper*.
- 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*.
- Pennington, K. (2020). Varying Serendipity between Inventors. *Working Paper*.
- Ramani, A. and N. Bloom (2021). The Donut Effect of COVID-19 on Cities. *Working Paper*.
- Rambachan, A. and J. Roth (2023). A more Credible Approach to Parallel Trends. *Review of Economic Studies* 90(5), 2555–2591.
- Roche, M. P. (2020). Taking Innovation to the Streets: Microgeography, Physical Structure, and Innovation. *Review of Economics and Statistics* 102(5), 912–928.
- Roche, M. P., A. Oettl, and C. Catalini (2024). Proximate (Co-) Working: Knowledge Spillovers and Social Interactions. *Management Science*.
- Roth, J. (2022). Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends. *American Economic Review: Insights* 4(3), 305–322.
- Roth, J., P. H. Sant’Anna, A. Bilinski, and J. Poe (2023). What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature. *Journal of Econometrics* 235(2), 2218–2244.
- Saxenian, A. (1996). *Regional Advantage: Cultural and Competition in Silicon Valley and Route 128*. Cambridge, MA: Harvard University Press.
- Schankerman, M. and A. Pakes (1986). Estimates of the Value of Patent Rights in European Countries during the Post-1950 Period. *The Economic Journal* 96(384), 1052–1076.
- Scherer, F. M. (1965). Firm Size, Market Structure, Opportunity, and the Output of Patented Inventions. *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.
- Toole, A. A., N. A. Pairolero, J. Frumkin, A. V. Giczy, M. F. Moezie, J. Andres, B. Fetterolf, G. Nickol, and G. Li (2023). Diagnosing Covid-19: A Perspective from U.S. Patenting Activity. Technical report, United States Patent and Trademark Office.
- 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.
- Verendel, V. (2023). Tracking Artificial Intelligence in Climate Inventions with Patent Data. *Nature Climate Change* 13(1), 40–47.
- Walker, J., C. Brewster, R. Fontinha, W. Haak-Saheem, S. Benigni, F. Lamperti, and D. Ribaudo (2022). The Unintended consequences of the Pandemic on Non-pandemic Research Activities. *Research Policy* 51(1), 104369.
- 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	2
B	Matching Inventors' Cities with County Sub-Divisions (CSDs)	25
C	Additional Analyses Addressing Potential Measurement Concerns	27
C.1	Alternative Measures of Population Density	27
C.2	Alternative Measures of Employment Density	28
C.3	Percentage Change in Face-to-face Interactions	30
D	Additional Analyses Accounting for Inventor Heterogeneity	46
D.1	Data and Variables	46
D.2	Empirical Results	48
E	Additional Analyses Addressing Confounding Mechanisms	57
F	Additional Analyses on Micro-level Mechanisms	69
F.1	Type of Collaboration	69
F.2	Team Stability and Formation	70
F.3	Technological Complexity	71

A Online Appendix Figures and Tables




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




Figure A2: Temporal Changes in Density Premium of Patent Innovativeness (The 99th Percentile)

Notes: This figure plots the estimated coefficients and the corresponding 95% confidence intervals from estimating Equation (3) 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 99th percentile value of the local novelty distribution.




Figure A3: Temporal Changes in Density Premium of Patent Innovativeness (The 95th Percentile)

Notes: This figure plots the estimated coefficients and the corresponding 95% confidence intervals from estimating Equation (3) 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 95th percentile value of the local novelty distribution.




Figure A4: Pretests (Positive Slope)

Notes: This figure shows the pretest results using methods from Roth (2022). The slope is identified using pretests having 50 or 80 percent power.




Figure A5: Pretests (Negative Slope)

Notes: This figure shows the pretest results using methods from Roth (2022). The slope is identified using pretests having 50 or 80 percent power. Slope is negative and has the same absolute value as Figure A4.




Figure A6: Robust Inference in Difference-in-differences

Notes: This figure shows the results using robust inference and sensitivity analysis for differences-in-differences developed in [Rambachan and Roth \(2023\)](#). The sensitivity analysis imposes that the slope of post-treatment violations of parallel trends cannot change by more than M relative to the slope of the pre-trend. The figure shows the effect for the first post-treatment period.




Figure A7: Robust Inference in Difference-in-differences

Notes: This figure shows the results using robust inference and sensitivity analysis for differences-in-differences developed in [Rambachan and Roth \(2023\)](#). The sensitivity analysis imposes that the slope of post-treatment violations of parallel trends cannot change by more than M relative to the slope of the pre-trend. The figure shows the average effect over the two post-treatment periods.




Figure A8: Temporal Changes in Density Gradient of the WFH Intensity (Weekends Removed)

Notes: The empirical specification follows Equation (4). It reports the day-to-day changes in the density gradient of the share of devices with a median travel distance of zero kilometers in Panel A, a median travel distance of less than 2 km in Panel B, those that are completely at home in Panel C, and those that are full-time at work in 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 A9: Temporal Changes in Density Gradient of the Intensity of Visits to Dining Venues (Weekends Removed)

Notes: The empirical specification follows Equation (4). It reports the day-to-day changes in the density gradient of the logarithm of one plus the number of visits to dining venues 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 th percentile	872.52	1,181.97	0	3,587	46,041
Innovation index at the 95 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 95th and the 99th 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.
$\mathbb{1}(\text{Top } 1)$	0.019	0.138	0	1	794,354
$\mathbb{1}(\text{Top } 5)$	0.058	0.234	0	1	794,354
$\mathbb{1}(\text{Top } 10)$	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) 1(Top 1)	(2) 1(Top 5)	(3) 1(Top 10)	(4) Technological Diffusion	(5) Spatial Diffusion
1(4 th decile)	0.000 (0.001)	0.003* (0.001)	0.009*** (0.002)	0.076*** (0.009)	35.120*** (5.504)
1(5 th decile)	0.001 (0.001)	0.006*** (0.002)	0.013*** (0.002)	0.108*** (0.010)	44.757*** (6.497)
1(6 th decile)	0.001 (0.001)	0.007*** (0.002)	0.014*** (0.003)	0.140*** (0.011)	43.780*** (6.483)
1(7 th decile)	0.002 (0.002)	0.009*** (0.003)	0.018*** (0.004)	0.157*** (0.011)	63.310*** (6.659)
1(8 th decile)	0.002 (0.002)	0.010*** (0.003)	0.021*** (0.005)	0.201*** (0.012)	63.191*** (6.912)
1(9 th decile)	0.003 (0.002)	0.014*** (0.004)	0.029*** (0.006)	0.268*** (0.015)	79.419*** (8.068)
1(10 th decile)	0.010*** (0.003)	0.031*** (0.006)	0.055*** (0.009)	0.478*** (0.029)	116.536*** (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 (1). Standard errors clustered at the CPC technology class by filing year level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: 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
<i>Panel A. Linear</i>				
$\log(\text{pop density})$	209.89*** (10.64)	-	72.23*** (5.41)	-
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-39.13*** (6.69)	-41.02*** (6.65)	-17.21*** (5.05)	-16.77*** (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
<i>Panel B. Quadratic</i>				
$\log(\text{pop density})$	189.49*** (9.24)	-	69.78*** (5.39)	-
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-36.82*** (6.70)	-40.89*** (6.65)	-16.93*** (5.05)	-16.77*** (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 (2). The dependent variable is the local degree of innovativeness measured as the 99th percentile value of the local novelty distribution (columns 1 and 2) or the 95th 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

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
<i>Panel A. Quintile</i>				
$\log(\text{pop density})$	70.15*** (7.69)	-	35.91*** (5.32)	-
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-33.25*** (6.57)	-40.27*** (6.56)	-17.69*** (5.04)	-18.32*** (5.09)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.26	0.57	0.05	0.39
Obs.	46,041	46,041	46,041	46,041
<i>Panel B. Decile</i>				
$\log(\text{pop density})$	63.77*** (7.66)	-	34.59*** (5.31)	-
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-34.48*** (6.55)	-41.74*** (6.55)	-18.95*** (5.03)	-20.12*** (5.10)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.27	0.58	0.06	0.39
Obs.	46,041	46,041	46,041	46,041
<i>Panel C. Every 5 Pctile</i>				
$\log(\text{pop density})$	62.45*** (7.65)	-	34.74*** (5.32)	-
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-34.71*** (6.54)	-42.90*** (6.55)	-19.52*** (5.03)	-21.00*** (5.10)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.27	0.58	0.06	0.39
Obs.	46,041	46,041	46,041	46,041

Notes: All regressions follow Equation (2). The dependent variable is the local degree of innovativeness measured as the 99th percentile value of the local novelty distribution (columns 1 and 2) or the 95th percentile value of the local novelty distribution (columns 3 and 4). In Panels A-C, we include a set of five, ten, and twenty indicators of each observation's rank in the distribution of patent count in each CSD in a year. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A6: Impact of COVID-19 on Density Premium of Patent Innovativeness (Poisson Regression and Logarithm of Innovativeness)

	(1) 99 th	(2) 99 th	(3) 95 th	(4) 95 th
<i>Panel A. Poisson Regression</i>				
$\log(\text{pop density})$	0.27*** (0.01)	- -	0.17*** (0.01)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	0.01 (0.01)	0.01 (0.01)	-0.00 (0.01)	0.00 (0.01)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Log likelihood	-28,710,197	-11,424,009	-17,340,451	-8,319,378
Obs.	46,041	46,041	46,041	46,041
<i>Panel B. Logarithm of Innovativeness</i>				
$\log(\text{pop density})$	0.39*** (0.01)	- -	0.26*** (0.01)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.01 (0.01)	-0.00 (0.01)	-0.01 (0.01)	0.00 (0.01)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.10	0.58	0.06	0.46
Obs.	45,134	45,134	45,133	45,133

Notes: All regressions follow Equation (2). The dependent variable is the local degree of innovativeness measured as the 99th percentile value of the local novelty distribution (columns 1 and 2) or the 95th percentile value of the local novelty distribution (columns 3 and 4). Panel A uses Poisson regressions and Panel B uses the outcome variable as the logarithm of local innovativeness. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A7: The Impact of COVID-19 on Density Premium of Patent Innovativeness (Month Level and Month-CPC Section Level)

	(1) 99 th	(2) 99 th	(3) 95 th	(4) 95 th
<i>Panel A. Month</i>				
$\log(\text{pop density})$	141.03*** (7.96)	- -	83.94*** (4.79)	- -
$\mathbb{1}(\text{after 2020m3}) \times \log(\text{pop density})$	-27.87*** (3.00)	-35.01*** (3.05)	-17.68*** (2.75)	-20.40*** (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
<i>Panel B. Month by CPC Section</i>				
$\log(\text{pop density})$	71.58*** (4.26)	- -	57.75*** (3.54)	- -
$\mathbb{1}(\text{after 2020m3}) \times \log(\text{pop density})$	-10.56*** (1.64)	-15.33*** (1.64)	-10.03*** (1.52)	-12.95*** (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 (2). The dependent variable is the local degree of innovativeness measured as the 99th percentile value of the local novelty distribution (columns 1 and 2) or the 95th 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A8: The Impact of COVID-19 on Density Premium of Patent Counts

	(1) Count	(2) Count	(3) Weighted Count	(4) Weighted Count
<i>Panel A. Year</i>				
$\log(\text{pop density})$	40.03*** (6.02)	- -	19.37*** (3.33)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	0.10 (0.63)	0.10 (0.63)	-0.41 (0.26)	-0.41 (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
<i>Panel B. Month</i>				
$\log(\text{pop density})$	3.34*** (0.50)	- -	1.61*** (0.28)	- -
$\mathbb{1}(\text{after 2020m3}) \times \log(\text{pop density})$	0.01 (0.06)	0.01 (0.06)	-0.03 (0.02)	-0.03 (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
<i>Panel C. Month by CPC Section</i>				
$\log(\text{pop density})$	0.47*** (0.04)	- -	0.23*** (0.02)	- -
$\mathbb{1}(\text{after 2020m3}) \times \log(\text{pop density})$	0.00 (0.01)	0.00 (0.01)	-0.00 (0.00)	-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 (2) 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A9: The Impact of COVID-19 on Density Premium of Different Types of Patent Applications (Cutoff at the 95th Percentile)

	(1) Count	(2) Count	(3) Weighted Count	(4) Weighted Count
<i>Panel A. $\geq 95^{th}$ percentile novelty</i>				
$\log(\text{pop density})$	2.69*** (0.33)	- -	1.08*** (0.15)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.60*** (0.06)	-0.60*** (0.06)	-0.19*** (0.03)	-0.19*** (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
<i>Panel B. $< 95^{th}$ percentile novelty</i>				
$\log(\text{pop density})$	37.34*** (5.75)	- -	18.29*** (3.21)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	0.70 (0.65)	0.70 (0.65)	-0.22 (0.27)	-0.22 (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 (2) except that the dependent variable is the count of patents with novelty measure above (below) the 95th 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A10: The Impact of COVID-19 on Density Premium of Different Types of Patent Applications (Cutoff at the 99th Percentile)

	(1) Count	(2) Count	(3) Weighted Count	(4) Weighted Count
<i>Panel A. $\geq 99^{th}$ percentile novelty</i>				
$\log(\text{pop density})$	0.64*** (0.08)	- -	0.22*** (0.03)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.17*** (0.03)	-0.17*** (0.03)	-0.05*** (0.01)	-0.05*** (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
<i>Panel B. $< 99^{th}$ percentile novelty</i>				
$\log(\text{pop density})$	39.38*** (5.96)	- -	19.15*** (3.30)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	0.27 (0.63)	0.27 (0.63)	-0.36 (0.27)	-0.36 (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 (2) except that the dependent variable is the count of patents with novelty measure above (below) the 99th 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A11: The Impact of COVID-19 on Density Premium of Different Types of Patent Applications (Cutoff at the 90th Percentile)

	(1) Count	(2) Count	(3) Weighted Count	(4) Weighted Count
<i>Panel A. $\geq 90^{th}$ percentile novelty</i>				
$\log(\text{pop density})$	4.92*** (0.60)	- -	2.10*** (0.28)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.87*** (0.08)	-0.87*** (0.08)	-0.31*** (0.04)	-0.31*** (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
<i>Panel B. $< 90^{th}$ percentile novelty</i>				
$\log(\text{pop density})$	35.11*** (5.50)	- -	17.27*** (3.08)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	0.97 (0.63)	0.97 (0.63)	-0.09 (0.26)	-0.09 (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 (2) except that the dependent variable is the count of patents above (below) the 90th 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A12: 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
$\log(\text{pop density})$	22.89*** (3.28)	- -	-1.83 (1.44)	- -	-3.16*** (0.58)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-1.41 (3.85)	-0.90 (3.89)	0.29 (2.15)	0.44 (2.22)	-1.34 (1.28)	-1.33 (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 (2). The dependent variable is the local degree of innovativeness measured as the 90th percentile value of the local novelty distribution (columns 1 and 2), the 75th percentile value of the local novelty distribution (columns 3 and 4), or the 50th percentile value of the local novelty distribution (columns 5 and 6). Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A13: Summary Statistics for Safegraph Data (Devices)

Variable				
<i>Panel A. 2019</i>	<i>Jan</i>	<i>Feb</i>	<i>Mar</i>	<i>Apr</i>
Ratio of devices...				
dist = 0	0.36 (0.09)	0.34 (0.09)	0.31 (0.09)	0.30 (0.08)
dist $\leq$ 2km	0.47 (0.11)	0.46 (0.11)	0.43 (0.11)	0.42 (0.10)
home device	0.36 (0.09)	0.34 (0.09)	0.31 (0.09)	0.30 (0.08)
work device	0.06 (0.04)	0.06 (0.04)	0.05 (0.04)	0.06 (0.04)
Obs.	817,475	738,471	817,696	791,309
<i>Panel B. 2020</i>	<i>Jan</i>	<i>Feb</i>	<i>Mar</i>	<i>Apr</i>
Ratio of devices...				
dist = 0	0.23 (0.08)	0.25 (0.08)	0.28 (0.11)	0.37 (0.10)
dist $\leq$ 2km	0.37 (0.11)	0.39 (0.10)	0.43 (0.13)	0.54 (0.11)
home device	0.23 (0.08)	0.25 (0.08)	0.28 (0.11)	0.37 (0.10)
work device	0.08 (0.05)	0.07 (0.05)	0.05 (0.04)	0.03 (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 zero kilometers, 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 A14: Summary Statistics for Safegraph Data (Visit and Visitors to Dining Venues)

Variable				
<i>Panel A. 2019</i>				
	<i>Jan</i>	<i>Feb</i>	<i>Mar</i>	<i>Apr</i>
$\log(1 + \text{daily number of visits})$	1.48	1.51	1.53	1.54
	(2.54)	(2.59)	(2.62)	(2.64)
$\log(1 + \text{daily number of visitors})$	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
<i>Panel B. 2020</i>				
	<i>Jan</i>	<i>Feb</i>	<i>Mar</i>	<i>Apr</i>
$\log(1 + \text{daily number of visits})$	1.57	1.56	1.43	1.31
	(2.67)	(2.66)	(2.47)	(2.28)
$\log(1 + \text{daily number of visitors})$	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 dining venues.

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.³⁰ 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.³¹ 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.”³² 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.³³

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

³⁰<https://www.usgs.gov/u.s.-board-on-geographic-names/download-gnis-data>.

³¹<https://www.nhgis.org/documentation/gis-data/place-points>.

³²<https://www.usgs.gov/us-board-on-geographic-names/how-do-i>.

³³For 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 Analyses Addressing Potential Measurement Concerns

C.1 Alternative Measures of Population Density

In our baseline analysis, population density is calculated as a CSD's population divided by the CSD's land area. To further account for the heterogeneity of population density across census tracts within a CSD, we construct alternative CSD-specific population density measures by taking the population-weighted average of census tract-specific population density.

We follow three approaches that vary slightly in how we assign census tracts that cross CSD boundaries to respective CSDs and how we count the census tract-specific population density. In the first approach, we assign a census tract overlapping multiple CSDs to the CSD in which the largest share of the population in the census tract resides. We determine the share of the census tract population residing in a CSD by summing up the number of census block residents that live in the respective census tract and the CSD. The specific formula is as follows:

$$CSDdensity_i = \sum_{k \in i} \frac{TractPop_k}{\sum_{k \in i} TractPop_k} \times TractDensity_k, \quad (C1)$$

where $TractPop_k$ represents census tract k 's population size; $TractDensity_k$ is calculated as the tract population divided by the corresponding land area in square kilometers. Note that because a census tract is assigned to the CSD which has most of its population and some CSDs have no tracts assigned to them, this results in a small reduction in sample size.

In the second and third approaches, we assign overlapping census tracts to multiple CSDs based on the portion of the overlapping population. We calculate the overlapping population summing up the respective census block-specific population and use it to form the weight when calculating population-weighted CSD-specific population density. The difference between the second and third approaches is that, in the second approach, we use the original undivided census tract-specific population density. In the third approach,

we recalculate the divided census tract-population density by taking the ratio of the overlapping population to the overlapping area.

The density formula in the second approach is as follows:

$$CSDdensity_i = \sum_{k \in i} \frac{SubTractPop_k}{CSDPop_i} \times TractDensity_k, \quad (C2)$$

where *SubTract* is the intersection of a CSD and a census tract for overlapping census tracts. In this way, the area of a CSD is the same as the sum of all fully contained tracts and sub-tracts and $CSDPop_i = \sum_{k \in i} SubTractPop_k$.

The density formula for the third approach is as follows:

$$CSDdensity_i = \sum_{k \in i} \frac{SubTractPop_k}{CSDPop_i} \times SubTractDensity_k, \quad (C3)$$

where *SubTractDensity_k* represents the population density of the overlapping area.

We re-run our baseline regressions using the above three alternative measures of population density and report the results in Table C1. Panels A-C correspond to the three different approaches described above. Throughout all panels, we obtain similar results as shown in the baseline Table 1.

C.2 Alternative Measures of Employment Density

In the USPTO, we are only able to capture the inventor's locations as the cities of their residences. To mitigate the concern that the residential CSDs may not match the inventors' workplace CSDs, we follow a two-step procedure to construct an alternative commuting flow weighted employment density measure by utilizing the LEHD Origin-Destination Employment Statistics (LODES) dataset, which provides information on the commuting volumes between the origin (residential CSD) and the destination (workplace CSD).

In the first step, we aggregate the number of people working at each destination CSD to form a CSD-specific employment count, which is further divided by the size of the area to form a CSD-specific employment density measure. To assess the validity

of our employment measure, we aggregate the CSD’s employment count to the state level and contrast it with the state-level employment count reported by the Bureau of Labor Statistics (BLS) in Figure C1. The figure shows that our measure of employment matches extremely well with the state-level employment publicly released by the BLS, thereby validating the soundness of our employment density measure.

In the second step, we construct the commuting flow weighted employment density for each CSD, taking into account the share of residents commuting to other employment CSDs and the destination CSD’s employment density. In this way, although the USPTO identifies residential CSDs, we still manage to link the weighted employment density at work locations to the corresponding innovation outcomes of people residing in a CSD but commuting to work at those CSDs.

Specifically, for each residential CSD, we compute the share of jobs across all CSDs in 2011. We opted for 2011 data due to missing information for Massachusetts in 2010. Utilizing these job shares as weights, we subsequently calculate the weighted workplace CSD density using the formula below:

$$WeightedWorkplaceDensity_i = \sum_{k \in K_i} \frac{\#CSDJobs_k}{\sum_{k \in K_i} \#CSDJobs_k} \times CSDDensity_k, \quad (C4)$$

where CSD_k belongs to K_i if there exists at least a job whose workplace is CSD_k and residence is CSD_i . We compute two measures of employment density: one using all jobs and another using only high-paid jobs. High-paid jobs are defined as those with monthly earnings exceeding 3,333 USD, which is the highest level available in the LODES dataset.

We obtain robust evidence as shown in Table C2. In Panel A, we construct the CSD-specific employment count by summing up the total number of workers commuting to the CSD for work. In Panel B, we adjust the CSD-specific employment count by only considering high-paid jobs as they could be more relevant for innovation activities. We find similar declines in the density premium of patent innovativeness after the onset of COVID-19.

In addition to the robustness check on our baseline results, we take further precautions

to ensure that similar measurement concerns do not invalidate our mechanism analysis. For the first set of mechanism analyses focusing on the change in proxies of WFH intensity, we rely on Social Distancing Metrics in which locations are identified as residential locations. In Figures C2 and C3, we link this WFH intensity measure to commuting flow weighted employment density and find robust evidence that areas with a higher employment density experience a larger increase in WFH intensity.

The concern about the mismatch between residential and work locations is slightly different when it comes to the mechanism analysis involving the Weekly Pattern dataset. The Weekly Pattern dataset provides information on the frequency of visits to restaurants and other dining places. If those visits are more likely to take place at employment CSDs, the identified locations in the data could better represent workplace locations as opposed to residential locations, which do not match the locations at which the innovation outcomes are measured. To mitigate this concern, we further link the weekly visitors reported in the Weekly Pattern dataset to those visitors' residential census block groups and aggregate the census block group-specific visits to the CSD level to create a new proxy for the intensity of informal interactions based on the individuals' residential CSDs. When we replace our original measure with this new measure, we obtain robust evidence as shown in Figures C4 and C5. Moreover, the pattern is still robust when we use the commuting flow weighted employment density as the density measure, as shown in Figures C6 and C7.

C.3 Percentage Change in Face-to-face Interactions

An alternative way of measuring the change in face-to-face interactions is to use the logarithm of the share of devices (as opposed to the share of devices directly). In this way, the effects can be interpreted as percentage changes from one period to another. With this alternative measure of the change in face-to-face interactions, we find robust evidence as reported in Figures C8 and C9. The evidence is also robust when we regress the logarithm of the share of devices on the commuting flow weighted employment density measure described in Section C.2. The results are shown in Figures C10 and C11. When

we repeat our mechanism analysis as discussed in Section 5, we find similar patterns as shown in Table C3.




Figure C1: State Employment (Verification)

Notes: This white bar represents number of jobs aggregated at the state level using LODES7 data in 2011. The black bar represents BLS State Employment in December 2011 from Current Employment Statistics (CES) archived data.




Figure C2: Temporal Changes in Density Gradient of the WFH Intensity (Commuting Flow Weighted Employment Density)

Notes: The empirical specification follows Equation (4). It reports the day-to-day changes in the density gradient of the share of devices with a median travel distance of zero kilometers in Panel A, a median travel distance of less than 2 km in Panel B, those that are completely at home in Panel C, and those that are full-time at work in 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 contrast to Figure 4, where we utilize the population density, Figure C2 employs the employment density for each CSD, weighted by commuting flow.




Figure C3: Temporal Changes in Density Gradient of the WFH Intensity (Commuting Flow Weighted Employment Density, Weekends Removed)

Notes: The empirical specification follows Equation (4). It reports the day-to-day changes in the density gradient of the share of devices with a median travel distance of zero kilometers in Panel A, a median travel distance of less than 2 km in Panel B, those that are completely at home in Panel C, and those that are full-time at work in 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. In contrast to Figure A8, where we utilize the population density, Figure C3 utilizes the employment density for each CSD, weighted by commuting flow.




Figure C4: Temporal Changes in Density Gradient of the Intensity of Visits to Dining Venues (Population Density of Visitors' Residential CSDs)

Notes: The empirical specification follows Equation (4). It reports the day-to-day changes in the density gradient of the logarithm of one plus the number of visits to dining venues 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. In contrast to Figure 5, where we utilize the population density of CSDs where restaurants and other dining places are located, Figure C4 uses the population density of visitors' residential CSDs.




Figure C5: Temporal Changes in Density Gradient of the Intensity of Visits to Dining Venues (Population Density of Visitors' Residential CSDs, Weekends Removed)

Notes: The empirical specification follows Equation (4). It reports the day-to-day changes in the density gradient of the logarithm of one plus the number of visits to dining venues 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. In contrast to Figure A9, where we utilize the population density of CSDs where restaurants and other dining places are located, Figure C5 uses the population density of visitors' residential CSDs.




Figure C6: Temporal Changes in Density Gradient of the Intensity of Visits to Dining Venues (Commuting Flow Weighted Employment Density of Visitors' Residential CSDs)

Notes: The empirical specification follows Equation (4). It reports the day-to-day changes in the density gradient of the logarithm of one plus the number of visits to dining venues 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. In contrast to Figure 5, which utilizes the population density of CSDs where restaurants and other dining places are located, Figure C6 uses the visitors' residential CSDs along with their employment density of visitors' residential CSDs, weighted by commuting flow.




Figure C7: Temporal Changes in Density Gradient of the Intensity of Visits to Dining Venues (Commuting Flow Weighted Employment Density of Visitors' Residential CSDs, Weekends Removed)

Notes: The empirical specification follows Equation (4). It reports the day-to-day changes in the density gradient of the logarithm of one plus the number of visits to dining venues 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. In contrast to Figure A9, which utilizes the population density of CSDs where restaurants and other dining places are located, Figure C7 uses the employment density of visitors' residential CSDs, weighted by commuting flow.




Figure C8: Temporal Changes in Density Gradient of the Logarithm of WFH Intensity

Notes: The empirical specification follows Equation (4). It reports the day-to-day changes in the density gradient of the logarithm of the share of 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 contrast to Figure 4, where we utilize the ratio of devices, Figure C8 utilizes the logarithm of the ratio of devices.




Figure C9: Temporal Changes in Density Gradient of the Logarithm of WFH Intensity (Weekends Removed)

Notes: The empirical specification follows Equation (4). It reports the day-to-day changes in the density gradient of the logarithm of the share of devices with a median travel distance of zero kilometers in Panel A, a median travel distance of less than 2 km in Panel B, those that are completely at home in Panel C, and those that are full-time at work in 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. In contrast to Figure A8, where we utilize the ratio of devices, Figure C9 utilizes the logarithm of the ratio of devices.




Figure C10: Temporal Changes in Density Gradient of the Logarithm of WFH Intensity (Employment Density)

Notes: The empirical specification follows Equation (4). It reports the day-to-day changes in the density gradient of the logarithm of the share of devices with a median travel distance of zero kilometers in Panel A, a median travel distance of less than 2 km in Panel B, those that are completely at home in Panel C, and those that are full-time at work in 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 contrast to Figure 4, where we utilize the ratio of devices and population density, Figure C10 utilizes the logarithm of the ratio of devices and the employment density.




Figure C11: Temporal Changes in Density Gradient of the Logarithm of WFH Intensity (Employment Density, Weekends Removed)

Notes: The empirical specification follows Equation (4). It reports the day-to-day changes in the density gradient of the logarithm of the share of devices with a median travel distance of zero kilometers in Panel A, a median travel distance of less than 2 km in Panel B, those that are completely at home in Panel C, and those that are full-time at work in 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. In contrast to Figure A8, where we utilize the ratio of devices and population density, Figure C11 utilizes the logarithm of the ratio of devices and the employment density.

Table C1: Impact of COVID-19 on Density Premium of Patent Innovativeness

	(1) 99 th	(2) 99 th	(3) 95 th	(4) 95 th
<i>Panel A. First alternative population density</i>				
$\log(\text{pop density})$	236.84*** (9.06)	- -	72.60*** (5.37)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-32.15*** (6.72)	-33.24*** (6.75)	-10.02** (4.98)	-9.54* (5.05)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.08	0.57	0.02	0.39
Obs.	44,303	44,303	44,303	44,303
<i>Panel B. Second alternative population density</i>				
$\log(\text{pop density})$	240.64*** (8.55)	- -	76.10*** (5.08)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-34.25*** (6.43)	-35.52*** (6.46)	-12.54*** (4.82)	-11.94** (4.89)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.09	0.57	0.02	0.39
Obs.	46,041	46,041	46,041	46,041
<i>Panel C. Third alternative population density</i>				
$\log(\text{pop density})$	225.18*** (9.30)	- -	68.21*** (5.45)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-29.23*** (6.81)	-30.37*** (6.86)	-9.59* (5.03)	-8.93* (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

Notes: All regressions follow Equation (2). The dependent variable is the local degree of innovativeness defined as the 99th percentile value of the local novelty distribution (columns 1 and 2) or the 95th percentile value of the local novelty distribution (columns 3 and 4). Panel A uses population weighted tract density measure where tracts are uniquely assigned to the CSD with the highest share of population, as outlined in Equation (C1). Panel B uses population weighted tract density measure where the population is counted at the sub-tract level within CSD while density is still constructed at the tract level, as described in Equation (C2). Panel C uses population weighted density measure where both population and density are constructed at the sub-tract level within each CSD, as calculated in Equation (C3). Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C2: Impact of COVID-19 on Density Premium of Patent Innovativeness Using Workplace Weighted Density

	(1) 99 th	(2) 99 th	(3) 95 th	(4) 95 th
<i>Panel A. All jobs</i>				
$\log(\text{workplace emp density})$	401.38*** (15.89)	- -	154.43*** (9.48)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{workplace emp density})$	-87.84*** (11.88)	-89.41*** (11.83)	-41.81*** (9.06)	-41.56*** (9.09)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.10	0.57	0.04	0.39
Obs.	46,041	46,041	46,041	46,041
<i>Panel B. High-paid jobs</i>				
$\log(\text{workplace emp density})$	335.68*** (13.04)	- -	133.13*** (7.78)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{workplace emp density})$	-74.76*** (9.84)	-76.24*** (9.78)	-36.18*** (7.51)	-36.11*** (7.54)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.10	0.57	0.04	0.39
Obs.	46,041	46,041	46,041	46,041

Notes: All regressions follow Equation (2). The dependent variable is the local degree of innovativeness defined as the 99th percentile value of the local novelty distribution (columns 1 and 2) or the 95th percentile value of the local novelty distribution (columns 3 and 4). Compared to Table 1, Table C2 uses employment density, as outlined in Equation (C4). Panel A includes all jobs and Panel B includes only high-paid jobs. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C3: The Decline in Density Premium of Patent Innovativeness Explained by Changes in the Logarithm of WFH Intensity

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. 99th</i>						
$\mathbb{1}(\text{after 2020}) \times$						
$\log(\text{pop density})$	-41.12*** (6.64)	-32.37*** (6.93)	-41.69*** (6.69)	-32.45*** (6.94)	-32.27*** (6.80)	-33.18*** (7.41)
$\Delta \text{ dist} = 0$	-	-228.18*** (56.27)	-	-	-	-620.80 (810.36)
$\Delta \text{ dist} \leq 2\text{km}$	-	-	-253.96*** (65.62)	-	-	-53.03 (121.57)
$\Delta \text{ home device}$	-	-	-	-223.77*** (55.92)	-	615.20 (797.79)
$\Delta \text{ work device}$	-	-	-	-	237.86*** (40.76)	216.45*** (49.54)
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
<i>Panel B. 95th</i>						
$\mathbb{1}(\text{after 2020}) \times$						
$\log(\text{pop density})$	-16.72*** (5.09)	-10.70** (5.17)	-17.15*** (5.12)	-10.81** (5.18)	-12.06** (5.13)	-9.80* (5.63)
$\Delta \text{ dist} = 0$	-	-157.98*** (43.30)	-	-	-	-831.69 (706.50)
$\Delta \text{ dist} \leq 2\text{km}$	-	-	-154.32*** (48.13)	-	-	26.21 (86.76)
$\Delta \text{ home device}$	-	-	-	-153.58*** (43.28)	-	731.68 (698.02)
$\Delta \text{ work device}$	-	-	-	-	128.49*** (30.93)	90.98** (35.60)
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 (6). The dependent variable is the local degree of innovativeness defined as the 99th percentile value of the local novelty distribution in Panel A and the 95th percentile value in Panel B. $\Delta \text{ dist} = 0$ represents the CSD-specific change in the logarithm of the share of mobile devices with a median travel distance of zero. $\Delta \text{ dist} \leq 2\text{km}$ represents the CSD-specific change in the logarithm of the share of mobile devices with median distance traveled less than or equal to two kilometers. $\Delta \text{ home device}$ represents the CSD-specific change in the logarithm of the share of mobile devices at home. $\Delta \text{ work device}$ represents the CSD-specific change in the logarithm of the share of mobile devices outside the home during work hours. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

D Additional Analyses 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 had 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).³⁴

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.

D.1 Data and Variables

The analysis relies on the patent application dataset provided by PatentsView.³⁵ 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

³⁴Similar 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.

³⁵The 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 team 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 in the top 1% or the top 5% in terms of its novelty. The cutoff value is taken as the values at the 99th and the 95th 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 D1. 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 filing 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.

D.2 Empirical Results

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

$$HN_{ilt} = \alpha + \beta \log(Density_l) \times \mathbb{1}(After_t) + \gamma_0 X_i + \gamma_1 X_i \times \mathbb{1}(After_t) + \delta Z_i + \eta_l + \theta_t + \varepsilon_{ilt}, \quad (D1)$$

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 D2 and D3 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 D2 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 D3 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 D2 and Table D3 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 D2 and Table D3,

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 to 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 D4. Columns 1-2 of Table D4 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.

Furthermore, we utilize the age information provided in [Kaltenberg et al. \(2021\)](#) and introduce a new dummy variable, “young,” indicating whether the patent application was

filed by an inventor under the age of 30.³⁶ The results of our analysis based on this new variable are presented in Table D5. As a reference, columns 1 and 3 report the baseline results, akin to column 1 of Tables D2 and D3, respectively. The decrease in the number of observations in Table D5, compared to that in the corresponding Tables D2 and D3, is due to the absence of inventors beyond the scope of the data provided by Kaltenberg et al. (2021). Correspondingly, the number of observations in columns 2 and 4 is smaller than that in columns 1 and 3. Nevertheless, our findings remain consistent across this alternative definition of young inventors, as shown in the table.

Finally, 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 in or after 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 D6. Our baseline findings on the reduced density premium in patent informativeness remain robust.

³⁶This study provides a measure of birth year for inventors listed on patents granted before 2018. The sample period of their study does not align with ours, but we can narrow down our focus to the subset of inventors who filed patent applications before 2018 and continued through the post-COVID era.

Table D1: 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 D2: 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)
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.001*** (0.000)	-0.001** (0.000)	-0.001*** (0.000)	-0.001** (0.000)
$\mathbb{1}(\text{female})$	-	0.006*** (0.001)	-	0.006*** (0.001)
$\mathbb{1}(\text{after 2020}) \times \mathbb{1}(\text{female})$	-	-0.004*** (0.001)	-	-0.004*** (0.001)
$\mathbb{1}(\text{young})$	-	-	-0.002*** (0.000)	-0.002*** (0.000)
$\mathbb{1}(\text{after 2020}) \times \mathbb{1}(\text{young})$	-	-	-0.001** (0.001)	-0.002** (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 (D1). 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table D3: 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)
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.004*** (0.001)	-0.003*** (0.001)	-0.004*** (0.001)	-0.003*** (0.001)
$\mathbb{1}(\text{female})$	-	0.027*** (0.003)	-	0.027*** (0.003)
$\mathbb{1}(\text{after 2020}) \times \mathbb{1}(\text{female})$	-	-0.012*** (0.002)	-	-0.011*** (0.002)
$\mathbb{1}(\text{young})$	-	-	-0.001 (0.001)	-0.002** (0.001)
$\mathbb{1}(\text{after 2020}) \times \mathbb{1}(\text{young})$	-	-	-0.008*** (0.002)	-0.008*** (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 (D1). 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table D4: 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%
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.001** (0.000)	-0.001** (0.000)	-0.003*** (0.001)	-0.003*** (0.001)
$\mathbb{1}(\text{female})$	0.006*** (0.001)	0.006*** (0.001)	0.028*** (0.003)	0.027*** (0.002)
$\mathbb{1}(\text{after 2020}) \times \mathbb{1}(\text{female})$	-0.004*** (0.001)	-0.004*** (0.001)	-0.011*** (0.002)	-0.011*** (0.002)
$\mathbb{1}(\text{young}(2005))$	-0.003*** (0.000)	- -	-0.006*** (0.001)	- -
$\mathbb{1}(\text{after 2020}) \times \mathbb{1}(\text{young}(2005))$	-0.002** (0.001)	- -	-0.007*** (0.002)	- -
$\mathbb{1}(\text{young}(2015))$	- -	0.000 (0.001)	- -	0.005*** (0.001)
$\mathbb{1}(\text{after 2020}) \times \mathbb{1}(\text{young}(2015))$	- -	-0.003*** (0.001)	- -	-0.011*** (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 (D1). 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table D5: The Impact of COVID-19 on Density Premium of the Likelihood of Being the Top 1% and 5% in Novelty (Controls for Young Inventors, Alternative Measures)

	(1) 99 th	(2) 99 th	(3) 95 th	(4) 95 th
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.001*** (0.000)	-0.002*** (0.001)	-0.004*** (0.001)	-0.005*** (0.001)
$\mathbb{1}(\text{young})$	- (0.001)	-0.006*** (0.001)	- (0.001)	-0.018*** (0.003)
$\mathbb{1}(\text{after 2020}) \times \mathbb{1}(\text{young})$	- (0.001)	-0.003** (0.001)	- (0.001)	-0.011** (0.004)
Team size	Yes	Yes	Yes	Yes
CSD FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Mean of DV	0.010	0.011	0.049	0.053
Adj. R-squared	0.020	0.021	0.040	0.042
Obs.	1,739,122	1,098,622	1,739,122	1,098,622

Notes: All regressions follow Equation (D1). The dependent variable in columns 1 and 2 represents innovativeness at the top 1%, while in columns 3 and 4, it represents innovativeness at the top 5%. The “young” variable, a new dummy constructed based on [Kaltenberg et al. \(2021\)](#), indicates whether there is an inventor under the age of 30. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table D6: 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%
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.001*** (0.000)	-0.001*** (0.000)	-0.004*** (0.001)	-0.004*** (0.001)
$\mathbb{1}(\text{mover})$	- (0.001)	0.001 (0.001)	- (0.001)	0.003** (0.001)
$\mathbb{1}(\text{after 2020}) \times \mathbb{1}(\text{mover})$	- (0.001)	-0.002*** (0.001)	- (0.001)	-0.006*** (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 (D1). 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 in or after 2020. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

E Additional Analyses Addressing Confounding Mechanisms

We conduct a set of robustness analyses to rule out potential confounding mechanisms related to concurrent changes in the landscape of technology driven by the pandemic, climate change, and shifts in the political environment. We also address additional concerns on whether the baseline results are mechanically caused by the small-density CSDs experiencing almost no change in patent innovativeness.

First, we provide evidence that our findings remain robust even when we exclude patent applications related to technologies that support working from home. To achieve this, we utilize the dictionary of terms related to WFH technologies from [Bloom et al. \(2021\)](#).³⁷ We exclude patent applications that contain one or more terms from the aforementioned dictionary in their titles or abstracts and report the results in columns 1 and 3 of Table [E1](#). The coefficients of the interaction terms barely change. One possible reason for this is that the proportion of patent applications related to WFH technologies is small, around 1%, as documented in [Bloom et al. \(2021\)](#), and these may not represent the most innovative inventions.

Second, we show that our baseline results are not driven by the boom of bio-innovation related to COVID-19 diagnostics during the pandemic. To achieve this goal, we identify COVID-19-related patent applications by applying the methods outlined by [Toole et al. \(2023\)](#), authored by the United States Patent and Trademark Office’s Office of the Chief Economist, the USPTO’s Patents Business Unit, and the Office of Policy and International Affairs. They first identify COVID-19 patent applications that contain specific keywords in the title and abstract.³⁸ We utilize their specified keywords to exclude

³⁷Specifically, the dictionary of WFH terms: telecommuting, telework, teleworking, working from home, mobile work, remote work, flexible workplace, work from home, mobile working, remote working, work remotely, working remotely, remote workplace, telecommuter, teleworker, home-sourced worker, home-sourced employee, work-at-home, work at home, telecommuting specialist, nomadic worker, nomadic employee, work-from-home, work-from-anywhere, video conference, video conferencing, virtual office, distance work, flexible work, virtual work, virtual office, virtual employee, home office, home-based office, home-based work, work from anywhere, working from anywhere, work-from-anywhere, digital workplace, video chat, video call, teleconference, teleconferencing, working from a remote location, and work from a remote location.

³⁸Specifically, the keywords to identify COVID-19 diagnostics patent applications include: “covid19”, “covid-19” or variations, “covid 19” or variations (with a space), “sars-cov-2” or variations, “sars cov-2” or variations (with a first space), “sars cov-2” or variations (with a second space), “sars cov-2” or variations (with multiple spaces), “covid”, “cov2” , “ncov2” or variations, variations of “severe acute respiratory

patents that contain at least one of these words in the title or abstract. The results, reported in columns 2 and 4 of Table E1, confirm that our findings remain robust even after excluding these COVID-19-related patent applications.

In addition, Toole et al. (2023) also identifies a COVID-19 diagnostic patent application—defined as any published patent application that aids in detecting and identifying the SARS-CoV-2 virus—and then identifies the top 10 Cooperative Patent Classification (CPC) subclasses containing COVID-19 diagnostic patent applications.³⁹ We exclude these subclasses related to COVID-19 diagnostics in Table E2. The results confirm that the decline in the urban density premium is robust, even when excluding patent applications that facilitate the detection and identification of the SARS-CoV-2 virus.

Third, there might be a concern related to the change in the political environment during the COVID-19. The inclusion of year fixed effects in the specification should account for any macroeconomic conditions that could affect innovations, including the presidential election. However, since the political environment might have affected states to different extent, we run specifications including state-year fixed effects. This specification also controls for potential confounders at the state-year level, such as local economic conditions and the severity of the pandemic. As shown in Table E3, the results remain robust.

Fourth, to address concerns regarding concurrent changes in environment and climate change-related innovation, we follow the literature by labeling patents as climate inventions if they are classified under the CPC class Y02 (e.g., Verendel, 2023). The USPTO website describes class Y02 (technologies or applications for mitigation or adaptation against climate change) as follows: “This class covers selected technologies, which control, reduce or prevent anthropogenic emissions of greenhouse gases [GHG], in the framework of the Kyoto Protocol and the Paris Agreement, and also technologies which

syndrome virus 2019”, “Coronavirus 2019” and variations, “Wuhan pneumonia” and variations, “Wuhan coronavirus” and variations, “wuhan-hu-1”, variations of “novel coronavirus”, and variations of “new coronavirus.”

³⁹G01N (Analyze materials by chemical/physical properties), C12Q (Measuring enzymes, nucleic acids, microorganisms), A61B (Diagnosis/surgery/identification), A61K (Preparations for medical, dental, toiletry purposes), C12N (Microorganisms/enzymes), C07K (Peptides), G16H (Health care informatics), B01L (Chemical/physical lab apparatus), A61P (Therapeutics of chemical/medicinal preparations), and G06T (Image data processing).

allow adapting to the adverse effects of climate change.”⁴⁰ The results, excluding patent applications that include Y02 among their CPC classes, are presented in Table E4 and demonstrate a robust decline in the density premium.

Fifth, to address concerns regarding potential changes in the nature of innovative activities across different types of patent applications following COVID-19, we have included controls for technology class-by-year or by year-month fixed effects to absorb technology-specific intertemporal changes, as reported in Table E5. Specifically, we control for technology class fixed effects in column 1, technology class-by-year fixed effects in column 2, and technology class-by-year-month fixed effects in column 3. These fixed effects control for any heterogeneous changes across technology classes. The dependent variable in Panel A is a dummy variable indicating whether a patent belongs to the top 1% of the novelty distribution, and in Panel B, it represents patents within the top 5%. Our analysis demonstrates that the observed decline in the density premium remains robust, even when accounting for potential temporal variations across different technologies.

Lastly, there might also be a concern regarding the potential influence of small-density CSDs with close to zero patent innovativeness throughout, thereby mechanically contributing to our baseline finding. We address this concern as follows. We first remove observations subject to the truncated nature of the measure in a robustness analysis based only on CSDs with positive innovativeness in all years from 2011 to 2019 (pre-COVID-19). Results are displayed in Table E6. We obtain robust evidence in this robustness check, which ensures that our findings are not affected by the CSDs with zero innovativeness to start with. In the data, the average of the CSD-specific patent innovativeness index based on the 99th percentile value within the lowest decile, for example, falls between 300 and 400. In Table E7, we present the means of the local degree of innovativeness index for CSDs across different deciles based on their population density before (2011-2019) and after (2020-2021) COVID-19. Panel A illustrates the mean of the innovativeness measure based on the 99th percentile value, while Panel B illustrates that based on the 95th percentile value. The table shows that the change in the density premium of patent

⁴⁰<https://www.uspto.gov/web/patents/classification/cpc/html/cpc-Y.html>.

innovativeness is not driven by the fact that small-density CSDs have close to zero innovativeness index (hence experiencing no decline) and other regions experience an equal decline in the innovativeness index. Nevertheless, to further addressing this concern, in Table E8 we display results from another robustness check in which we exclude CSDs in the bottom decile of population density, and show that our results are not driven by those small CSDs.

Table E1: The Impact of COVID-19 on Density Premium of the Likelihood of Being the Top 1% and 5% in Novelty (Excluding Patent Applications Including Words Related to WFH or COVID-19 Technologies)

	(1) 99 th , ex WFH	(2) 99 th , ex COVID-19	(3) 95 th , ex WFH	(4) 95 th , ex COVID-19
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.001*** (0.000)	-0.001*** (0.000)	-0.004*** (0.001)	-0.004*** (0.001)
Team size	Yes	Yes	Yes	Yes
CSD FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.020	0.020	0.040	0.040
Obs.	1,737,491	1,738,504	1,737,491	1,738,504

Notes: All regressions follow Equation (D1). 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). Columns 1 and 3 exclude patent applications that advance Working from Home (WFH) technologies from Bloom et al. (2021). Columns 2 and 4 exclude COVID-19 diagnostic patent applications identified by Toole et al. (2023). Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E2: The Impact of COVID-19 on Density Premium of the Likelihood of Being the Top 1% and 5% in Novelty (Exclude Top 10 CPC Subclasses related to COVID-19 diagnostics)

	(1) 99 th	(2) 95 th
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.001*** (0.000)	-0.003*** (0.001)
Team size	Yes	Yes
CSD FE	Yes	Yes
Year FE	Yes	Yes
Adj. R-squared	0.023	0.034
Obs.	1,458,473	1,458,473

Notes: All regressions follow Equation (D1). The dependent variable is a dummy variable indicating whether a patent belongs to the top 1% (column 1) or the top 5% of the novelty distribution (column 2). This table excludes the top 10 CPC subclasses that contain COVID-19 diagnostic patent applications from [Toole et al. \(2023\)](#). Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E3: The Impact of COVID-19 on Density Premium of the Likelihood of Being the Top 1% and 5% in Novelty (State by Year Fixed Effect)

	(1) 99 th	(2) 95 th
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.001*** (0.000)	-0.003*** (0.001)
Team size	Yes	Yes
CSD FE	Yes	Yes
State by Year FE	Yes	Yes
Adj. R-squared	0.020	0.040
Obs.	1,739,122	1,739,122

Notes: All regressions follow Equation (D1) but controls for state by year fixed effects. The dependent variable is a dummy variable indicating whether a patent belongs to the top 1% (column 1) or the top 5% of the novelty distribution (column 2). Standard errors clustered at the CSD level are reported in parentheses.

Table E4: The Impact of COVID-19 on Density Premium of the Likelihood of Being the Top 1% and 5% in Novelty (Exclude CPC Class Y02)

	(1) 99 th	(2) 95 th
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.001*** (0.000)	-0.003*** (0.001)
Team size	Yes	Yes
CSD FE	Yes	Yes
Year FE	Yes	Yes
Adj. R-squared	0.020	0.040
Obs.	1,586,124	1,586,124

Notes: All regressions follow Equation (D1). The dependent variable is a dummy variable indicating whether a patent belongs to the top 1% (column 1) or the top 5% of the novelty distribution (column 2). This table excludes patent applications that include Y02 among their CPC classes. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E5: The Impact of COVID-19 on Density Premium of the Likelihood of Being the Top 1% and 5% in Novelty (Class, Class by Year, Class by Month Fixed Effect)

	(1)	(2)	(3)
<i>Panel A. Top 1%</i>			
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.001*** (0.000)	-0.001** (0.000)	-0.001** (0.000)
Team size	Yes	Yes	Yes
CSD FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
CPC Class FE	Yes	No	No
Class by Year FE	No	Yes	No
Class by Month FE	No	No	Yes
Adj. R-squared	0.051	0.052	0.049
Obs.	1,739,071	1,739,064	1,738,426
<i>Panel B. Top 5%</i>			
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.003*** (0.001)	-0.003*** (0.001)	-0.002*** (0.001)
Team size	Yes	Yes	Yes
CSD FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
CPC Class FE	Yes	No	No
Class by Year FE	No	Yes	No
Class by Month FE	No	No	Yes
Adj. R-squared	0.130	0.131	0.131
Obs.	1,739,071	1,739,064	1,738,426

Notes: All regressions follow Equation (D1). The dependent variable in Panel A is a dummy variable indicating whether a patent belongs to the top 1%, and the dependent variable in Panel B is the top 5% of the novelty distribution. Columns 1 to 3 include technology class fixed effects, technology class by year fixed effects, and technology class by year-month fixed effects, respectively. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E6: Impact of COVID-19 on Density Premium of Patent Innovativeness (CSDs with Positive Innovativeness Index from 2011 to 2019)

	(1) 99 th	(2) 99 th	(3) 95 th	(4) 95 th
$\log(\text{pop density})$	217.40*** (10.03)	- -	67.98*** (5.93)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-36.71*** (7.42)	-36.45*** (7.41)	-13.66** (5.61)	-12.90** (5.63)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.06	0.56	0.02	0.39
Obs.	41,147	41,147	41,147	41,147

Notes: All regressions follow Equation (2). The dependent variable is the local degree of innovativeness defined as the 99th percentile value of the local novelty distribution (columns 1 and 2) or the 95th percentile value of the local novelty distribution (columns 3 and 4). This table includes only cities whose minimum value of the innovativeness index from 2011 to 2019 is above zero. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E7: Mean Value of Innovative index by Deciles

	Number of CSD	Mean (2011-2019)	Mean (2020-2021)
<i>Panel A. 99th Innovativeness Index</i>			
Group 1	421	396.366	334.182
Group 2	420	480.963	341.573
Group 3	420	627.683	463.098
Group 4	420	765.497	585.848
Group 5	420	1,018.711	802.409
Group 6	420	1,007.415	830.073
Group 7	420	942.736	777.617
Group 8	420	1,085.292	878.068
Group 9	420	1,226.523	933.634
Group 10	420	1,509.856	1,211.307
<i>Panel B. 95th Innovativeness Index</i>			
Group 1	421	301.767	245.942
Group 2	420	332.023	236.735
Group 3	420	379.747	296.293
Group 4	420	432.542	322.683
Group 5	420	532.725	385.178
Group 6	420	462.418	369.823
Group 7	420	440.697	376.390
Group 8	420	487.818	367.291
Group 9	420	549.629	398.766
Group 10	420	726.108	565.857

Notes: This table displays the mean value of the innovativeness index for patent applications filed before COVID-19 (2011-2019) and after COVID-19 (2020-2021).

Table E8: Impact of COVID-19 on Density Premium of Patent Innovativeness (Exclude CSDs in the Bottom Decile of Density)

	(1) 99 th	(2) 99 th	(3) 95 th	(4) 95 th
$\log(\text{pop density})$	252.30*** (12.66)	- -	83.75*** (7.57)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-37.20*** (9.05)	-37.46*** (9.05)	-15.81** (6.84)	-15.36** (6.86)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.06	0.57	0.02	0.40
Obs.	41,463	41,463	41,463	41,463

Notes: All regressions follow Equation (2). The dependent variable is the local degree of innovativeness defined as the 99th percentile value of the local novelty distribution (columns 1 and 2) or the 95th percentile value of the local novelty distribution (columns 3 and 4). This table excludes CSDs within the bottom decile of population density. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

F Additional Analyses on Micro-level Mechanisms

In this section, we present additional analyses to uncover more detailed micro-level mechanisms through which population density affects the novelty of patents.

F.1 Type of Collaboration

In our first analysis, we focus on the novelty of collaborated patents to examine the relative importance of in-person interactions internal versus external to the collaboration team. We start by splitting the sample into two categories for patents with multiple inventors: (1) Close-distance collaboration applications, which refer to patent applications where at least two inventors are from the same CSD, and (2) Far-distance collaboration applications, which refer to patent applications with all inventors from different CSDs. The two groups have a comparable number of observations, comprising 618,738 close-distance collaboration applications and 511,830 far-distance collaboration applications.

For close-distance collaborations, knowledge could be primarily created through frequent face-to-face interactions among collaborators, complemented by necessary online communication tools. In contrast, far-distance collaborations likely rely primarily on online communications via phone, text, email, and video with infrequent face-to-face communications via travel by airplane or car. For both types of collaborations, knowledge could also be acquired from interactions with individuals external to the inventing team. These interactions are typically expected to be in-person and to occur within the inventors' respective CSDs. Hence, the key distinction between close-distance and far-distance collaborations lies in the frequency of face-to-face interactions between inventors within the same team.

This argument has two main testable implications. If the main driver of novelty of collaborated patents lies in the in-person interactions within the inventing team, we should expect to observe that, after the onset of COVID-19, (1) novelty decreases more for close-distance than for far-distance collaborations, and (2) the *density premium* in novelty decreases more for close-distance than for far-distance collaborations. By contrast, if the main driver of novelty lies in idea exchanges with people external to the inventing

team, we should expect to observe no differences between close-distance and far-distance collaborations in response to COVID-19.

We test these implications by adding additional interaction terms to our patent application level regression and report the results in Table F1. The first row shows that the impact of COVID-19 on the probability of being in the top 1% or the top 5% in novelty does not differ significantly between close-distance and far-distance collaboration applications. The third row shows that the change in the density premium after the onset of COVID-19 was not significantly different between close-distance and far-distance collaboration applications.

Overall, this evidence suggests that in-person interactions with individuals external to the inventing team are likely to drive novelty in collaborated inventions more than in-person interactions among collaborators.

F.2 Team Stability and Formation

In this section, we explore whether COVID-19 disrupted the collaboration of existing teams and impeded the formation of new teams.

We start by defining a team as a unique combination of inventors. Subsequently, we expand and reformat the dataset to include entries for each team from its inception in our records. This augmented dataset comprises over 4,453,245 observations. We then introduce a binary variable, “work together”, which takes a value of one if team k collaborates in year t , and zero otherwise. Leveraging the panel structure at the team-year level, we incorporate team-fixed effects to control for any time-invariant characteristics unique to each team.

The findings are presented in Table F2. The analysis reveals that teams with at least two inventors located in the same CSD exhibit a lower propensity to sustain patenting compared to far-distance teams after the onset of COVID-19. This suggests that close-distance teams, who rely more on in-person interactions to continue collaboration, were more negatively affected by COVID-19 in the probability of continuing working together compared to far-distance teams.

In addition, we examine whether high-density areas foster the formation of new teams and whether such benefit decreases following the onset of COVID-19. We show in Table F3 that the number of new teams formed between inventors within the same CSD increases with the population density of the CSD and the gradient decreases after COVID-19. The evidence is in line with our hypothesis that high-density areas create abundant in-person meeting opportunities that bring inventors together to create new ideas and collaborate. The decrease in the likelihood of forming new collaborations could be one of the channels contributing to the decline in the density premium of patent innovativeness.

F.3 Technological Complexity

In our final analysis, we explore the heterogeneity across technology classes characterized by different extents of knowledge complexity. To this end, we classify technology classes based on the complexity score introduced by Broekel (2019). This approach models technologies as combinatorial networks and derives a measure of technological complexity based on the diversity of sub-network topologies in these networks. Broekel (2019) publishes the complexity scores for 646 four-digit CPC classes, with positive scores for 633 CPC classes and zero scores for 13 classes, leaving the remaining 9 classes with missing scores. The most complex CPC class pertains to the propulsion of electrically-propelled vehicles.

By aligning our data with the list of complexity scores specific to each four-digit CPC class, we obtained technological complexity values for all patent applications within the 646 subclasses. We use this data to conduct the following analysis.

Firstly, we classify applications as more complex if their corresponding complexity values rank in the top 50% four-digit CPC subclasses (top 323 four-digit CPC subclass) and as less complex if they rank in the bottom 50% four-digit CPC subclasses (bottom 323 four-digit CPC subclass). The findings, presented in Table F4, reveal a positive coefficient for the complexity dummy, suggesting that patents exhibiting higher technological complexity are more likely to rank in the top 1% or 5% in terms of novelty. However, the negative coefficient for the interaction term indicates that, post-COVID-19, these

complex patents are less likely to rank among the top patents compared to others. Given the potential correlation between complexity and team size, we present results both with and without controlling for team size. The outcomes remain largely consistent across both specifications. Furthermore, we employ the continuous complexity measure directly from Broekel (2019) and report the results in Table F5. Consistent with the findings in Table F4, we observe that higher technological complexity is associated with a greater likelihood of ranking among the top patents and this likelihood decreases substantially post-COVID-19.

Secondly, we explore whether patent classes with higher complexity exhibit more co-authors or an increased average geographical distance between co-authors. Findings are presented in Table F6. Columns 1 to 3 utilize various measures of inventor team size as outcome variables. Column 1 employs the logarithm of the number of inventors per patent application, column 2 creates a dummy variable indicating whether the number of inventors exceeds 3 (approximately the mean team size), and column 3 is a binary indicator for patents with only one inventor. The coefficients suggest that patent applications in more complex technological classes tend to have larger team sizes, a higher likelihood of having more than three inventors, and are less likely to have only one inventor. However, we do not observe a reduction in the number of co-authors for patents in more complex classes after the onset of COVID-19. This result can be interpreted as suggesting that while team sizes might decrease for more complex technological classes due to certain factors, they may also increase for reasons such as the desire to push the boundaries of knowledge by involving more inventors.

To capture the distance between coauthors, we adopt the criterion outlined in Section F.1, wherein close-distance collaboration applications refer to patent applications where a minimum of two inventors are from the same CSD. Column 4 of Table F6 presents the findings using this dummy variable, reflecting close-distance collaboration applications as the outcome variable. Since we are comparing close-distance with far-distance collaboration applications, column 4 excludes single inventor applications. We find that patents in more complex technological classes are more likely to involve close-distance collaboration.

However, we do not detect a significant change in close-distance collaboration for these patents.

Furthermore, we employ the continuous complexity measure directly from [Broekel \(2019\)](#) and report the results in Table [F7](#). The findings are similar to the ones in Table [F6](#).

We further analyze the relationship with respect to population density after splitting the sample into more or less complex technological classes. The results are presented in Table [F8](#). The table indicates that the density gradient in the probability of being in the top 1% or the top 5% in patent novelty significantly decreases after COVID-19 only for more complex patents but not for less complex patents. The evidence is in line with the idea that the creation of more complex innovation relies more on in-person interactions facilitated by urban density and, hence, is more negatively impacted by COVID-19 than that of less complex innovation.

One possible concern pertains to the disparity in the number of patent applications between the two groups, with 195,365 observations in the low-complexity group and 1,515,926 observations in the high-complexity group. To address this issue, we introduced an alternative dummy variable to denote complexity. This variable identifies patent applications as complex if their complexity measure exceeds the average complexity score calculated for all applications from 2011 to 2019, conducted at the patent application level. This adjustment ensures comparability across observations in the two groups. Results obtained using this new dummy variable are presented in Table [F9](#). Despite differences in sample splitting methods, Table [F9](#) yields similar findings to those in Table [F8](#).

We also use this alternative dummy variable to split the sample based on the average complexity index calculated for all applications from 2011 to 2019 at the patent application level. We firstly examine whether the loss in novelty is concentrated in more complex patent classes in Table [F10](#). Despite differences in sample splitting techniques, both Tables [F4](#) and [F10](#) reveal consistent patterns regarding more novelty loss in more complex patent classes. We then examine the relationship between more complex patent

classes and either the number of inventors or the distance between co-authors in Table F11, which, despite differences in sample splitting methods, shows findings consistent with those in Table F6.

Table F1: The Impact of COVID-19 on Density Premium of the Likelihood of Being the Top 1% and 5% in Novelty

	(1) Top 1%	(2) Top 5%
$\mathbb{1}(\text{after 2020}) \times \mathbb{1}(\text{close collaboration})$	-0.001 (0.005)	-0.009 (0.011)
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.001** (0.001)	-0.005*** (0.002)
$\mathbb{1}(\text{after 2020}) \times \mathbb{1}(\text{close collaboration}) \times \log(\text{pop density})$	-0.000 (0.001)	0.000 (0.001)
Team size	Yes	Yes
CSD FE	Yes	Yes
Year FE	Yes	Yes
Adj. R-squared	0.024	0.045
Obs.	1,131,051	1,131,051

Notes: All regressions follow Equation (D1). The specification includes all interactions among the three variables, but we do not report all of them. The dependent variable is a dummy variable indicating whether a patent belongs to the top 1% or the top 5% of the novelty distribution. Close collaboration applications are patent applications in which at least two inventors are from the same CSD. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table F2: The Impact of COVID-19 on team stability

	(1) Full	(2) Exclude teams formed after 2020	(3) Exclude teams formed after 2020 or Movers
$\mathbb{1}(\text{after 2020}) \times \mathbb{1}(\text{close collaboration})$	-0.016*** (0.001)	-0.016*** (0.001)	-0.016*** (0.001)
Team FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Adj. R-squared	0.19	0.19	0.19
Obs.	4,453,245	4,361,839	4,191,474

Notes: This table presents results using data at the team-year level. The dependent variable is a binary variable, “work together,” which takes a value of one if team k collaborates in year t , and zero otherwise. Column 1 includes all teams available from 2011 to 2021. Column 2 reports results excluding teams formed after 2020. Column 3 further restricts teams that do not move. Close collaboration teams are defined as teams in which at least two inventors reside in the same CSD. We control for team fixed effects and year fixed effects. Standard errors, clustered at the team level, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table F3: Impact of COVID-19 on the Number of New Close-distance Teams

	(1)	(2)
$\log(\text{pop density})$	6.51*** (1.44)	- -
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-1.32*** (0.30)	-1.47*** (0.33)
CSD FE	No	Yes
Year FE	Yes	Yes
Adj. R-squared	0.01	0.97
Obs.	46,041	46,041

Notes: All regressions follow Equation (2), with the dependent variable being the number of new close-distance teams. New close-distance teams are defined as teams in which at least two inventors are from the same CSD, and the team submits their patent application for the first time in a given year. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table F4: The Impact of COVID-19 on the Likelihood of Being the Top 1% and 5% in Novelty (By Complexity)

	(1) 99 th	(2) 99 th	(3) 95 th	(4) 95 th
$\mathbb{1}(\text{Tech Complexity})$	0.007*** (0.000)	0.008*** (0.001)	0.034*** (0.002)	0.037*** (0.002)
$\mathbb{1}(\text{Tech Complexity}) \times \mathbb{1}(\text{after 2020})$	-0.002*** (0.001)	-0.002*** (0.001)	-0.008*** (0.002)	-0.008*** (0.002)
Team size	Yes	No	Yes	No
CSD FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.021	0.016	0.043	0.037
Obs.	1,711,523	1,711,523	1,711,523	1,711,523

Notes: All regressions follow Equation (D1). 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). We classify applications as more complex if their corresponding complexity values rank in the top 50% four-digit CPC classes and as less complex if they rank in the bottom 50% four-digit CPC subclasses. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table F5: The Impact of COVID-19 on the Likelihood of Being the Top 1% and 5% in Novelty (Continuous Measure of Complexity)

	(1) 99 th	(2) 99 th	(3) 95 th	(4) 95 th
Tech Complexity	0.005*** (0.000)	0.005*** (0.000)	0.019*** (0.001)	0.020*** (0.001)
Tech Complexity $\times \mathbb{1}(\text{after 2020})$	-0.001*** (0.000)	-0.001*** (0.000)	-0.004*** (0.001)	-0.003*** (0.001)
Team size	Yes	No	Yes	No
CSD FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.023	0.018	0.050	0.045
Obs.	1,711,523	1,711,523	1,711,523	1,711,523

Notes: All regressions follow Equation (D1). 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). This table uses continuous measure of technological complexity. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table F6: The Impact of COVID-19 on Team Composition (By Complexity)

	(1) ln(TS)	(2) TS > 3	(3) Single	(4) Close
$\mathbb{1}(\text{Tech Complexity})$	0.159*** (0.011)	0.104*** (0.007)	-0.088*** (0.007)	0.023*** (0.005)
$\mathbb{1}(\text{Tech Complexity}) \times \mathbb{1}(\text{after 2020})$	0.000 (0.008)	0.000 (0.006)	0.007 (0.005)	0.004 (0.007)
CSD FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.152	0.113	0.145	0.166
Obs.	1,711,523	1,711,523	1,711,523	1,111,894

Notes: All regressions follow Equation (D1) with different dependent variables. Columns 1 to 3 utilize various measures of inventor team size as outcome variables. Column 1 employs the logarithm of the number of inventors per patent application, column 2 creates a dummy variable indicating whether the number of inventors exceeds 3 (approximately the mean team size), and column 3 is a binary indicator for patents with only one inventor. Column 4 creates a dummy variable indicating close-distance collaboration applications. We classify applications as more complex if their corresponding complexity values rank in the top 50% four-digit CPC classes and as less complex if they rank in the bottom 50% four-digit CPC classes. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table F7: The Impact of COVID-19 on Team Composition (Continuous Measure of Complexity)

	(1) ln(TS)	(2) TS > 3	(3) Single	(4) Close
Tech Complexity	0.051*** (0.003)	0.033*** (0.002)	-0.027*** (0.002)	0.008*** (0.002)
Tech Complexity $\times \mathbb{1}(\text{after 2020})$	0.003 (0.004)	0.001 (0.003)	0.002 (0.002)	0.003 (0.002)
CSD FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.154	0.115	0.146	0.166
Obs.	1,711,523	1,711,523	1,711,523	1,111,894

Notes: All regressions follow Equation (D1) with different dependent variables. Columns 1 to 3 utilize various measures of inventor team size as outcome variables. Column 1 employs the logarithm of the number of inventors per patent application, column 2 creates a dummy variable indicating whether the number of inventors exceeds 3 (approximately the mean team size), and column 3 is a binary indicator for patents with only one inventor. Column 4 creates a dummy variable indicating close-distance collaboration applications. This table uses continuous measure of technological complexity. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table F8: The Impact of COVID-19 on Density Premium of the Likelihood of Being the Top 1% and 5% in Novelty (By Complexity)

	(1) 99 th , Complexity = 0	(2) 99 th , Complexity = 1	(3) 95 th , Complexity = 0	(4) 95 th , Complexity = 1
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.000 (0.000)	-0.001*** (0.000)	-0.001 (0.001)	-0.004*** (0.001)
Team size	Yes	Yes	Yes	Yes
CSD FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	-0.003	0.022	0.010	0.043
Obs.	195,365	1,515,926	195,365	1,515,926

Notes: All regressions follow Equation (D1) with different dependent variables. 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). Columns 1 and 3 contain samples of less complex applications, while columns 2 and 4 contain samples of more complex applications. We classify applications as more complex if their corresponding complexity values rank in the top 50% four-digit CPC classes and as less complex if they rank in the bottom 50% four-digit CPC classes. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table F9: The Impact of COVID-19 on Density Premium of the Likelihood of Being the Top 1% and 5% in Novelty (An Alternative Dummy Variable Indicating Complexity)

	(1) 99 th , Complexity = 0	(2) 99 th , Complexity = 1	(3) 95 th , Complexity = 0	(4) 95 th , Complexity = 1
$\mathbb{1}(\text{after 2020}) \times \log(\text{pop density})$	-0.000** (0.000)	-0.001*** (0.001)	-0.002*** (0.001)	-0.005*** (0.001)
Team size	Yes	Yes	Yes	Yes
CSD FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.009	0.029	0.018	0.055
Obs.	845,055	866,271	845,055	866,271

Notes: All regressions follow Equation (D1) with different dependent variables. 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). Columns 1 and 3 contain samples of less complex applications, while columns 2 and 4 contain samples of more complex applications. We split the sample according to the average complexity index for all applications from 2011 to 2019. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table F10: The Impact of COVID-19 on the Likelihood of Being the Top 1% and 5% in Novelty (An Alternative Dummy Variable Indicating Complexity)

	(1) 99 th	(2) 99 th	(3) 95 th	(4) 95 th
$\mathbb{1}(\text{Tech Complexity})$	0.011*** (0.001)	0.012*** (0.001)	0.045*** (0.004)	0.046*** (0.004)
$\mathbb{1}(\text{Tech Complexity}) \times \mathbb{1}(\text{after 2020})$	-0.003*** (0.001)	-0.003*** (0.001)	-0.007*** (0.002)	-0.007*** (0.002)
Team size	Yes	No	Yes	No
CSD FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.023	0.019	0.050	0.045
Obs.	1,711,523	1,711,523	1,711,523	1,711,523

Notes: All regressions follow Equation (D1). 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). The dummy variable about complexity is defined based on the average complexity index for all applications from 2011 to 2019. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table F11: The Impact of COVID-19 on Team Composition (An Alternative Dummy Variable Indicating Complexity)

	(1) ln(TS)	(2) TS > 3	(3) Single	(4) Close
$\mathbb{1}(\text{Tech Complexity})$	0.072*** (0.007)	0.044*** (0.005)	-0.037*** (0.004)	0.011*** (0.004)
$\mathbb{1}(\text{Tech Complexity}) \times \mathbb{1}(\text{after 2020})$	0.002 (0.007)	0.001 (0.005)	0.001 (0.004)	0.005 (0.004)
CSD FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.149	0.111	0.143	0.166
Obs.	1,711,523	1,711,523	1,711,523	1,111,894

Notes: All regressions follow Equation (D1) with different dependent variables. Columns 1 to 3 utilize various measures of inventor team size as outcome variables. Column 1 employs the logarithm of the number of inventors per patent application, column 2 creates a dummy variable indicating whether the number of inventors exceeds 3 (approximately the mean team size), and column 3 is a binary indicator for patents with only one inventor. The dummy variable about complexity is defined based on the average complexity index for all applications from 2011 to 2019. Standard errors clustered at the CSD level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.