

NBER WORKING PAPER SERIES

INNOVATION SPILLOVERS ACROSS U.S. TECH CLUSTERS

Xavier Giroud
Ernest Liu
Holger Mueller

Working Paper 32677
<http://www.nber.org/papers/w32677>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2024

We thank Nick Bloom, Sabrina Howell, Simone Lenzu, Quinn Maingi, and seminar participants at Princeton, Columbia, NYU, Ohio State, Tulane, LBS Summer Finance Symposium, NYU Spatial Economics Conference, and the University of Chicago Conference on Empirical Finance for helpful comments and suggestions. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 1908. (CBDRB-FY24-P1908-R11479). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Xavier Giroud, Ernest Liu, and Holger Mueller. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Innovation Spillovers across U.S. Tech Clusters
Xavier Giroud, Ernest Liu, and Holger Mueller
NBER Working Paper No. 32677
July 2024
JEL No. G30,O30,R30

ABSTRACT

The vast majority of U.S. inventors work for firms that also have inventors and plants in other tech clusters. Using merged USPTO–U.S. Census Bureau plant-level data, we show that larger tech clusters not only make local inventors more productive but also raise the productivity of inventors and plants in other clusters, which are connected to the focal cluster through their parent firms' networks of innovating plants. Cross-cluster innovation spillovers do not depend on the physical distance between clusters, and plants cite disproportionately more patents from other firms in connected clusters, across large physical distances. To rationalize these findings, and to inform policy, we develop a tractable model of spatial innovation that features both within- and cross-cluster innovation spillovers. Based on our model, we derive a sufficient statistic for the wedge between the social and private returns to innovation in a given location. Taking the model to the data, we rank all U.S. tech clusters according to this wedge. While larger tech clusters exhibit a greater social-private innovation wedge, this is not because of local knowledge spillovers, but because they are well-connected to other clusters through firms' networks of innovating plants. In counterfactual exercises, we show that an increase in the interconnectedness of U.S. tech clusters raises the social-private innovation wedge in (almost) all locations, but especially in tech clusters that are large and well-connected to other clusters.

Xavier Giroud
Columbia Business School
Finance Division
665 W 130th Street, Kravis Hall 771
New York, NY 10027
and NBER
xavier.giroud@gsb.columbia.edu

Holger Mueller
Stern School of Business
New York University
44 West Fourth Street
Suite 9-190
New York, NY 10012-1126
and NBER
hmueller@stern.nyu.edu

Ernest Liu
Department of Economics
Princeton University
20 Washington Road
Princeton, NJ 08540
and NBER
ernestliu@princeton.edu

1 Introduction

Innovative activity is unevenly distributed across space. In the United States, a few large tech clusters, like those in San Francisco, Boston, New York, Los Angeles, San Diego, or Seattle, account for a disproportionate share of innovative output.¹ These large tech clusters not only produce more patents, but they also produce more patents per inventor. Importantly, this is not because more productive inventors select into larger clusters, but rather because larger clusters make the inventors more productive (Moretti, 2021). While determining the exact sources of this productivity gain is arguably difficult, the literature has pointed to Marshallian agglomeration economies, in particular, local knowledge spillovers, given that knowledge is a key input into the innovation process. Examples of knowledge spillovers in the context of innovation range from anecdotal evidence, such as the famous Tom Wolfe quote in Saxenian’s (1994) comparative study of Silicon Valley versus Boston’s Route 128, to a large empirical literature beginning with Jaffe, Trajtenberg, and Henderson (1993) showing that patent citations are geographically localized.² Perhaps the most direct and compelling evidence comes from a fascinating recent study by Atkin, Chen, and Popov (2022), who use smartphone geolocation data to identify face-to-face meetings between workers of different firms in Silicon Valley. The study finds that these face-to-face meetings lead to significantly higher subsequent mutual patent citations, which is consistent with local knowledge spillovers.

A lesser known fact about U.S. patenting activity is that 80.7% of inventors work for firms that also have inventors in other tech clusters.³ If knowledge diffuses locally across inventors of different firms, one would expect that these inventors pass on this knowledge to other productive units within the *same* firm, including inventors in other clusters, thereby generating “cross-cluster innovation spillovers.” Within the firm, knowledge diffusion does not need to rely on face-to-face interactions but can take place through conventional channels, such as emails, internal memos, or video conferencing. Also, as knowledge is non-rival within the firm, it can be used by other productive units at little or

¹See, e.g., Audretsch and Feldman (2004), Carlino and Kerr (2015), and Kerr and Robert-Nicoud (2020).

²“Every year there was some place, the Wagon Wheel, Chez Yvonne, Rickey’s, the Roundhouse, where members of this esoteric fraternity, the young men and women of the semiconductor industry, would head after work to have a drink and gossip and brag and trade war stories about phase jitters, phantom circuits, bubble memories, pulse trains, bounceless contacts, burst modes, leapfrog tests, p-n junctions, sleeping sickness modes, slow-death episodes, RAMs, NAKs, MOSes, PCMs, PROMs, PROM blowers, PROM blasters, and teramagnitudes, meaning multiples of a million millions” (Wolfe, 1983, quoted in Saxenian, 1994, p. 33).

³Based on the universe of United States Patent and Trademark Office (USPTO) patent data.

no additional cost, thus creating significant economies of scale and scope.⁴

This paper provides robust evidence of cross-cluster innovation spillovers. Larger tech clusters not only make local inventors more productive, but they also raise the productivity of inventors in other clusters, who are connected to the focal cluster because they work for a firm that has inventors in both clusters. Crucially, firms only benefit from a larger cluster if they also have inventors in the cluster. If a firm has a plant in a cluster but no inventors, then a larger cluster does not benefit the firm's local plant, nor does it benefit other plants (or inventors) of the firm. On the other hand, if the firm has a plant with inventors in a cluster, then a larger cluster still benefits other (distant) plants of the firm, by raising plant-level productivity, even if the (receiving) plants do not have inventors. Inventors thus effectively act as "antennas," who process local knowledge and pass it on to other productive units within the firm. Receiving plants, on the other hand, do not necessarily need to have inventors, suggesting that some of the shared knowledge is directly beneficial to the firm beyond raising inventors' productivity.

Our interpretation of inventors acting as "antennas" is consistent with Cohen and Levinthal's (1989) notion of "absorptive capacity." In an influential paper, the authors posit that *"firms invest in R&D not only to pursue directly new process and product innovation, but also to [...] exploit externally available information"* (p. 593). This externally available information explicitly includes local knowledge spillovers; in order to benefit from these, firms invest in absorptive capacity:⁵

"[E]conomists have assumed that technological knowledge which is in the public domain is a public good. Like a radio signal or smoke pollution, its effects are thought to be costlessly realised by all firms located within the neighbourhood of the emission. [...] We suggest that if these costs are relatively small, it is by virtue of the considerable R&D already conducted by the firms in the vicinity of the 'emission'; the firm has already invested in the development of its absorptive capacity in the relevant field" (p. 570).

⁴Ding (2023) estimates a quantitative model of joint production that provides evidence in support of the commonly made assumption that knowledge inputs are non-rival within the firm. Markusen's (1984) model of economies in multi-plant firms is an early example of the idea that knowledge is non-rival within the firm: *"Once an innovation is made, it can be incorporated into any number of additional plants without reducing the marginal product of that innovation in existing plants"* (p. 207).

⁵In Cohen and Levinthal's (1989) model, firms invest in R&D so as to absorb external knowledge spillovers. In practice, one can think of a firm's "absorptive capacity" as consisting of inventors and other employees involved in R&D.

To explore the role of inventors and plants in fostering cross-cluster innovation spillovers, we merge USPTO patent data with confidential plant-level data from the U.S. Census Bureau. Building on (and expanding) the empirical research design in Moretti (2021), we identify cross-cluster spillovers through within-plant variation in the size of “connected clusters”—the number of *other* firms’ inventors in the plant’s research field in *other* cities where the plant’s parent firm also has plants with inventors. That is, we examine if, e.g., inventors and plants in Silicon Valley are more productive if other firms have more inventors in Boston’s Route 128, to which the focal (Silicon Valley) inventors and plants are connected through their parent firms’ networks of innovating plants. Threats to identification come from omitted factors that may simultaneously affect inventor- or plant-level productivity and the in- or outflow of other firms’ inventors in connected clusters. For example, given that the firm’s network of innovating plants is non-random, cities in which the firm has innovating plants may be subject to correlated (city-field specific) technology shocks. As a result, plant-level productivity may increase, while at the same time clusters to which the plant is connected may experience an inflow of inventors. Our empirical research design accounts for these and other omitted factors by including granular city \times field \times year fixed effects, besides plant fixed effects. Effectively, we are comparing plants in the same city, research field, and year, which are connected to different clusters, in different cities, because they belong to different parent firms with different geographical networks of innovating plants.

Comparing plants in the same city, research field, and year may not be enough if productivity shocks occur at a more granular level. To account for this possibility, we propose an instrumental variable (IV) research design that isolates plausibly exogenous variation in the size of a plant’s connected clusters based on the expansion of other firms in other cities that do not overlap with the geographical network of the plant’s parent firm. The idea is that (within-plant) changes in the number of inventors in other cities employed by other firms that also have inventors in the plant’s connected clusters are predictive of changes in the number of other firms’ inventors in the plant’s connected clusters but uncorrelated with unobserved shocks that directly affect the plant’s productivity.

In our baseline specification, the elasticity of plant-level total factor productivity (TFP) with respect to the total size of connected clusters is 0.012, and the corresponding elasticity for inventor-level productivity (patents per inventor), aggregated to the plant level, is 0.021. These elasticities are economically meaningful. Suppose we rank plants

based on the size of their connected clusters, and consider a plant at the 25th percentile of the corresponding distribution. If we replaced this plant’s connected clusters with those of a plant at the 75th percentile, plant-level inventor productivity would increase by 8.2% and plant-level TFP would increase by 4.7%. Accordingly, plants at the top of the (connected-cluster size) distribution—i.e., plants which are either connected to many clusters, to large clusters, or both—experience significant productivity advantages due to cross-cluster innovation spillovers. Our IV estimates are slightly noisier but similar to the OLS estimates (0.014 and 0.023, respectively), suggesting that unobserved shocks to plant-level productivity are not a major source of bias in our empirical specification.

Inventors play a key role in fostering cross-cluster innovation spillovers. When we form “placebo connected clusters” based on other cities in which the parent firm has plants but no inventors (“non-innovating plants”), we find no evidence of cross-cluster innovation spillovers. Conversely, we find that non-innovating plants do benefit from spillovers originating from clusters in which the parent firm has innovating plants, albeit the magnitude of the effect is smaller: for non-innovating plants, the elasticity of plant-level TFP with respect to the size of connected clusters is 0.006, which is about half of the elasticity for innovating plants. Hence, some of the information that is disseminated via cross-cluster spillovers is relevant for productivity more broadly, that is, beyond raising inventors’ productivity.

Plants only benefits from exposure to other clusters if the parent firm has inventors in those clusters. In the language of Cohen and Levinthal (1989), inventors provide firms with “absorptive capacity” to collect and process local knowledge, and ultimately to pass on this knowledge to other productive units within the firm. We provide two additional pieces of evidence consistent with knowledge diffusion as a potential mechanism:

1) Knowledge flows within the firm, across different clusters, can take place through conventional channels, such as emails, internal memos, or video conferencing. Therefore, if the mechanism is knowledge diffusion within the firm, we would expect cross-cluster innovation spillovers not to decay with physical distance. We confirm this in our data: all our estimates remain stable and significant if we exclude connected clusters within a 100, 250, or 500 mile radius around the focal plant.

2) As the size of a plant’s connected clusters increases, the plant not only cites more patents (by other firms) from those clusters—which could simply be because it innovates more, and therefore cites more patents from everywhere—but the *share* of citations to

patents from connected clusters increases. Hence, plants disproportionately cite patents from clusters they are connected to through their firm’s innovation network, across large physical distances, which is consistent with knowledge diffusion.

To rationalize our empirical findings, and to inform policy, we develop a tractable model of spatial innovation that features both within- and cross-cluster innovation spillovers. In the model, firms can invest in innovation to improve plants’ productivities. While firms internalize spillovers on other plants of the firm, they do not internalize spillovers on other firms. Cross-cluster spillovers arise for two reasons. First, a plant’s innovation investment spills over to other local plants, and from there to other plants of those firms, via those firms’ plant-level networks. Second, a plant’s innovation investment spills over to other plants of the same firm, in other clusters, and from there to other firms’ plants in those clusters. Within- and cross-cluster spillovers interact in a recursive fashion, generating higher-order spillovers that propagate throughout the economy.

We first characterize the government’s optimal innovation policy.⁶ In the absence of cross-cluster spillovers, the optimal innovation policy involves a uniform subsidy to all innovating plants. With cross-cluster spillovers, however, a uniform subsidy is not optimal. Instead, the optimal subsidy is plant-specific and disproportionately targets plants that are well-connected to other clusters, which are in turn well-connected to other clusters, and so on. To inform policy which locations under-invest the most in innovation, we derive a model-based sufficient statistic—the “social-private innovation wedge”—that captures the wedge between the social and private returns to innovation in a given location. This sufficient statistic can be measured in the data using information on how a given location is connected to other locations through firms’ innovation networks (from our merged USPTO–U.S. Census Bureau micro data) in combination with our reduced-form within- and cross-cluster innovation spillover elasticities.

In the final part of the paper, we take our model to the data. We first rank all U.S. tech clusters based on their social-private innovation wedge. While larger clusters exhibit a greater wedge, this is not because of local knowledge spillovers. In fact, if it was only for local knowledge spillovers, the social-private innovation wedge would be invariant to cluster size.⁷ Instead, what is driving the location ranking is the extent to which clusters

⁶Spillovers are a main justification for innovation policy. As Bloom, Van Reenen, and Williams (2019) note: “*Knowledge spillovers are the central market failure on which economists have focused when justifying government intervention in innovation*” (p. 166).

⁷This invariance echoes a well-known result in urban economics (going back to Glaeser and Gottlieb,

are connected to other clusters through firms’ networks of innovating plants. Hence, the Silicon Valley cluster has a high social-private innovation wedge not because it generates large local knowledge spillovers, but because it is connected to many other clusters, which are in turn connected to many other clusters, etc. We conclude our quantitative analysis with a counterfactual exercise in which we increase the interconnectedness of U.S. tech clusters. While this raises the social-private innovation wedge in (almost) all clusters—exacerbating the under-investment problem—it especially raises the wedge in clusters that are large and well-connected to other clusters.

Our paper contributes to several strands of literature. First, it contributes to the empirical literature on innovation spillovers. In an important contribution, Moretti (2021) develops an empirical research design to identify innovation spillovers based on within-inventor variation in local cluster size. Our research design builds on Moretti’s. However, our focus is not on individual inventors but on plants and their parent firms’ innovation networks. Importantly, our focus is not on within-cluster spillovers but on innovation spillovers across tech clusters. Both Moretti’s and our paper study innovation spillovers originating from other firms’ *inventors*. A related literature studies innovation spillovers originating from other firms’ *R&D* (e.g., Griliches, 1979; Jaffee, 1986; Hall, Mairesse, and Mohnen, 2010; Bloom, Schankerman, and Van Reenen, 2013; Arqu -Castells and Spulber, 2022).⁸ The typical approach in that literature is to regress firm outcomes on some weighted average of other firms’ R&D capital stocks (“spillover pool”). In an important paper, Bloom, Schankerman, and Van Reenen (2013) empirically separate between the technology spillover effect of other firms’ R&D and the product market rivalry effect. Most papers in the R&D spillover literature abstract from geographical considerations. Notable exceptions are Bloom, Schankerman, and Van Reenen (2013, Section 6.2) and Lychagin et al. (2016), who consider geographically distance-weighted R&D spillover pools based on the locations of firm inventors, and Griffith, Harrison, and Van Reenen (2006), who show that UK firms with a higher share of US-based inventors benefit relatively more from US R&D growth.⁹

2008, 2009) that place-based redistributions are not welfare-enhancing unless the elasticity of productivity to agglomeration differs across locations. On the efficiency costs of place-based redistributions, see, e.g., Kline and Moretti (2014), Neumark and Simpson (2015), Gaubert (2018), and Gaubert et al. (2020).

⁸Liu and Ma (2023) embed cross-sector R&D spillovers in an endogenous growth model to study the optimal allocation of R&D resources. Empirically, they show that a sector’s innovation output benefits from past innovations in upstream sectors linked through a global innovation network.

⁹Audretsch and Feldman (1996) find that more R&D-intensive industries exhibit a higher geographic

Second, many empirical studies (including those on R&D spillovers) examine spillovers *between* firms. By contrast, only relatively few studies analyze spillovers, or the propagation of shocks, through *within-firm* networks (e.g., Cravino and Levchenko, 2017; Boehm, Flaaen, and Pandalai-Nayar, 2019; Giroud and Mueller, 2019; Bilir and Morales, 2020; Bena, Dinc, and Erel, 2022; Giroud et al., 2023; Ding, 2023; Biermann and Huber, 2024). Relative to those studies, our paper shows that within-firm networks of innovating plants foster innovation spillovers across U.S. tech clusters. Crucially, plant-level networks alone are not sufficient to generate cross-cluster innovation spillovers; plants also need to have inventors, who absorb and process local knowledge and then pass it on to other inventors and plants within the firm.

Finally, our paper contributes to the literature on tech clusters. Not only do larger tech clusters produce more patents, but they also exhibit higher inventor productivity (Moretti, 2021) and higher growth in patenting activity over time (Bettencourt, Lobo, and Strumsky, 2007; Chattergoon and Kerr, 2022). Many researchers rationalize these productivity advantages with Marshallian agglomeration economies, especially local knowledge spillovers (e.g., Audretsch and Feldman, 2004; Carlino and Kerr, 2015; Kerr and Robert-Nicoud, 2020).¹⁰ Our findings suggest that inventors and plants in larger tech clusters are not only more productive because of local knowledge spillovers, but also because they receive knowledge spillovers from other tech clusters, to which they are connected via firms’ networks of innovating plants. In our data, innovating plants in the ten largest tech clusters are connected to 7.5 times more, and 8.5 times larger, other tech clusters than are innovating plants in the ten smallest tech clusters.¹¹

The rest of this paper is organized as follows. Section 2 describes the data and sample. Section 3 considers within-cluster spillovers. Section 4 considers cross-cluster spillovers; it contains our main results, robustness exercises, results for non-innovating plants, and IV estimates. Section 5 explores potential mechanisms. Section 6 develops the model and derives the social-private innovation wedge. Section 7 ranks all U.S. tech clusters based on this wedge and performs counterfactual exercises. Section 8 concludes.

concentration of both manufacturing and innovative activity, which is consistent with R&D spillovers.

¹⁰Some researchers view knowledge spillovers as an integral part of what defines a tech cluster: “*Our definition of tech clusters emphasizes settings with a frontier edge, and many companies seek insights on emerging possibilities, either through first access to codified knowledge or to tacit knowledge that cannot be written down easily*” (Kerr and Robert-Nicoud, 2020, p. 58).

¹¹The ten largest and smallest tech clusters (by number of inventors) are listed in Table 8.

2 Data

Data Sources

We combine data from the United States Patent and Trademark Office (USPTO) patent database with confidential establishment-level data from the U.S. Census Bureau’s Longitudinal Business Database (LBD), the Census of Manufactures (CMF), and the Annual Survey of Manufactures (ASM). The patent data include the application date, the date when the patent is granted, the names of inventors and assignees (e.g., firms, universities), other patents cited, research field, and the inventors’ home address. The LBD covers all business establishments in the United States with at least one paid employee and contains longitudinal establishment identifiers along with information on employment, payroll, industry, location, and firm affiliation. The CMF is conducted every five years (“Census years”) and covers all manufacturing plants in the United States with at least one paid employee. The ASM is conducted annually in all non-Census years and covers a subset of the plants covered by the CMF. Plants with at least 250 employees are included in every ASM year, whereas plants with fewer employees are randomly (re-)sampled every five years. Although the ASM is technically referred to as a “survey,” reporting is mandatory, and fines are levied for misreporting. The CMF and ASM contain information about key plant-level variables, such as shipments, assets, material inputs, employment, payroll, capital expenditures, industry, and location.

Sample

We merge the USPTO patent data with the LBD for the years 1976 to 2018. “Year” refers to the year of the patent application, not the year in which the patent is granted. The key challenge is to merge firms in both data sets. For the years 2000 to 2018, we use the bridge created by the U.S. Census Bureau (Dreisigmeier et al., 2018); it is the best available bridge to date. For the years 1976 to 1999, we use the bridge by Kerr and Fu (2008). This bridge is only available for firms in the NBER patent database; we thus extend it to all assignees in the USPTO database following the steps in Section 3 of their paper. We assign patents to research fields using 1-digit Cooperative Patent Classification (CPC) codes. Finally, as in Moretti (2021), we assign patents to Bureau of Economic Analysis (BEA) economic areas (“cities”) based on the inventor’s (not the assignee’s) residential address. BEA economic

areas are similar to metropolitan statistical areas (MSAs), albeit some of them are larger. A main advantage of using BEA economic areas, as opposed to MSAs, is that they provide a complete partition of the United States. Figure A.1 in the Online Appendix shows a map of the United States with all 179 BEA economic areas.

We next assign firm inventors to establishments in the LBD. We first aggregate the patent-level data at the inventor-year level. If a patent has multiple inventors, we assign each inventor an equal fraction of the patent. If an inventor has patents in multiple fields in the same year, we use the modal field. We then match inventors to establishments by locating the firm establishment that is nearest to the inventor’s residential address. If an inventor has multiple addresses in the same year, we use the modal address. That said, it is possible that inventors get matched to different establishments in different years if they change their home address or, likewise, their firm affiliation. Finally, we drop matches if the firm has no establishment in the inventor’s city, as it is unclear where the inventor’s place of work might be in that case.

To compute measures of establishment-level productivity, we link establishments in the LBD to the ASM and CMF (using the common identifier LBDNUM). Our sample thus consists of manufacturing plants; establishments which are not in the ASM or CMF are dropped. In some cases, this implies that inventor-establishment matches are dropped even though the firm has a manufacturing plant in the same city, because that manufacturing plant is not the “nearest match.” In robustness exercises, we show that our results are similar if instead of dropping these matches we assign inventors to the firm’s nearest manufacturing plant in the same city. Also in robustness exercises, we provide results for all inventor-establishment matches in the LBD—i.e., without linking establishments in the LBD to the ASM and CMF—with the caveat that in this case we can only compute measures of inventor productivity (aggregated to the establishment level). Finally, we aggregate observations to the plant-year level by counting the number of inventors and patents associated with a plant in a given year. If a plant has patents in multiple fields in the same year, we define the plant’s field based on its modal field, and we only consider patents in that field.¹² In robustness exercises, we show that our results are similar if we consider all fields in which the plant has patents.

Applying these filters yields a sample of 134,000 plant-year observations. We refer to this sample as “within-cluster spillover” sample. Since inventors and patents are always

¹²In our main analysis sample, 86.8% of a plant’s patents are in its (modal) field.

jointly observed in the USPTO data, all plant-year observations include both inventors and patents. In our main analysis sample, we exclude plant-year observations if the firm has no other plant(s) with inventors and patents in the same field in a different city in the given year. This sample has 57,000 plant-year observations; we refer to it as our “cross-cluster spillover” sample. Accordingly, the within-cluster spillover sample includes all manufacturing plants with matched inventors and patents, whereas the cross-cluster spillover sample only includes plants if the firm has at least one other plant with matched inventors and patents in the same field but in a different city in the same year.

Cluster Size

As in Moretti (2021), a cluster represents a combination of city and research field. At the plant level, a cluster comprises all inventors of other firms in the plant’s city and research field, i.e., excluding the plant’s own inventors as well as any inventors working for the plant’s parent firm.¹³ In our empirical analysis, we identify innovation spillovers through within-plant variation in cluster size. In this regard, a concern is that cluster size may be mismeasured, as inventors in the USPTO patent data are only observed in years when patents are filed, even though they may also work in the plant’s cluster in between patent filings. As we do not know with certainty if this is the case, we choose not to interpolate the data in “missing” (i.e., non-patenting) years. In robustness exercises, we show that our results are similar if we employ a broader notion of cluster size where it is assumed that inventors work in a cluster also in between patent filings.

Descriptive Statistics

Table 1 shows descriptive statistics for plants and firms in our main analysis sample with 57,000 plant-year observations. The typical plant has 699 employees and \$275 million in shipments (i.e., sales). Also, it has 4.99 inventors and produces 10.87 patents per year.¹⁴

¹³We form clusters based on all inventors in the USPTO patent data, not just based on the inventors in our analysis sample. The number of inventors in a city depends on demand and supply factors, including business (demand) and personal (supply) taxes (Moretti and Wilson, 2017). This includes feedback effects: as the number of inventors in a city increases, so do place-based costs, like wages, land prices, and rents, which may in turn affect the number of inventors (Gruber, Moretti, and Wilson, 2023). As long as these factors vary at the city \times year (or even city \times field \times year) level, they will be absorbed by our fixed effects.

¹⁴This is the total number of patents per plant and year, not the number of patents adjusted for co-authorship. That number is 3.36, implying 0.67 (“fractional”) patents per inventor and year. By comparison,

Finally, the typical plant is “connected” to 3.22 other clusters—other cities in which the parent firm also has plants (in the same field) with inventors—and these other clusters have in total 1,915 inventors (again, in the plant’s field) working at other firms. There is substantial variation in the number of connected clusters. Plants in the top decile of the distribution are connected to 11.3 other clusters, and these other clusters have 8,563 inventors of other firms, whereas plants in the bottom decile are only connected to one other cluster, and that cluster has 32 inventors of other firms.

Plants’ parent firms have 6,853 employees and \$2,524 million in shipments. They have 17.51 plants, which are spread across 15.23 counties, 11.77 cities, and 9.16 states. Note that these are *all* plants of these firms, not just plants with inventors (“innovating plants”). At the firm level, the ratio of innovating plants to the total number of plants (“inventor share”) is 0.56. This ratio is high for two reasons. First, to be included in our main analysis sample, firms must have at least two innovating plants (in the same field but in different cities). Second, smaller firms with only a few plants have a high inventor share by construction. For example, firms with only two plants have an inventor share of 100%, and those with three plants have an inventor share of either 66.6% or 100%. Since all firms receive equal weight, this pushes the average inventor share upward. Alternatively, we could simply divide the total number of innovating plants across all firms in our sample by the total number of plants. Doing so yields a much smaller inventor share of 0.35.

Innovating plants are geographically dispersed within the firm. For the typical firm, the ratio of counties where the firm has innovating plants to the total number of counties where it has plants is 0.55. The corresponding ratio at the city and state level is 0.52 and 0.49, respectively. Hence, firms have innovating plants in about half of the counties, cities, or states in which they have plants. Again, these are averages across all firms, which are high for the reasons discussed above. If we instead divide the total number of counties, cities, or states in which firms have innovating plants by the total number of counties, cities, or states in which they have plants, we obtain smaller ratios of 0.35, 0.32, and 0.32, respectively. Regardless of which method we use, however, these results strongly suggest that firms do not concentrate their inventors in only one or two regions.

in Moretti’s (2021) sample of “star inventors”—who are at or above the 90th percentile of patent output—the corresponding number is 1.08 patents per inventor and year.

3 Within-Cluster Innovation Spillovers

To provide a benchmark, and to illustrate the empirical research design, we first analyze the effect of local cluster size on local plants' innovation and productivity. A larger local cluster means that there are more inventors (of other firms) in the plant's city and research field, generating more scope for local knowledge spillovers. We estimate:

$$y_{icft} = \gamma_i + \gamma_{ct} + \gamma_{ft} + \gamma_{cf} + \beta \text{Cluster Size}_{icft} + \varepsilon_{icft}, \quad (1)$$

where y_{icft} is either inventor productivity, aggregated to the plant level, or plant-level total factor productivity (TFP), i denotes plants, c denotes cities, f denotes research fields, and t denotes years, $\text{Cluster Size}_{icft}$ is the number of other firms' inventors in the plant's city and research field in a given year (in logs), and γ_i , γ_{ct} , γ_{ft} , and γ_{cf} are plant, city \times year, field \times year, and city \times field fixed effects, respectively. Plant-level inventor productivity is the ratio (in logs) of the number of patents to the number of inventors at the plant level. This ratio is always non-zero in our sample as patents and inventors are jointly observed in the USPTO patent data and we do not interpolate the data in "missing" (i.e., non-patenting) years. TFP is the estimated residual from a plant-level regression of output on capital, labor, and material inputs (all in logs). To allow for different factor intensities across industries and over time, we estimate the regression separately for each 3-digit SIC code industry and year. Accordingly, TFP can be interpreted as the relative productivity of a plant within a given industry and year. Observations are weighted by plant employment. Standard errors are double clustered at the city and year level.

Our empirical strategy of identifying innovation spillovers through within-unit variation in cluster size builds on Moretti (2021), who focuses on individual inventors as the unit of analysis. In our case, the unit of analysis is a plant.¹⁵ Cluster size varies at the city \times field \times year level, allowing us to include granular city \times year, field \times year, and city \times field fixed effects, besides plant fixed effects. The city \times year fixed effects absorb local shocks (across all fields) that may cause an inflow of inventors in the plant's city while at the same time raising plant-level productivity. Similarly, the field \times year fixed effects absorb technology shocks (across all cities) that may lead to an inflow of inventors in the plant's research field while at the same time raising plant-level productivity. Finally, the

¹⁵Moretti (2004) uses a related empirical research design that exploits within-plant variation in local human capital to study the effects on plant-level productivity.

city \times field fixed effects absorb time-invariant heterogeneity that may explain why some research fields are more productive (and attract more inventors) in certain cities.

Table 2 shows the results. In columns (1) and (2), we replace the city \times year, field \times year, and city \times field fixed effects with field and year fixed effects. (City fixed effects are collinear with the plant fixed effects as plants cannot physically change their location.) In columns (3) and (4), we estimate the full specification from equation (1) with all interacted fixed effects. As is shown, the estimates in columns (1) and (2) are similar to those in columns (3) and (4), suggesting that the estimated elasticities are not very sensitive to shocks at either the city or field level. In column (3), the elasticity of inventor productivity at the plant level with respect to cluster size is 0.076, which is close to the corresponding elasticity of 0.0676 at the individual inventor level in Moretti (2021, column (8), Table 3), despite differences in both sample size and composition resulting from different levels of aggregation and different sample selection criteria. In column (4), the elasticity of plant-level TFP with respect to cluster size is 0.023, suggesting that the same economic forces that make plant inventors more productive also increase overall plant-level productivity.¹⁶ Finally, in column (5), we consider the effect of cluster size on plant-level productivity for plants without inventors (“non-innovating plants”). For these plants, “field” is not defined, so we form local clusters based on the *firm’s* main field—the modal field based on all firm patents in a given year.¹⁷ As is shown, for non-innovating plants, local cluster size has no significant effect on plant-level TFP. Hence, a plant only benefits from the local presence of other firms’ inventors if the plant itself has inventors.

4 Cross-Cluster Innovation Spillovers

If inventors pass on local knowledge to other inventors and plants of the firm, including inventors and plants in other clusters, then a larger local cluster should not only affect innovation and productivity at local plants but also at distant plants of the same firm.

¹⁶Since inventor- and plant-level productivity are both outcome variables, this interpretation is consistent with our empirical specification. That said, it is very well possible that plant-level productivity increases *because* the plant innovates more. For instance, in Schumpeterian growth models, innovation raises the quality of input goods; in those models, quality is synonymous with productivity. Also, innovation has been shown to affect firm- and worker-level outcomes that may in turn feed back into firm-level productivity (e.g., Balasubramanian and Sivadasan, 2011; Kline et al., 2019).

¹⁷In our within-cluster (cross-cluster) spillover sample, 83.3% (91.1%) of a firm’s patents are in its main (i.e., modal) field.

To examine whether plants of the same firm, in other clusters, benefit from a larger local cluster, we estimate the following variant of equation (1):

$$y_{icft} = \gamma_i + \gamma_{cft} + \beta \text{Cluster Size}_{icft} + \varepsilon_{icft}, \quad (2)$$

where $\text{Cluster Size}_{icft}$ is the total size of all clusters, as defined in equation (1), in other cities where the plant’s parent firm has innovating plants (in the plant’s field) in a given year (in logs). For expositional simplicity, we refer to those other clusters as “connected clusters.” Crucially, connected clusters do not include cities in which the firm has plants but no inventors. As we have shown previously, non-innovating plants do not benefit from the local presence of other firms’ inventors, so we would not expect them to pass on local knowledge to other inventors and plants of the firm.

Cross-cluster innovation spillovers are identified through within-plant variation in the number of other firms’ inventors in connected clusters. As is shown below, our results are robust if we focus on connected clusters that are more than 500 miles away from the focal plant. Importantly, within a given city, field, and year, there is ample variation in the total size of plants’ connected clusters. Accordingly, we can include granular city \times field \times year fixed effects—denoted by γ_{cft} in equation (2)—allowing us to separate cross-cluster innovation spillovers from productivity shocks that are common to the locations of the firm’s innovating plants. For example, given that the firm’s network of innovating plants is non-random, cities in which the firm has innovating plants may be subject to correlated city-field specific shocks. As a result, the plant’s productivity may increase, while at the same time clusters to which the plant is connected may experience an inflow of (other firms’) inventors. Effectively, we are comparing plants in the same city, research field, and year, which are connected to different clusters, in other cities, because they belong to different parent firms with different spatial networks of innovating plants.

Main Results

Table 3 provides strong and robust evidence of cross-cluster innovation spillovers. As we did in Table 2, we successively add more granular fixed effects, starting with field and year fixed effects (columns (1) and (2)), then including city \times field, city \times year, and field \times year fixed effects (columns (3) and (4)), and finally including city \times field \times year fixed effects (columns (5) and (6)), besides plant fixed effects. The latter is our tightest specification and

will be our baseline specification going forward. As can be seen, the point estimates are stable and significant across different fixed-effect specifications, suggesting that, similar to what we had found in the within-cluster analysis, the estimated elasticities are not very sensitive to shocks at the city or field (or even city \times field) level.

In column (5), the elasticity of plant-level inventor productivity with respect to the total size of connected clusters is 0.021, and in column (6), the corresponding elasticity for plant-level TFP is 0.012. These elasticities are between one third (inventor productivity) and one half (TFP) of the corresponding within-cluster elasticities in columns (3) and (4) of Table 2, suggesting that some but not all of the local information is beneficial to other inventors or plants within the firm. (Alternatively, some information may be “soft” and therefore difficult to pass on.) To interpret the magnitudes, suppose we rank all plants based on their value of $ClusterSize_{icft}$, and consider a plant at the 25th percentile of the corresponding distribution. If we replaced this plant’s connected clusters with those of a plant at the 75th percentile, plant-level inventor productivity would increase by 8.2% and plant-level TFP would increase by 4.7%. Hence, plants at the top of the distribution—i.e., plants which are either connected to many clusters, to large clusters, or both—experience significant productivity advantages due to cross-cluster innovation spillovers.

Non-Innovating Plants

As was previously shown in column (5) of Table 2, non-innovating plants do not benefit from the local presence of other firms’ inventors, so we would not expect these plants to pass on local knowledge to other inventors and plants within the firm. To verify this conjecture, we link (innovating) plants to “placebo clusters” in cities where the parent firm has plants but no inventors. As this placebo test only requires firms to have at least one innovating plant—as opposed to our main analysis sample, which requires firms to have at least two innovating plants in the same field in different cities—the number of observations becomes larger than in our main analysis sample. Table 4 presents the results. As columns (1) and (2) show, a larger cluster in cities where the firm has plants but no inventors—even if the cluster is in the focal plant’s research field—has no significant effect on either plant-level inventor productivity or plant-level TFP. We may therefore conclude that cross-cluster innovation spillovers do not simply originate from any cluster in which the firm has plants but only from clusters in which it has innovating plants.

Do non-innovating plants benefit from spillovers originating from clusters in cities where the parent firm has innovating plants? Theoretically, the answer is unclear. If cross-cluster spillovers only raise plant-level productivity because they spur innovation, then non-innovating plants should not benefit from such spillovers. On the other hand, if cross-cluster spillovers disseminate information that is relevant for productivity more broadly, then even non-innovating plants may benefit. We explore this hypothesis in column (3). As “field” is not defined for plants without inventors, we use again the firm’s main field to form connected clusters, analogous to column (5) of Table 2. As is shown, non-innovating plants benefit from cross-cluster spillovers, albeit less than innovating plants; the elasticity of plant-level TFP with respect to the total size of connected clusters is one half of the corresponding elasticity in column (6) of Table 3. Hence, some of the information that is disseminated through cross-cluster spillovers is relevant for productivity more broadly, that is, beyond raising inventors’ productivity.

Robustness

Table A.1 of the Online Appendix contains robustness exercises. In our main sample, we drop inventor-establishment matches if the establishment is not a manufacturing plant, even if the firm has a nearby manufacturing plant in the same city. In columns (1) and (2), instead of dropping these matches, we assign inventors to the firm’s nearest manufacturing plant in the same city. As is shown, the results are virtually unchanged. Second, in our main sample, the plant’s research field is based on the modal field across all its patents in a given year. This has two implications: i) we only consider patents in that (modal) field, and ii) for consistency, the plant’s connected clusters only consist of other firms’ inventors in that field. In columns (3) and (4), we consider all fields in which the plant has patents, and connected clusters consist of other firms’ inventors in all those fields. As is shown, the results are similar to our baseline results. Third, in the USPTO patent data, inventors are only observed in years when patents are filed, even though they may also work in the cluster in between patent filings. In columns (5) and (6), we employ a broader concept of cluster size where it is assumed that inventors work in a cluster also in between patent filings (that is, provided they do not file patents elsewhere during that time). Again, the results are similar to our baseline results.

Our last robustness exercise, in column (7), concerns industry coverage. Our main

sample is restricted to manufacturing establishments in the ASM and CMF, which allows us to compute plant-level TFP. We are interested in plant-level TFP for two reasons. First, for innovating plants, it provides us with an additional measure of productivity besides inventor productivity. Second, for non-innovating plants, it provides us with *a* measure of productivity, as inventor productivity is not available for those plants. That said, while many innovating firms are in manufacturing industries, others are not.¹⁸ In column (7), we drop the sample requirement that LBD establishments must have a link to the ASM and CMF. Accordingly, we consider all industries, not just manufacturing industries. All other data filters and sample selection criteria remain the same. Naturally, this means we can only examine the effects on inventor productivity (aggregated to the establishment level) but not on TFP. As is shown, the result is similar to our baseline result in column (5) of Table 3, despite a broader industry coverage and a much larger sample size.

Instrumental Variable Estimation

Our empirical research design accounts for granular shocks at the city \times field \times year level that may cause an inflow of other firms' inventors in connected clusters while at the same time raising the focal plant's productivity. However, equation (2) may lead to biased estimates if the shocks occur at a more granular level. To account for this possibility, we propose an instrumental variable (IV) research design that isolates plausibly exogenous variation in the size of a plant's connected clusters based on the expansion of other firms in other cities that do not overlap with the geographical network of the plant's parent firm. The idea is that (within-plant) changes in the number of inventors in other cities employed by other firms that also have inventors in the plant's connected clusters are predictive of changes in the number of other firms' inventors in the plant's connected clusters but uncorrelated with unobserved shocks that may directly affect the plant's productivity, conditional on plant and city \times field \times year fixed effects.¹⁹

¹⁸Manufacturing industries with a high share of U.S. patenting output include, e.g., pharmaceuticals and medicines, medical equipment, aerospace products, automobiles, communications equipment, and semiconductors. Altogether, manufacturing industries account for 66.3% of U.S. patents issued in 2013 (Table 60, Business R&D and Innovation: 2013, Detailed Statistical Tables NSF 16-313, National Science Foundation, National Center for Science and Engineering Statistics, 2016).

¹⁹Moretti (2021) uses a similar IV strategy to instrument for local cluster size. Closely related instruments have also been used in many other literatures, e.g., in the literature on multinationals (MNCs), where local expansion by MNCs is instrumented using the MNCs' expansion in other regions or countries (e.g., Setzler and Tintelnot, 2021; Alfaro-Ureña, Manelici, and Vasquez, 2021). A key difference between our IV strategy

To illustrate, consider a plant in Silicon Valley that is connected to the San Diego and Seattle clusters through its parent firm’s network of innovating plants. Other firms that have inventors in the San Diego and Seattle clusters also have inventors in the Boston, New York, and Los Angeles clusters (where the focal plant’s parent firm has no physical presence). The first-stage regression exploits changes in the number of these firms’ inventors in the Boston, New York, and Los Angeles clusters to predict (within-plant) changes in the number of other firms’ inventors in the San Diego and Seattle clusters; instrument relevance derives from the two being correlated (e.g., because of firm-wide expansions), while instrument exogeneity requires that the growth in the number of other firms’ inventors in the Boston, New York, and Los Angeles clusters is uncorrelated with unobserved productivity shocks to the (focal) Silicon Valley plant.

Table 5 presents the results. As column (1) shows, the number of other firms’ inventors in other cities (in which the plant’s parent firm has no presence) is highly predictive of the size of the plant’s connected clusters. The instrument is both relevant and strong; the corresponding first-stage F-statistic is 70.6. Columns (2) and (3) show 2SLS estimates for plant-level inventor productivity and TFP, respectively. As is shown, the estimates are similar to the corresponding OLS estimates (naturally, the IV estimates are noisier), suggesting that unobserved shocks to plant-level productivity, conditional on the fixed effects, are not a major source of bias in our baseline empirical specification.

Number of Plant Inventors

In our theoretical framework, we assume that the effect of cluster size on plant-level productivity scales up with the plant’s “innovation investment” (e.g., inventors, R&D). To support this assumption empirically, we interact $ClusterSize_{icft}$ with the number of inventors (in logs) at the plant level. The results are shown in Table A.2 of the Online Appendix. As our model features both within- and cross-cluster innovation spillovers, we show empirical results for both. In both cases, we consider the number of inventors at the focal plant (columns (1) to (4)); in the case of cross-cluster spillovers, we also consider the number of inventors at the firm’s (innovating) plants in connected clusters (columns (5) and (6)). There are two main takeaways from Table A.2. First, having more inventors in

and others is that in our cross-location spillover setting the location of the focal unit is not only different from the location of the instrument but also from the location of the instrumented unit, which relaxes the assumption required for instrument exogeneity.

a plant is associated with higher plant-level inventor productivity and higher plant-level TFP, which is consistent with peer effects. (We are obviously careful not to make causal statements here.) Second, and importantly, the interaction term between $ClusterSize_{icft}$ and the number of plant inventors is positive, suggesting that the two are complements, which is consistent with our modeling assumption in Section 6.

Aggregation to the Cluster Level

Our plant-level regression makes comparisons between plants based on their differential exposures to other clusters. In Table A.3 of the Online Appendix, we make comparisons between *clusters* by aggregating plant-year observations to the cluster-year level. Inventor productivity at the cluster level is the ratio (in logs) of the total number of patents to the total number of inventors based on all plants in the cluster. TFP at the cluster level is the employment-weighted average value of TFP across all plants in the cluster. To compute the total size of a cluster’s “connected clusters,” we either add up (columns (1) and (2)) or compute the average value of (columns (3) and (4)) $ClusterSize_{icft}$ across all plants in the cluster.²⁰ As is shown, our plant-level results aggregate well to the cluster level: the elasticities of inventor productivity and TFP at the cluster level are similar to their corresponding elasticities at the plant level.

5 Mechanism

As we have shown, firms only benefit from the local presence of other firms’ inventors if the firm itself has inventors in a cluster. In the language of Cohen and Levinthal (1989), inventors provide firms with “absorptive capacity” to collect and process local knowledge, and ultimately to pass on this knowledge to other inventors and plants within the firm. We now present two additional pieces of evidence consistent with knowledge diffusion, one based on geographical proximity and the other based on backward patent citations.

²⁰In columns (1) and (2), if n plants in a cluster are connected to the same distant cluster, this distant cluster appears n times in the summation, thus receiving greater weight. In columns (3) and (4), averages are simple averages, but the results are virtually identical if we use employment-weighted averages.

Geographical Distance

While local knowledge spillovers, across different firms, rely on face-to-face interactions, knowledge diffusion within the firm, across different geographic clusters, can take place through conventional channels, such as emails, internal memos, or video conferencing. Accordingly, if the driving force behind cross-cluster innovation spillovers is knowledge diffusion within the firm, we would expect the spillovers not to decay with geographical distance. To examine this hypothesis, we exclude all connected clusters within a 100, 250, or 500 mile radius around the location of the focal plant. As some plants have all of their cluster connections within these radii, the sample size gradually shrinks. The results are shown in Table 6. As can be seen, all estimates are stable if we exclude nearby connected clusters, which is consistent with knowledge diffusion.

Patent Citations

Patent citations are widely viewed as evidence of knowledge flows. Hence, if the driving force behind cross-cluster innovation spillovers is knowledge diffusion, we would expect plants to cite more patents from clusters to which they are connected as the size of those clusters increases. That being said, merely looking at patent citation counts is not informative. This is because, as we have previously shown, an increase in the size of a plant's connected clusters raises the plant's inventor productivity and therefore its patent output.²¹ But if the plant produces more patents, this means it will also cite more patents. Even if the plant were to cite patents equally from all clusters, this would mean it also cites more patents from connected clusters, simply because it cites more patents in general. Accordingly, we focus not on patent citation counts but on patent citation *shares*. Specifically, we examine how an increase in the size of a plant's connected clusters affects the share of the plant's citations to patents from those clusters.

In Table 7, the dependent variable is the share of the plant's citations to other firms' patents from connected clusters relative to all its citations to other firms' patents (in the plant's field). In columns (1) to (3), we successively add more granular fixed effects. As is shown, the estimates are positive and stable across all specifications. Hence, plants cite disproportionately more patents from connected clusters, across large physical distances,

²¹As Table A.4 of the Online Appendix shows, the results for patent output mirror those for inventor productivity.

as the size of those clusters increases. In column (4), we drop citations to patents filed in the same year to account for the contemporaneous growth in patents from connected clusters. As is shown, our results are virtually unchanged.²² Finally, in column (5), we consider citations to patents from “placebo clusters” in which the firm has plants but no inventors, similar to what we did in columns (1) and (2) of Table 4. As is shown, an increase in the size of those clusters has no significant effect on their share of the plant’s patent citations.

6 Theoretical Framework

To rationalize our empirical findings, and to inform policy, we develop a theoretical framework in which firms make innovation investments in the presence of both within- and cross-cluster innovation spillovers. A key objective of our model is the derivation of a sufficient statistic of the wedge between the social and private returns to innovation in a given location. In Section 7, we will take our model to the data and use this wedge to show which U.S. tech clusters under-invest the most in innovation, as well as to perform counterfactual exercises in which we strengthen the interconnectedness of tech clusters.

6.1 Setup

There are N locations. Each location has a representative consumer and a continuum of plants producing differentiated varieties that are sold to all locations subject to iceberg trade costs. Each plant belongs to a firm, and each firm owns a set of plants across different locations. Plants can invest in innovation to improve their productivities.²³ We consider a setting that features both within-location innovation spillovers—a plant may receive knowledge spillovers from other plants in the same location—and cross-location (within-firm) innovation spillovers—a plant may receive knowledge spillovers from other

²²Our results are also similar if we drop citations to patents filed in the last two, five, or ten years.

²³We take the plant-level networks of firms as given and focus on the decision to invest in innovation, both along the intensive margin (main analysis) and extensive margin (Online Appendix C.2). In principle, one could envision a richer setting where firms must additionally decide whether and where to open new plants or close down existing ones. For recent modeling advances along these lines, see, e.g., Arkolakis, Eckert, and Shi (2023), who solve the plant location choice problem when there are either positive or negative complementarities between plants, and Oberfield et al. (2024), who solve the limit problem in which the firm chooses a density rather than a discrete set of plants. We have verified that, in our empirical setting, firms’ plant-level networks do not adjust in response to the source of variation considered in this paper. Specifically, as the size of a plant’s local cluster or its connected clusters increases, the parent firm does not adjust its plant-level network structure by either opening new plants or closing down existing ones.

plants of the same firm, in other locations. Within- and cross-location spillovers may interact, generating higher-order spillovers: knowledge may spill over to other plants of the same firm, in other locations, from there it may spill over to other firms' plants in those locations, and so on.

Preferences. Each location n has a representative consumer with preferences:

$$U_n = \ln Y_n - L_n, \quad Y_n \equiv \left[\sum_i \int y_{ni}(\nu)^{\frac{\sigma-1}{\sigma}} d\nu \right]^{\frac{\sigma}{\sigma-1}}, \quad (3)$$

where Y_n is a consumption aggregator, $y_{ni}(\nu)$ is the quantity of variety ν produced in location i that is consumed in location n , L_n is the labor supplied in location n , and $\sigma > 1$ is the elasticity of substitution.

The consumer's budget constraint is:

$$\sum_i \int p_{ni}(\nu) y_{ni}(\nu) d\nu = w_n L_n + \Pi_n - T_n, \quad (4)$$

where $p_{ni}(\nu)$ is the price of variety ν produced in location i that is sold to location n , w_n is the wage rate, which we normalize to one, Π_n are firm profits rebated to the consumer in location n , and T_n is a lump-sum tax.

The solution to the consumer's problem of maximizing (3) subject to (4) yields a demand curve for each variety:

$$p_{ni}(\nu) y_{ni}(\nu) = \frac{p_{ni}(\nu)^{1-\sigma}}{\sum_i \int p_{ni}(\nu)^{1-\sigma} d\nu}, \quad (5)$$

which implies total expenditures by consumer n across all varieties sum up to one:

$$\sum_i \int p_{ni}(\nu) y_{ni}(\nu) d\nu = 1. \quad (6)$$

Remark. We normalize the population of consumers in each location to one for expositional convenience. As is shown in Online Appendix C.1, our analysis readily extends to a setting with heterogeneous location size.

Production and Trade. Plants are monopolistically competitive and feature constant-returns-to-scale production using labor. Given marginal cost $\kappa_i(v)$ in labor units for a plant in location i producing variety v , the price of its product that is sold to location n is:

$$p_{ni}(v) = \frac{\sigma}{\sigma - 1} \tau_{ni} \kappa_i(v), \quad (7)$$

where $\frac{\sigma}{\sigma - 1}$ is the markup and $\tau_{ni} \geq 1$ is a bilateral iceberg trade cost for goods produced in location i that are sold to location n . The plant's variable profit from selling goods to location n is:

$$\pi_{ni}(v) = \frac{1}{\sigma} p_{ni}(v) y_{ni}(v). \quad (8)$$

Substituting (5) and (7) in (8) and summing across all locations n which the plant sells to, the plant's total variable profit is:

$$\pi_i(v) = b_i \kappa_i(v)^{1-\sigma},$$

where b_i is a (location-specific) “market access” demand shifter for plants in location i :

$$b_i \equiv \sum_n \frac{1}{\sigma} \frac{\tau_{ni}^{1-\sigma}}{\sum_m \int_v (\kappa_m(v) \tau_{nm})^{1-\sigma} dv}. \quad (9)$$

Location i 's market access is decreasing in the trade cost of selling goods to other locations, τ_{ni} , and increasing in the production and trade costs of producers from other locations.

Firm Networks, Innovation Investments, and Knowledge Spillovers. Each plant belongs to a firm, and each firm owns a set of plants across different locations. We refer to a plant's productivity as its *knowledge*, which is inversely related to marginal cost. A plant can invest in innovation to improve its knowledge and thus reduce its marginal cost. Motivated by our reduced-form evidence, we consider two types of knowledge spillovers. First, a plant benefits from the knowledge of other local plants. Second, a plant benefits from the knowledge of other plants of the same firm, in other locations.

In our baseline model, we assume no heterogeneity across plants besides their locations and the firms they belong to. We can therefore use the set of plant locations J to denote firms with plants in locations $n \in J$. Likewise, we use nJ to denote the plant in location $n \in J$ that belongs to firm J . Also, we abstract from the extensive margin of

innovation and focus on the intensive margin: we assume all plants innovate to some degree in equilibrium. In Online Appendix C.2, we extend our setting by adding heterogeneous productivities and endogenous innovation decisions along the extensive margin, so that only a subset of the plants innovates in equilibrium, and we derive our main sufficient statistic—the social-private innovation wedge—in that extended setting.

Let k_{nJ} denote the knowledge of plant nJ . We specify:

$$k_{nJ} \equiv \prod_{i \in J} \left(h_{iJ}^\alpha K_i^\delta \right)^{\omega_{ni}}, \quad (10)$$

where h_{iJ} denotes plant iJ 's innovation investment and K_i is an aggregator of the knowledge of all plants in location i :

$$K_i \equiv \sum_{J \ni i} N_J k_{iJ}, \quad (11)$$

where N_J is the measure of firms with plants in location set J .

Equations (10) and (11) have five key implications. First, plants directly benefit from their own innovation investment as well as from the innovation investments of other plants of the same firm. Second, plants receive local knowledge spillovers, which are increasing in their own innovation investment. Third, plants receive knowledge spillovers from other locations, which are increasing in their parent firm's innovation investments in those (other) locations. The second and third implications are consistent with the evidence in Table A.2 of the Online Appendix. Fourth, the knowledge spillovers which plants receive, both locally and from other locations, are increasing in the innovation investments of other firms. Finally, given the recursive formulation in (10) and (11), there is, in principle, an infinite loop of knowledge spillovers, where knowledge spills over to a firm's local plant, from there it spills over to other plants of the same firm, in other locations, from there it spills over to other firms' plants in those locations, and so on.

A key object in our model is the strength of knowledge spillovers across locations governed by the elasticities ω_{ni} . Motivated by our reduced-form evidence that spillovers across clusters are weaker than within-cluster spillovers but do not decay with distance,

we specify:

$$\omega_{ni} = \begin{cases} 1 & \text{if } n = i, \\ \omega < 1 & \text{otherwise.} \end{cases}$$

A plant's marginal cost is inversely related to its knowledge: $\kappa_{nJ}^{1-\sigma} = k_{nJ}$. We can thus write the variable profit of plant nJ as:

$$\pi_{nJ} = b_n \prod_{i \in J} \left(h_{iJ}^\alpha K_i^\delta \right)^{\omega_{ni}}, \quad (12)$$

where b_n is the market access in location n , as defined in (9), which can be re-written as:

$$b_i \equiv \sum_n \frac{1}{\sigma} \frac{\tau_{ni}^{1-\sigma}}{\sum_m \sum_{J \ni m} N_J k_{mJ} \tau_{nm}^{1-\sigma}}. \quad (13)$$

Given the demand elasticity σ (c.f. equation 5), plant-level revenues are $r_{nJ} = \sigma \pi_{nJ}$.

The firm internalizes the impact of investing in innovation at any of its plants on the knowledge of any of its other plants and makes investment decisions accordingly to maximize the total profits of all of its plants. However, the firm does not internalize the impact of its innovation investment on location-wide knowledge K_n and thus on the productivities of other firms' plants.

Formally, firm J 's problem is to choose innovation investments $\{h_{nJ}\}$ across all of its plants to maximize the total profits of these plants, net of the cost of the innovation investments:

$$\hat{\pi}_J \equiv \max_{\{h_{nJ}\}} \sum_n \left[b_n \prod_{i \in J} \left(h_{iJ}^\alpha K_i^\delta \right)^{\omega_{ni}} - \frac{c_n h_{nJ}}{t_{nJ}} \right]. \quad (14)$$

The first term inside the bracket captures plant nJ 's variable profits from selling its products, c_n is the marginal cost of innovation investments in labor units, and t_{nJ} is a proportional subsidy to plant nJ 's innovation which the firm takes as given.

We assume all profits are rebated back to consumers, satisfying:

$$\sum_n \Pi_n = \sum_J N_J \hat{\pi}_J. \quad (15)$$

Government Budget. Innovation subsidies are provided by the government, which must balance its budget by levying lump-sum taxes:

$$\sum_n \left\{ \sum_{J \ni n} N_J \left(\frac{c_n h_{nJ}}{t_{nJ}} - c_n h_{nJ} \right) + T_n \right\} = 0. \quad (16)$$

Market Clearing Conditions. Each plant's output must equal the sum of quantities delivered to consumers in all locations. Because plants produce linearly from labor, the labor market clearing condition is:

$$\sum_{iJ} k_{iJ}^{\frac{1}{1-\sigma}} \sum_n y_{niJ} t_{ni} + \sum_{nJ} c_n h_{nJ} = L, \quad (17)$$

where (with a slight abuse of notation) we use y_{niJ} to denote the consumption in location n of the variety produced by plant iJ .

Definition 1. Given innovation subsidies $\{t_{nJ}\}$, an equilibrium is the collection of innovation investments $\{h_{iJ}\}$ that solve the firm's problem (14), where a plant's knowledge k_{nJ} and location-wide knowledge K_i are recursively defined by (10) and (11), the market access b_n satisfies (13), the price of plant iJ 's output that is sold in location n is $\frac{\sigma}{\sigma-1} \tau_{ni} k_{iJ}^{\frac{1}{1-\sigma}}$, the consumer in each location chooses consumption and labor supply to maximize (3) subject to (4), lump-sum taxes are levied to balance the government's budget according to (16), profits are rebated to consumers according to (15), and the labor market clears according to (17).

Definition 2. An equilibrium without government subsidies (i.e., $t_{nJ} \equiv 1 \forall nJ$) is referred to as *Laissez-faire* equilibrium.

Lemma 1. Let $\mathcal{J} \equiv \max |J|$ denote the maximum number of plants belonging to a firm. A sufficient condition for equilibrium uniqueness is $(\alpha + \delta) (1 + (\mathcal{J} - 1) \omega) < 1$.

Proof. See Online Appendix B.1. □

Social Welfare. The utilitarian social welfare is the sum of consumer utilities across locations:

$$\mathcal{W} \equiv \sum_n U_n.$$

Substituting (6) in (4), summing across locations, and substituting out Π_n and T_n using (15) and (16), we can write social welfare as a function of the innovation investments $\{h_{nJ}\}$:

$$\mathcal{W} \equiv \sum_n \left(\ln Y_n - N_J \sum_{J \ni n} c_{nJ} h_{nJ} \right) - \frac{\sigma - 1}{\sigma} N. \quad (18)$$

Equation (18) implies that social welfare is equal to the value of the aggregate consumption bundle net of innovation costs minus a constant. To facilitate the interpretation of our welfare analysis, it is useful to define an alternative social welfare function $\hat{\mathcal{W}}$ that only depends on firms' variable profits and innovation costs:

$$\hat{\mathcal{W}} \equiv \sum_n \left(\frac{\sigma}{\sigma - 1} \sum_{J \ni n} N_J b_n \prod_{i \in J} \left(h_{iJ}^\alpha K_i^\delta \right)^{\omega_{ni}} - \sum_{J \ni n} N_J c_{nJ} h_{nJ} \right). \quad (19)$$

Using the definition of variable profits in equation (12), we can write $\hat{\mathcal{W}}$ conveniently as:

$$\hat{\mathcal{W}} \equiv \sum_n \sum_{J \ni n} N_J \left(\frac{\sigma}{\sigma - 1} \pi_{nJ} - c_{nJ} h_{nJ} \right).$$

Lemma 2. $\frac{d\mathcal{W}}{dh_{nJ}} = \frac{d\hat{\mathcal{W}}}{dh_{nJ}}|_{\text{holding } \{b_n\} \text{ constant}} \text{ for all } nJ$.

Proof. See Online Appendix B.2. □

Lemma 2 implies that, in equilibrium, the marginal impact of innovation investments on social welfare is proportional to the impact on the sum of variable profits ($\frac{d \sum_n \ln Y_n}{dh_{kJ'}} = \frac{\sigma}{\sigma - 1} \frac{d \sum_n \sum_J N_J \pi_{nJ}}{dh_{kJ'}}$) taking the demand shifters $\{b_n\}$ in each location as given. This enables us to examine the impact of plant-level innovation on social welfare by simply analyzing the impact on firms' profits. Juxtaposing the welfare function $\hat{\mathcal{W}}$ with the objective function of the firm's problem (14), we are now ready to discuss the sources of inefficiencies that arise in a Laissez-faire equilibrium.

6.2 Equilibrium

Sources of Inefficiencies. The Laissez-faire equilibrium (without subsidies, i.e., $t_{nJ} \equiv 1$) features three sources of inefficiencies. First, firms only internalize the impact of their innovation investments on profits and ignore the impact on consumer surplus. This

inefficiency is captured by the coefficient $\frac{\sigma}{\sigma-1}$ in the social welfare function (19) and, by Lemma 2, can be corrected through a uniform innovation subsidy across all plants.

Second, firms only internalize the impact of their innovation investments on their own profits and ignore the impact on other firms' profits. This second inefficiency manifests itself in two ways. First, a plant's innovation investment generates local knowledge spillovers on other plants in the same location. Second, a plant's innovation investment affects plants in other locations, because i) the local knowledge spillover diffuses across space through the plant-level networks of *other* firms, and ii) a plant's innovation investment affects the knowledge of other plants of the *same* firm, in other locations, from where it spills over to other firms' plants in those locations. This second inefficiency can be corrected through innovation policies that subsidize plants or firms heterogeneously depending on their location and "connectedness" to other locations.

Finally, due to markups, firms under-produce given their marginal costs. This inefficiency is standard in monopolistic settings and can be corrected through a uniform production subsidy across all plants. In our analysis, this inefficiency plays no role. For expositional simplicity, our equilibrium definition takes markups as given, and we make no attempt to correct for them. Instead, we focus on innovation subsidies and their impact on social welfare.

Social and Private Value of Innovation. To firm J , the private value of investing in innovation in plant iJ , defined as the semi-elasticity of firm-wide profits with respect to plant-level innovation h_{iJ} , is:

$$\beta_{iJ} \equiv \frac{\partial \sum_{n \in J} \pi_{nJ}}{\partial \ln h_{iJ}} = \alpha \sum_{n \in J} \pi_{nJ} \omega_{ni} \mathbf{1}_{ni}^J, \quad (20)$$

where $\{n \in J\}$ is the set of locations in which firm J has plants, and $\mathbf{1}_{ni}^J$ is an indicator for firm J having plants in both locations n and i . In equilibrium, by the first-order condition of the firm's problem (14), the private value of innovation must equal the firm's total private expenditure on innovation:

$$\beta_{iJ} = \frac{c_i h_{iJ}}{t_{iJ}} \quad \text{in equilibrium.} \quad (21)$$

The social value of investing in innovation in plant iJ , defined as the semi-elasticity

of social welfare with respect to plant-level innovation, is:

$$\begin{aligned}
\gamma_{iJ} &\equiv N_J^{-1} \frac{d\hat{W}}{d \ln h_{iJ}} \Big|_{\{b_n\}} \\
&= \frac{\sigma}{\sigma-1} N_J^{-1} \left(N_J \frac{\partial \sum_{n \in J} \pi_{nJ}}{\partial \ln h_{iJ}} + \sum_m \frac{\partial \sum_{J'} N_{J'} \sum_{n \in J'} \pi_{nJ'}}{\partial \ln K_m} \frac{d \ln K_m}{d \ln h_{iJ}} \right) \\
&= \frac{\sigma}{\sigma-1} N_J^{-1} \left(\alpha \sum_{n \in J} N_J \pi_{nJ} \omega_{ni} \mathbf{1}_{ni}^J + \delta \sum_{J'} N_{J'} \sum_{n \in J'} \sum_m \pi_{nJ'} \omega_{nm} \mathbf{1}_{nm}^{J'} \frac{d \ln K_m}{d \ln h_{iJ}} \right). \quad (22)
\end{aligned}$$

The first term inside the parenthesis is proportional to the private value of innovation; it captures the impact of plant-level innovation h_{iJ} on the profit of firm J (multiplied by $\frac{\sigma}{\sigma-1}$ to account for the impact on consumer surplus). The second term inside the parenthesis is new. It is proportional to the total indirect effect of plant-level innovation on the profits of all firms through its impact on location-wide knowledge $\{K_m\}$. A social planner seeking to maximize social welfare will choose innovation investments such that the social value coincides with the total social cost of innovation:

$$\gamma_{iJ} = h_{iJ} c_i. \quad (23)$$

The Impact of Plant-Level Innovation on Location-Wide Knowledge. We now derive the impact of plant-level innovation h_{iJ} on location-wide knowledge K_n , both in the plant's own location and other locations. To this end, we substitute (10) in (11) and re-write location-wide knowledge K_n as:

$$K_n \equiv \sum_{J \ni n} N_J k_{nJ} = \sum_{J \ni n} N_J \prod_{i \in J} \left(h_{iJ}^\alpha K_i^\delta \right)^{\omega_{ni}}.$$

Totally differentiating and substituting using $r_{nJ} = \sigma b_n \prod_{i \in J} \left(h_{iJ}^\alpha K_i^\delta \right)^{\omega_{ni}}$ we obtain:

$$\frac{d \ln K_n}{d \ln h_{iJ}} = S_{nJ} \alpha \omega_{ni} \mathbf{1}_{ni}^J + \sum_{J' \ni n} S_{nJ'} \sum_{q \in J'} \delta \omega_{nq} \frac{d \ln K_q}{d \ln h_{iJ}}, \quad (24)$$

where $S_{nJ} \equiv \frac{N_J r_{nJ}}{\sum_{J' \ni n} N_{J'} r_{nJ'}}$ denotes plant nJ 's sales as a share of total sales by all plants in location n .

Equation (24) decomposes the effect of plant-level innovation h_{ij} on location-wide knowledge K_n into direct and indirect network effects. A plant's innovation investment improves both its own knowledge and the knowledge of other plants of the same firm, which contributes to location-wide knowledge in the plant's own location and the other plants' locations, including location n . This direct spillover effect is captured by the first term on the right-hand side. An increase in location-wide knowledge in any of these locations in turn improves the knowledge of any plants in these locations, which in turn improves the location-wide knowledge in any of the locations these plants are connected to through their firms' plant-level networks, including location n . The second term on the right-hand side captures this higher-order (indirect) spillover effect. We could further decompose and expand $\frac{d \ln K_q}{d \ln h_{ij}}$ into direct and indirect spillover effects. Instead of doing so, we provide a succinct summary of the infinite rounds of network spillover effects in matrix notation below. For now, we merely note that plant iJ 's innovation investment has a larger impact on location-wide knowledge K_n if firm J 's plant in location n accounts for a large share of local sales, or if plant iJ has a large impact on location-wide knowledge in other locations q with many plants there being connected to large plants in location n .

Optimal Innovation Subsidy. Comparing the equilibrium and social optimum first-order conditions, (21) and (23), and recognizing that plant-level revenues are proportional to profits ($r_{nJ} = \sigma \pi_{nJ}$), we can now characterize the optimal subsidy.

Proposition 1. *The optimal innovation subsidy $\left\{ \arg \max_{\{t_{ij}\}} \hat{\mathcal{W}} \right\}$ is plant-specific and satisfies*

$$t_{ij} = \frac{\sigma}{\sigma - 1} \frac{\left[\alpha N_J \sum_{n \in J} r_{nJ} \omega_{ni} \mathbf{1}_{ni}^J + \delta \sum_{J'} N_{J'} \sum_{n \in J'} \sum_m r_{nJ'} \omega_{nm} \mathbf{1}_{nm}^{J'} \frac{d \ln K_m}{d \ln h_{ij}} \right]}{\alpha N_J \sum_{n \in J} r_{nJ} \omega_{ni} \mathbf{1}_{ni}^J}. \quad (25)$$

Proposition 1 has two key implications. First, the optimal subsidy is generally not uniform, neither across all plants nor across all plants in a location. Rather, it is plant-specific. Two plants in the same location may receive different optimal subsidies because they are differentially connected to other locations through their firms' plant-level networks. Ceteris paribus, a plant should receive a higher subsidy if its innovation leads to a larger increase in location-wide knowledge ($\frac{d \ln K_n}{d \ln h_{ij}}$ as characterized in (24)) in locations that are disproportionately connected to large (i.e., high-revenue) plants through firms'

plant-level networks.

There are special cases when the optimal subsidy is uniform across all plants: when either cross-location spillovers are zero ($\omega_{nm} = 0$ for all $n \neq m$) or when all firms are single-location firms ($\mathbf{1}_{nm}^J = 0$ for all $n \neq m$). In those special cases, the optimal subsidy (25) simplifies intuitively to:

$$t_{iJ} = \frac{\sigma}{\sigma - 1} \frac{1}{1 - \delta} \text{ for all } iJ, \quad (26)$$

where $\frac{\sigma}{\sigma - 1}$ corrects for the fact that plants under-invest in innovation because their profits do not capture the full social surplus, and $\frac{1}{1 - \delta}$ corrects for the fact that plants under-invest because they ignore local knowledge spillovers on other plants through their effect on K_i . On the other hand, in the general case when firms have plants in multiple locations and cross-location spillovers are non-zero, as we find in our data, the optimal subsidy will not be uniform but plant-specific.

Second, Proposition 1 implies that the optimal plant-specific subsidy can be expressed using only a few elasticities ($\sigma, \alpha, \delta, \omega$), the network of plants ($\mathbf{1}_{ni}^J$), and plant-level revenues as sufficient statistics. In particular, model primitives such as the costs of innovation investments and bilateral trade costs, or endogenous equilibrium objects such as prices, marginal costs, and location-wide knowledge, do not appear in equation (25).²⁴

6.3 Social versus Private Value of Innovation

A much studied object in the literature is the wedge between the social and private value of innovation. Our theoretical framework allows us to characterize this wedge at the location level accounting for all sources of inefficiencies in our model. We begin by defining the private and social value of innovation in a given location.

Definition 3. The *private value of innovation* in location i , β_i , is defined as the sum of plant-level private values, $\beta_i \equiv \sum_{J \ni i} \beta_{iJ} N_J$. The *social value of innovation* in location i , γ_i , is defined as the sum of plant-level social values, $\gamma_i \equiv \sum_{J \ni i} \gamma_{iJ} N_J$.

Definition 3 aggregates plant-level private and social values to the location level. At

²⁴If one can jointly measure the sufficient statistics—the right-hand side of equation (25)—and the plant-level subsidies in the data, one can test whether the subsidies in the data are optimal. However, as plant-level revenues depend on the subsidies, one cannot treat observed revenues as exogenous data inputs to prescribe optimal subsidies.

the plant level, γ_{iJ}/β_{iJ} captures the wedge between social and private returns to innovation. At the location level, γ_i/β_i captures the average wedge between social and private returns across all plants in a location, weighted by plant-level innovation expenditures:

$$\frac{\gamma_i}{\beta_i} = \sum_{J \ni i} \frac{\gamma_{iJ}}{\beta_{iJ}} \frac{N_J h_{iJ}}{\sum_{J' \ni i} N_{J'} h_{iJ'}}.$$

We refer to γ_i/β_i as the “*social-private innovation wedge*.” A given location has a higher wedge if private incentives lead to larger under-investment in innovation relative to the social return.

We use boldface notation to denote vectors and matrices. Let \mathbf{I} denote the identity matrix, and let $\mathbf{\Omega} \equiv [\omega_{ni}]$ denote the matrix with entries ω_{ni} . Let $s_{ni} \equiv \frac{\sum_J N_J r_{nJ} \mathbf{1}_{ni}^J}{\sum_J N_J r_{nJ}}$ denote the revenue-weighted fraction of plants in location n that are connected to location i through plant-level networks, and let $\mathbf{S} \equiv [s_{ni}]$ denote the corresponding matrix form. Let $\mathbf{r} \equiv [\sum_J N_J r_{iJ}]$ denote the column vector of total plant-level revenues in each location, and let $\boldsymbol{\gamma} \equiv [\gamma_i]$ and $\boldsymbol{\beta} \equiv [\beta_i]$. Our next proposition characterizes $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ using sufficient statistics that can be readily measured in the data.

Proposition 2. (i) $\boldsymbol{\beta}' = \frac{\alpha}{\sigma} \mathbf{r}' (\mathbf{\Omega} \circ \mathbf{S})$; (ii) $\boldsymbol{\gamma}' = \frac{\sigma}{\sigma-1} \boldsymbol{\beta}' (\mathbf{I} - \delta \mathbf{\Omega} \circ \mathbf{S})^{-1}$, where \circ denotes the Hadamard product.

Proof. See Online Appendix B.3. □

The first part of Proposition 2 characterizes the private value of innovation in a given location. Consider first again the special case when either cross-location spillovers are zero (so that $\mathbf{\Omega}$ is the identity matrix \mathbf{I}) or all firms are single-location firms (so that \mathbf{S} is the identity matrix). In either case, the (Hadamard) product $\mathbf{\Omega} \circ \mathbf{S}$ is equal to the identity matrix. Proposition 2 then implies that the private value of innovation in a given location is $\beta_i = \frac{\alpha}{\sigma} r_i$. By contrast, when firms have plants in multiple locations ($\mathbf{S} \neq \mathbf{I}$) and cross-location spillovers are non-zero ($\mathbf{\Omega} \neq \mathbf{I}$), the first part of Proposition 2 implies that $\beta_n = \alpha \sum_i \pi_i \omega_{ni} s_{ni}$. Thus, firms have stronger private incentives to invest in innovation in a given location if the investments generate larger spillovers (ω_{ni}) on other locations to which they have strong connections (s_{ni}) and which are highly profitable (π_i).

As for the second part of Proposition 2, we can rewrite γ as:

$$\begin{aligned}\gamma' &= \frac{\sigma}{\sigma-1} \beta' (I - \delta \Omega \circ S)^{-1} \\ &= \frac{\sigma}{\sigma-1} \beta' (I + \delta \Omega \circ S + (\delta \Omega \circ S)^2 + (\delta \Omega \circ S)^3 + \dots).\end{aligned}\quad (27)$$

The scalar $\frac{\sigma}{\sigma-1}$ captures that—as in the definition of \hat{W} —the impact of innovation investments on social welfare is $\frac{\sigma}{\sigma-1}$ times firm profits. The Leontief-inverse $(I - \delta \Omega \circ S)^{-1}$ captures the network effects of knowledge spillovers that propagate across locations through the plant-level networks of multi-location firms. To understand this network propagation effect, consider again first the special case when either cross-location spillovers are zero or all firms are single-location firms. Equation (27) then implies that the social-private innovation wedge is constant across all locations and equal to:

$$\frac{\gamma_i}{\beta_i} = \frac{\sigma}{\sigma-1} \frac{1}{1-\delta}, \quad (28)$$

which coincides with the optimal innovation subsidy in (26).

In the general case when firms have plants in multiple locations and cross-location spillovers are non-zero, firms *additionally* under-invest in innovation because they also ignore the impacts of their innovation investments on other firms' plants in other locations, via cross-location knowledge spillovers. These spatial spillover effects are captured by the cross-location elasticity matrix Ω and are more socially valuable when a greater (revenue-weighted) fraction of plants in a given location are connected to other locations with stronger private incentives, as reflected by those locations' values of β_i .²⁵

It is worth emphasizing that both the vector \mathbf{r} , which measures the total revenue across plants in each location, and the matrix \mathbf{S} , which provides the revenue-weighted fraction of plants in each location that are connected to each other location, can be measured in the data. Hence, by calibrating only two parameters, δ and ω , we can recover the relative rankings of individual locations based on their social-private innovation wedge (the scalars σ and α are not needed for such a ranking).

²⁵In Online Appendix C.2, we derive the analogue of Proposition 2 in an extended version of our model with heterogeneous productivities and endogenous innovation decisions along the extensive margin, so that ultimately only a subset of the plants innovates in equilibrium.

6.4 A Three-Location Example

Consider a setting with three locations $\{a, b, c\}$ and six firms. Each location has three plants. The firm ownership of the plants and their distribution across locations are as follows:

		Firm					
		1	2	3	4	5	6
Location	a	✓	✓	✓			
	b	✓	✓		✓		
	c			✓		✓	✓

For simplicity, assume all plants have the same revenue. Given the distribution and firm ownership of the plants, the matrix S , which shows the revenue-weighted fraction of plants in each location that are connected to each other location, is:

$$S = \begin{bmatrix} 1 & 2/3 & 1/3 \\ 2/3 & 1 & 0 \\ 1/3 & 0 & 1 \end{bmatrix}.$$

By construction, all three locations have the same number of plants and the same revenue. The locations differ only in their connectedness through the firms' networks of plants. Specifically, locations a and b are connected through plants of firms 1 and 2; locations a and c are connected through plants of firm 3; locations b and c are not directly connected.

We can finally rank the three locations based on their social-private innovation wedges as follows:

$$\frac{\gamma_a}{\beta_a} > \frac{\gamma_b}{\beta_b} > \frac{\gamma_c}{\beta_c}.$$

The social-private innovation wedge is largest in location a and smallest in location c . Location a has the highest wedge because innovation there creates spillovers to both locations b and c . Location b has a higher wedge than location c because the spillovers from location b to location a occur through two firms, 1 and 2, whereas the spillovers from location c to location a only occur through a single firm, 3.

7 The Social-Private Innovation Wedge

7.1 Taking the Model to the Data

Our theoretical framework provides a characterization of the wedge between the social and private returns to innovation at the location level (“social-private innovation wedge”). This wedge can be described in terms of a simple sufficient statistic, γ_n/β_n , that can be measured in the data. In Section 7.2, we rank all U.S. tech clusters based on this wedge, and in Section 7.3, we examine how an increase in the interconnectedness of U.S. tech clusters affects the social-private innovation wedge in each location. In this section, we describe how we compute this wedge using our merged USPTO–U.S. Census Bureau micro data in combination with the reduced-form elasticities from Sections 3 and 4.

There are 179 locations (“cities”). Thus, in Proposition 2, $\gamma' = \frac{\sigma}{\sigma-1}\beta'(\mathbf{I} - \delta\Omega \circ \mathbf{S})^{-1}$ and $\beta' = \frac{\alpha}{\sigma}\mathbf{r}'(\Omega \circ \mathbf{S})$ are 179×1 vectors with entries γ_n and β_n , respectively. We compute γ and β separately for each year and take average values across all years. (The scalars σ and α are the same for all locations and drop out when computing the ratio between any two wedges.) In equation (10), δ captures the strength of local knowledge spillovers (recall that $\omega_{ni} = 1$ if $n = i$). For this, we use the estimated within-cluster elasticity from column (3) of Table 2. Similarly, $\delta\omega$ captures the strength of cross-location knowledge spillovers. For this, we use the estimated cross-cluster elasticity from column (5) of Table 3. Both elasticities pertain to plant-level inventor productivity; we obtain similar results if we instead use the corresponding plant-level TFP elasticities. To obtain ω (to populate the matrix Ω), we divide $\delta\omega$ by δ . In reduced form, this means we divide the cross-cluster elasticity by the within-cluster elasticity. The matrix Ω is a 179×179 matrix with diagonals of 1 and off-diagonals of ω (see Section 6.1); the latter are scaled by the average number of connected clusters, which is 3.22 (see Table 1). If a location n has no connected clusters in location i , we set the corresponding $\{n, i\}$ entry in the Ω matrix to zero.

The matrix \mathbf{S} is a 179×179 matrix with entries $s_{ni} \equiv \frac{\sum_J N_J r_{nJ} \mathbf{1}_{ni}^J}{\sum_J N_J r_{nJ}}$ representing the revenue-weighted fraction of innovating plants in location n that have a connected cluster in location i . To compute this matrix, we first construct an auxiliary location \times firm matrix similar to the one in the three-location example in Section 6.4, where each entry r_{nJ} represents the total value of plant-level revenues of firm J in location n . We obtain plant-level revenues directly from the Census Bureau data (“shipments” in the ASM and CMF). Summing over all entries in a given row (i.e., location) n yields the total value of

plant-level revenues in location n , $\sum_J N_J r_{nJ}$. Similarly, summing over only those entries in row n that are associated with firms J that also have entries in row i yields $\sum_J N_J r_{nJ} \mathbf{1}_{ni}^J$. Dividing $\sum_J N_J r_{nJ} \mathbf{1}_{ni}^J$ by $\sum_J N_J r_{nJ}$ yields the entry s_{ni} in the \mathbf{S} matrix. Also, $\sum_J N_J r_{nJ}$ is the entry in row n of the vector \mathbf{r} , which is a 179×1 vector representing the total value of plant-level revenues in each location.

7.2 Ranking U.S. Tech Clusters

We rank all 179 cities based on their values of γ_n/β_n . Since this is too large a ranking to report, Table 8 only shows the top and bottom ten clusters with the highest and lowest values of γ_n/β_n , respectively. For comparison, the table also shows the top and bottom ten clusters ranked by size (number of inventors). What stands out, first and foremost, is that larger clusters have a higher social-private innovation wedge. Three out of the top ten clusters ranked by this wedge are also in the top ten ranked by size (marked in red in Table 8): Los Angeles-Long Beach-Riverside, San Jose-San Francisco-Oakland, and Chicago-Naperville-Michigan City. Moreover, seven out of the top ten clusters with the highest wedge are in the top decile ranked by size, and all ten are in the top quartile.

While there is a clear association between the social-private innovation wedge and cluster size, it is not perfect. Indeed, seven out of the top ten clusters ranked by this wedge are *not* in the top ten ranked by size. Thus, cluster size per se is not what is driving the ranking. In fact, cluster size per se is unrelated to it. Granted, larger clusters have more local knowledge spillovers and thus a higher social return to innovation. But firms' private incentives to invest in innovation are also greater in larger clusters, in proportion to the social return. As a result, if it was only for local knowledge spillovers, the social-private innovation wedge would be *invariant* to cluster size (c.f. equation (28) showing that in the absence of cross-location spillovers, the social-private innovation wedge is constant across locations). This invariance echoes a well-known result in urban economics, going back to Glaeser and Gottlieb (2008, 2009), stating that place-based redistributions are not welfare-enhancing unless the elasticity of productivity to agglomeration differs across locations. Importantly, what matters for the optimality of policy intervention is the *elasticity* with respect to agglomeration, not the size of the local economy.

What is instead driving the location ranking in Table 8—and breaking the Glaeser-Gottlieb logic in our case—is the extent to which clusters are connected to other

(especially productive and well-connected) clusters. In the presence of cross-cluster innovation spillovers, firms' innovative activities raise the productivity of other firms' inventors and plants in connected clusters. Since firms do not internalize these spillovers, clusters that are well-connected to other clusters have a high social-private innovation wedge. In our data, the correlation between γ_n/β_n and a simple empirical proxy for a cluster's connectedness—the average number of connected clusters across all plants in the cluster—is 41.1%.²⁶ Accordingly, innovation policy should focus on well-connected clusters, from where innovations spill over to other, well-connected clusters, and so on.

7.3 Increasing the Interconnectedness of U.S. Tech Clusters

As clusters are becoming more interconnected, the social returns to innovation increase. But so do firms' private incentives to innovate. The relative effect, captured by γ_n/β_n , is theoretically ambiguous. To explore the effect of an increase in the interconnectedness of clusters on the social-private innovation wedge, we again take our model to the data. Specifically, we counterfactually increase all off-diagonal elements of the \mathbf{S} matrix by 10% and compute the percentage change in γ_n/β_n for all 179 cities. (Recall that an element of the \mathbf{S} matrix represents the revenue-weighted fraction of innovating plants in location n that have a connected cluster in location i .)

While the effect of an increase in interconnectedness is theoretically ambiguous, it is relatively clear in the data: a 10% increase in the off-diagonals of the \mathbf{S} matrix raises γ_n/β_n by 2.6% on average. Hence, as clusters are becoming more connected with each other, the gap between the social and private returns to innovation widens, exacerbating the under-investment problem. However, while γ_n/β_n increases by 2.6% on average, there is significant heterogeneity across locations. At the 25th percentile, γ_n/β_n only increases by 0.8%, and a few clusters even experience a slight *decrease* in γ_n/β_n . On the other hand, clusters like Chicago-Naperville-Michigan City or Washington-Baltimore-Northern Virginia experience substantial increases of 7.2% and 6.6%, respectively. The median increase is 1.9%, which is well below the average.

Table 9 displays the top and bottom ten clusters with the highest and lowest percentage changes in γ_n/β_n , respectively. For comparison, the table also shows the

²⁶As well-connected clusters are typically also larger, there is also a positive correlation between γ_n/β_n and a cluster's own size. However, this correlation is only 16.6% and thus much weaker than the correlation with a cluster's connectedness.

top and bottom ten clusters ranked by size (number of inventors). Similar to Table 8, there is a clear visual association between the change in the social-private innovation wedge and cluster size. Two out of the top ten clusters with the highest increase are also in the top ten ranked by size (marked in red): Chicago-Naperville-Michigan City and Boston-Worcester-Manchester. Across all 179 cities, the correlation between the percentage change in γ_n/β_n and cluster size is 28.3%, and the correlation with the average number of connected clusters across all plants in a cluster—a simple proxy for a cluster’s connectedness—is 22.4%. Accordingly, an increase in the interconnectedness of clusters especially raises the social-private innovation wedge in large and well-connected clusters.

8 Conclusion

Within the firm, knowledge diffusion does not require face-to-face interactions but can take place through conventional channels, such as emails, internal memos, or video conferencing. Accordingly, if knowledge spreads locally across inventors of different firms, we would expect that these inventors pass on this knowledge to other inventors and plants within the same firm, across different locations. Hence, we would expect the same economic forces that generate local innovation spillovers within a tech cluster also to generate innovation spillovers across tech clusters, as long as the clusters are connected with each other through firms’ networks of innovating plants.

We test this hypothesis using merged USPTO–U.S. Census Bureau micro data on innovating (and non-innovating) plants. Cross-cluster innovation spillovers at the plant level are identified through within-plant variation in the size of the plant’s “connected clusters”—the number of other firms’ inventors in the plant’s research field in other cities in which the plant’s parent firm also has innovating plants—in combination with granular city \times field \times year fixed effects. We additionally employ an IV research design that isolates plausibly exogenous variation in the size of a plant’s connected clusters based on the expansion of other firms in other cities that do not overlap with the geographical network of the plant’s parent firm. We find strong and robust evidence of cross-cluster innovation spillovers. Moreover, and consistent with a knowledge diffusion channel, these spillovers do not decay with physical distance, and plants cite disproportionately more patents by other firms in connected clusters as the size of those clusters increases.

Firms are unlikely to fully internalize the impacts of their innovative activities on

other firms’ innovation and productivity. To inform policy, we develop a tractable model of spatial innovation that features both within- and cross-cluster innovation spillovers. We show that the optimal government policy is plant-specific and disproportionately targets plants that are well-connected to other clusters, especially clusters that are in turn well-connected to other clusters, thereby generating the biggest “bang for the buck.” Importantly, we derive a sufficient statistic—the “social-private innovation wedge”—that captures the wedge between the social and private returns to innovation in a location. Taking our model to the data, we rank all U.S. tech clusters based on their value of this statistic. We find that larger clusters exhibit a greater gap between the social and private returns to innovation; however, this is not because of local knowledge spillovers, but because larger clusters are more connected to other clusters through firms’ networks of innovating plants. In a counterfactual exercise, we strengthen the interconnectedness of U.S. tech clusters. We find that the social-private innovation wedge increases almost everywhere, but especially so in large and well-connected clusters.

References

- Alfaro-Ureña, Alonso, Isabela Manelici, and José Vasquez, 2021, The Effects of Multinationals on Workers: Evidence from Costa Rican Microdata, mimeo, LSE.
- Arkolakis, Costas, Eckert, Fabian Eckert, and Rowan Shi, 2023, Combinatorial Discrete Choice, mimeo, Yale University.
- Arqué-Castells, Pere, and Daniel Spulber, 2022, Measuring the Private and Social Returns to R&D: Unintended Spillovers versus Technology Markets, *Journal of Political Economy* 130, 1860–1918.
- Atkin, David, Keith Chen, and Anton Popov, 2022, The Returns to Face-to-Face Interactions: Knowledge Spillovers in Silicon Valley, NBER Working Paper 30147.
- Audretsch, David, and Maryann Feldman, 1996, R&D Spillovers and the Geography of Innovation and Production, *American Economic Review* 86, 630–640.
- Audretsch, David, and Maryann Feldman, 2004, Knowledge Spillovers and the Geography of Innovation, in: *Handbook of Regional and Urban Economics*, Volume 4, Vernon Henderson and Jacques Thisse (eds.), 2713–2739. Amsterdam: Elsevier.

- Balasubramanian, Natarajan, and Jagadeesh Sivadasan, 2011, What Happens When Firms Patent? New Evidence from U.S. Economic Census Data, *Review of Economics and Statistics* 93, 126–146.
- Bena, Jan, Serdar Dinc, and Isil Erel, 2022, The International Propagation of Economic Downturns through Multinational Companies: The Real Economy Channel, *Journal of Financial Economics* 146, 277–304.
- Bettencourt, Luis, José Lobo, and Deborah Strumsky, 2007, Invention in the City: Increasing Returns to Patenting as a Scaling Function of Metropolitan Size, *Research Policy* 36, 107–120.
- Biermann, Marcus, and Kilian Huber, 2024, Tracing the International Transmission of a Crisis through Multinational Firms, *Journal of Finance*, forthcoming.
- Bilir, Kamran, and Eduardo Morales, 2020, Innovation in the Global Firm, *Journal of Political Economy* 128, 1566–1625.
- Bloom, Nicholas, Mark Schankerman, and John Van Reenen, 2013, Identifying Technology Spillovers and Product Market Rivalry, *Econometrica* 81, 1347–1393.
- Bloom, Nicholas, John Van Reenen, and Heidi Williams, 2019, A Toolkit of Policies to Promote Innovation, *Journal of Economic Perspectives* 33, 163–184.
- Boehm, Christoph, Aaron Flaaen, and Nitya Pandalai-Nayar, 2019, Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake, *Review of Economics and Statistics* 101, 60–75.
- Carlino, Gerald, and William Kerr, 2015, Agglomeration and Innovation, in: *Handbook of Regional and Urban Economics*, Volume 5, Gilles Duranton, Vernon Henderson, and William Strange (eds.), 349–404. Amsterdam: Elsevier.
- Chattergoon, Brad, and William Kerr, 2022, Winner Takes All? Tech Clusters, Population Centers, and the Spatial Transformation of U.S. Invention, *Research Policy* 51, 104418.
- Cohen, Wesley, and Daniel Levinthal, 1989, The Two Faces of R&D, *Economic Journal* 99, 569–596.

- Cravino, Javier, and Andrei Levchenko, 2017, Multinational Firms and International Business Cycle Transmission, *Quarterly Journal of Economics* 132, 921–962.
- Ding, Xiang, 2023, Industry Linkages from Joint Production, mimeo, Georgetown University.
- Dreisigmeyer, David, Nathan Goldschlag, Marina Krylova, Wei Ouyang, and Elisabeth Perlman, 2018, Building a Better Bridge: Improving Patent Assignee-Firm Links, CES Technical Notes, CES-TN-2018-01, U.S. Census Bureau.
- Gaubert, Cecile, 2018, Firm Sorting and Agglomeration, *American Economic Review* 108, 3117–3153.
- Gaubert, Cecile, Patrick Kline, Damian Vergara, and Danny Yagan, 2020, Place-Based Redistribution, mimeo, UC Berkeley.
- Giroud, Xavier, and Holger Mueller, 2019, Firm’s Internal Networks and Local Economic Shocks, *American Economic Review* 109, 3617–3649.
- Giroud, Xavier, Simone Lenzu, Quinn Maingi, and Holger Mueller, 2023, Propagation and Amplification of Local Productivity Spillovers, mimeo, Columbia University.
- Glaeser, Edward, and Joshua Gottlieb, 2008, The Economics of Place-Making Policies, *Brookings Papers on Economic Activity*, Spring, 155–239.
- Glaeser, Edward, and Joshua Gottlieb, 2009, The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States, *Journal of Economic Literature* 47, 983–1028.
- Griffith, Rachel, Rupert Harrison, and John Van Reenen, 2006, How Special is the Special Relationship: Using the Impact of US R&D Spillovers on British Firms as a Test of Technology Sourcing, *American Economic Review* 96, 1859–1875.
- Griliches, Zvi, 1979, Issues in Assessing the Contribution of Research and Development to Productivity Growth, *Bell Journal of Economics* 10, 92–116.
- Gruber, Jonathan, Enrico Moretti, and Daniel Wilson, 2023, Place-Based Productivity and Costs in Science, in: *Entrepreneurship and Innovation Policy and the Economy*,

- Volume 2, Benjamin Jones and Josh Lerner (eds.), 167–184, Chicago: University of Chicago Press.
- Hall, Bronwyn, Jacques Mairesse, and Pierre Mohnen, 2010, Measuring the Returns to R&D, in: *Handbook of the Economics of Innovation*, Volume 2, Bronwyn Hall and Nathan Rosenberg (eds.), 1033–1082. Amsterdam: Elsevier.
- Jaffee, Adam, 1986, Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits, and Market Value, *American Economic Review* 76, 984–1001.
- Jaffe, Adam, Manuel Trajtenberg, and Rebecca Henderson, 1993, Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations, *Quarterly Journal of Economics* 108, 577–598.
- Kerr, William, and Shihe Fu, 2008, The Survey of Industrial R&D—Patent Database Link Project, *Journal of Technology Transfer* 33, 173–186.
- Kerr, William, and Frederic Robert-Nicoud, 2020, Tech Clusters, *Journal of Economic Perspectives* 34, 50–76.
- Kline, Patrick, and Enrico Moretti, 2014, People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs, *Annual Review of Economics* 6, 629–662.
- Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar, 2019, Who Profits from Patents? Rent Sharing at Innovative Firms, *Quarterly Journal of Economics* 134, 1343–1404.
- Liu, Ernest, and Song Ma, 2023, Innovation Networks and R&D Allocation, mimeo, Princeton University.
- Lychagin, Sergey, Joris Pinkse, Margaret Slade, and John Van Reenen, 2016, Spillovers in Space: Does Geography Matter? *Journal of Industrial Economics* 64, 295–335.
- Markusen, James, 1984, Multinationals, Multi-Plant Economies, and the Gains from Trade, *Journal of International Economics* 16, 205–226.
- Moretti, Enrico, 2004, Workers' Education, Spillovers, and Productivity: Evidence from Plant-Level Production Functions, *American Economic Review* 94, 656–690.

- Moretti, Enrico, 2021, The Effect of High-Tech Clusters on the Productivity of Top Inventors, *American Economic Review* 111, 3328–3375.
- Moretti, Enrico, and Daniel Wilson, 2017, The Effect of State Taxes on the Geographical Location of Top Earners: Evidence from Star Scientists, *American Economic Review* 107, 1858–1903.
- Neumark, David, and Helen Simpson, 2015, Place-Based Policies, in: *Handbook of Regional and Urban Economics*, Volume 5, Gilles Duranton, Vernon Henderson, and William Strange (eds.), 1197–1287. Amsterdam: Elsevier.
- Oberfield, Ezra, Esteban Rossi-Hansberg, Pierre-Daniel Sarte, and Nicholas Trachter, 2024, Plants in Space, *Journal of Political Economy* 132, 867–909.
- Saxenian, AnnaLee, 1994, *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*, Harvard University Press.
- Setzler, Bradley, and Felix Tintelnot, 2021, The Effects of Foreign Multinationals on Workers and Firms in the United States, *Quarterly Journal of Economics* 136, 1943–1991.

Table 1: Descriptive Statistics

This table presents plant- and firm-level statistics for the main analysis sample consisting of 57,000 plant-year observations and 11,000 firm-year observations. Patents is the total number of patents (in the plant's field) per plant and year. Connected clusters is the number of other cities in which the firm has plants (in the plant's field) with inventors. Connected cluster size is the total number of other firms' inventors (in the plant's field) across all connected clusters. Cities are BEA economic areas. Fields are based on 1-digit CPC codes. Inventor share at the firm-plant level is the ratio of the firm's plants with inventors to all plants of the firm, averaged across all firms. Inventor share at the firm-county (city, state) level is the ratio of counties (cities, states) in which the firm has plants with inventors to all counties (cities, states) in which the firm has plants, averaged across all firms. The sample period is from 1976 to 2018.

	(1)	(2)
	Mean	Std. dev.
<i>Plants:</i>		
Employees	699	1,623
Shipments (\$M)	275	1,670
Inventors	4.99	13.67
Patents	10.87	43.07
Connected Clusters	3.22	2.81
Connected Cluster Size	1,915	2,748
<i>Firms:</i>		
Employees	6,853	13,280
Shipments (\$M)	2,524	6,730
Plants	17.51	23.19
Counties	15.23	18.47
Cities	11.77	11.72
States	9.16	7.33
<i>Inventor Share:</i>		
Plants	0.56	0.31
Counties	0.55	0.31
Cities	0.52	0.33
States	0.49	0.34

Table 2: Within-Cluster Innovation Spillovers

This table presents variants of the regression in equation (1) in which the dependent variable is either the ratio (in logs) of the number of patents to the number of inventors (all in the plant's field) at the plant level or plant-level TFP. Cluster size is the number of other firms' inventors in the plant's city and field. In column (5), the sample consists of plants without inventors. For these plants, fixed effects are based on the firm's main field, and cluster size is the total number of other firms' inventors in the plant's city in the firm's main field. Observations are weighted by plant employment. Standard errors are double clustered at the city and year level. The sample period is from 1976 to 2018. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Log(Pat/Inv)	TFP	Log(Pat/Inv)	TFP	TFP
	(1)	(2)	(3)	(4)	(5)
Log(Cluster Size)	0.078*** (0.010)	0.029*** (0.004)	0.076*** (0.013)	0.023*** (0.007)	0.003 (0.003)
Field FE	Yes	Yes	-	-	-
Year FE	Yes	Yes	-	-	-
Field × Year FE	No	No	Yes	Yes	Yes
City × Year FE	No	No	Yes	Yes	Yes
City × Field FE	No	No	Yes	Yes	Yes
Plant FE	Yes	Yes	Yes	Yes	Yes
Observations	134,000	134,000	134,000	134,000	259,000
R-squared	0.39	0.67	0.44	0.70	0.74

Table 3: Cross-Cluster Innovation Spillovers

This table presents variants of the regression in equation (2) in which the dependent variable is either the ratio (in logs) of the number of patents to the number of inventors (all in the plant's field) at the plant level or plant-level TFP. (Connected) cluster size is the total number of other firms' inventors (in the plant's field) across all other cities in which the firm has plants (in the plant's field) with inventors. Observations are weighted by plant employment. Standard errors are double clustered at the city and year level. The sample period is from 1976 to 2018. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Log(Pat/Inv)	TFP	Log(Pat/Inv)	TFP	Log(Pat/Inv)	TFP
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Cluster Size)	0.024*** (0.005)	0.010*** (0.003)	0.023*** (0.007)	0.012*** (0.004)	0.021*** (0.007)	0.012*** (0.004)
Field FE	Yes	Yes	-	-	-	-
Year FE	Yes	Yes	-	-	-	-
Field × Year FE	No	No	Yes	Yes	-	-
City × Year FE	No	No	Yes	Yes	-	-
City × Field FE	No	No	Yes	Yes	-	-
City × Field × Year FE	No	No	No	No	Yes	Yes
Plant FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	57,000	57,000	57,000	57,000	57,000	57,000
R-squared	0.45	0.69	0.52	0.75	0.59	0.79

Table 4: Non-Innovating Plants

This table presents variants of the regressions in columns (5) and (6) of Table 3. In columns (1) and (2), (connected) cluster size is the total number of other firms' inventors (in the plant's field) across all other cities in which the firm has plants but no inventors. In column (3), the sample consists of plants without inventors. For these plants, fixed effects are based on the firm's main field, and (connected) cluster size is the total number of other firms' inventors (in the firm's main field) across all other cities in which the firm has plants (in the firm's main field) with inventors. Observations are weighted by plant employment. Standard errors are double clustered at the city and year level. The sample period is from 1976 to 2018. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Log(Pat/Inv)	TFP	TFP
	(1)	(2)	(3)
Log(Cluster Size)	0.002 (0.008)	0.001 (0.005)	0.006** (0.003)
City × Field × Year FE	Yes	Yes	Yes
Plant FE	Yes	Yes	Yes
Observations	80,000	80,000	227,000
R-squared	0.52	0.75	0.82

Table 5: Instrumental Variable Estimation

This table presents variants of the regressions in columns (5) and (6) of Table 3 estimated by 2SLS. The excluded instrument, Z , is the number of inventors (in the plant's field) in other cities where the firm has no presence and which are employed by other firms that also have inventors in the plant's connected clusters. Column (1) shows the first stage; columns (2) and (3) shows 2SLS estimates. Observations are weighted by plant employment. Standard errors are double clustered at the city and year level. The sample period is from 1976 to 2018. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Log(Cluster Size)	Log(Pat/Inv)	TFP
	(1)	(2)	(3)
Z	0.107*** (0.013)		
Log(Cluster Size)		0.023*** (0.011)	0.014*** (0.007)
City \times Field \times Year FE	Yes	Yes	Yes
Plant FE	Yes	Yes	Yes
Observations	57,000	57,000	57,000
R-squared	0.85	0.59	0.79
F-statistic	70.60		

Table 6: Geographical Distance

This table presents variants of the regressions in columns (5) and (6) of Table 3 in which connected clusters within a 100, 250, or 500 mile radius around the plant are excluded. Observations are weighted by plant employment. Standard errors are double clustered at the city and year level. The sample period is from 1976 to 2018. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Log(Pat/Inv)			TFP		
	(1)	(2)	(3)	(4)	(5)	(6)
	Excluding clusters within 100 miles	Excluding clusters within 250 miles	Excluding clusters within 500 miles	Excluding clusters within 100 miles	Excluding clusters within 250 miles	Excluding clusters within 500 miles
Log(Cluster Size)	0.021*** (0.007)	0.020*** (0.007)	0.019** (0.008)	0.011*** (0.004)	0.011** (0.005)	0.011** (0.005)
City × Field × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Plant FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	53,000	46,000	35,000	53,000	46,000	35,000
R-squared	0.60	0.60	0.61	0.79	0.80	0.82

Table 7: Patent Citations

This table presents variants of the regressions in Table 3 and columns (1) and (2) of Table 4 in which the dependent variable is the share of the plant's citations to other firms' patents from connected clusters relative to all its citations to other firms' patents (in the plant's field). In column (4), citations to patents filed in year t are excluded. In column (5), (connected) cluster size is the total number of other firms' inventors (in the plant's field) across all other cities in which the firm has plants but no inventors. Observations are weighted by plant employment. Standard errors are double clustered at the city and year level. The sample period is from 1976 to 2018. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Patent Citation Share				
	(1)	(2)	(3)	(4)	(5)
Log(Cluster Size)	0.018** (0.007)	0.022*** (0.008)	0.019** (0.008)	0.019** (0.008)	0.002 (0.010)
Field FE	Yes	-	-	-	-
Year FE	Yes	-	-	-	-
Field \times Year FE	No	Yes	-	-	-
City \times Year FE	No	Yes	-	-	-
City \times Field FE	No	Yes	-	-	-
City \times Field \times Year FE	No	No	Yes	Yes	Yes
Plant FE	Yes	Yes	Yes	Yes	Yes
Observations	57,000	57,000	57,000	57,000	80,000
R-squared	0.55	0.63	0.69	0.69	0.58

Table 8: Ranking U.S. Tech Clusters

This table ranks U.S. tech clusters based on their social-private innovation wedge, γ_n/β_n (left) and cluster size (right). Clusters are the 179 BEA economic areas (“cities”). Cluster size is the number of inventors in a city. Clusters are listed in ascending order.

Top 10 (Social-Private Innovation Wedge)	Top 10 (Cluster Size)
Atlanta-Sandy Springs-Gainesville, GA-AL Denver-Aurora-Boulder, CO Washington-Baltimore-Northern Virginia, DC Los Angeles-Long Beach-Riverside, CA San Jose-San Francisco-Oakland, CA Philadelphia-Camden-Vineland, PA-NJ-DE-MD Chicago-Naperville-Michigan City, IL-IN-WI Phoenix-Mesa-Scottsdale, AZ Raleigh-Durham-Cary, NC Houston-Baytown-Huntsville, TX	San Jose-San Francisco-Oakland, CA New York-Newark-Bridgeport, NY-NJ-CT-PA Minneapolis-St. Paul-St. Cloud, MN-WI Detroit-Warren-Flint, MI Boston-Worcester-Manchester, MA-NH Los Angeles-Long Beach-Riverside, CA Chicago-Naperville-Michigan City, IL-IN-WI Portland-Vancouver-Beaverton, OR-WA Dallas-Fort Worth, TX Cleveland-Akron-Elyria, OH
Bottom 10 (Social-Private Innovation Wedge)	Bottom 10 (Cluster Size)
Dothan-Enterprise-Ozark, AL Cape Girardeau-Jackson, MO-IL Farmington, NM Pendleton-Hermiston, OR Alpena, MI Bismarck, ND Minot, ND Bangor, ME Scotts Bluff, NE Great Falls, MT	San Angelo, TX Rapid City, SD Pendleton-Hermiston, OR Honolulu, HI Abilene, TX Aberdeen, SD Minot, ND Great Falls, MT Anchorage, AK Santa Fe-Espanola, NM

Table 9: Increasing the Interconnectedness of U.S. Tech Clusters

This table ranks U.S. tech clusters based on the percentage change in their social-private innovation wedge following a 10% increase in the off-diagonal elements of the **S** matrix (left) and cluster size (right). Clusters are the 179 BEA economic areas (“cities”). Cluster size is the number of inventors in a city. Clusters are listed in ascending order.

Top 10 (Percentage Change in Wedge)	Top 10 (Cluster Size)
<p>Chicago-Naperville-Michigan City, IL-IN-WI (7.2%)</p> <p>Washington-Baltimore-Northern Virginia, DC (6.6%)</p> <p>Austin-Round Rock, TX (6.2%)</p> <p>Boston-Worcester-Manchester, MA-NH (6.1%)</p> <p>Indianapolis-Anderson-Columbus, IN (6.1%)</p> <p>Houston-Baytown-Huntsville, TX (5.8%)</p> <p>San Diego-Carlsbad-San Marcos, CA (5.6%)</p> <p>Rochester-Batavia-Seneca Falls, NY (5.5%)</p> <p>Seattle-Tacoma-Olympia, WA (5.3%)</p> <p>Pittsburgh-New Castle, PA (5.3%)</p>	<p>San Jose-San Francisco-Oakland, CA</p> <p>New York-Newark-Bridgeport, NY-NJ-CT-PA</p> <p>Minneapolis-St. Paul-St. Cloud, MN-WI</p> <p>Detroit-Warren-Flint, MI</p> <p>Boston-Worcester-Manchester, MA-NH</p> <p>Los Angeles-Long Beach-Riverside, CA</p> <p>Chicago-Naperville-Michigan City, IL-IN-WI</p> <p>Portland-Vancouver-Beaverton, OR-WA</p> <p>Dallas-Fort Worth, TX</p> <p>Cleveland-Akron-Elyria, OH</p>
Bottom 10 (Percentage Change in Wedge)	Bottom 10 (Cluster Size)
<p>Amarillo, TX (0.4%)</p> <p>Flagstaff, AZ (0.4%)</p> <p>Abilene, TX (0.4%)</p> <p>Kearney, NE (0.3%)</p> <p>Alpena, MI (0.2%)</p> <p>Gulfport-Biloxi-Pascagoula, MS (0.0%)</p> <p>Jonesboro, AR (-0.0%)</p> <p>Fort Smith, AR-OK (-0.1%)</p> <p>Paducah, KY-IL (-0.2%)</p> <p>Pueblo, CO (-0.2%)</p>	<p>San Angelo, TX</p> <p>Rapid City, SD</p> <p>Pendleton-Hermiston, OR</p> <p>Honolulu, HI</p> <p>Abilene, TX</p> <p>Aberdeen, SD</p> <p>Minot, ND</p> <p>Great Falls, MT</p> <p>Anchorage, AK</p> <p>Santa Fe-Espanola, NM</p>

Innovation Spillovers across U.S. Tech Clusters

Xavier Giroud Ernest Liu Holger Mueller

Online Appendix

A Additional Figures and Tables

Figure A.1: BEA Economic Areas




Table A.1: Robustness

This table presents variants of the regressions in columns (5) and (6) of Table 3. In columns (1) and (2), inventors are assigned to the nearest firm plant in the inventor's city. In columns (3) and (4), the plant's field is based on all fields in which it has patents, and its patent output, inventors, and connected clusters are all based on those fields. In columns (5) and (6), a broader concept of cluster size is employed, where it is assumed that inventors work in a cluster also in between patent filings. In column (7), the sample includes all LBD establishments with inventor matches. Observations are weighted by plant employment. Standard errors are double clustered at the city and year level. The sample period is from 1976 to 2018. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Log(Pat/Inv)	TFP	Log(Pat/Inv)	TFP	Log(Pat/Inv)	TFP	Log(Pat/Inv)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log(Cluster Size)	0.022*** (0.005)	0.011*** (0.004)	0.018*** (0.006)	0.010** (0.004)	0.018*** (0.006)	0.010*** (0.003)	0.018*** (0.003)
City × Field × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	70,000	70,000	57,000	57,000	57,000	57,000	383,000
R-squared	0.56	0.77	0.61	0.80	0.59	0.79	0.71

Table A.2: Number of Plant Inventors

This table presents variants of the regressions in columns (3) and (4) of Table 2 and columns (5) and (6) of Table 3 in which (connected) cluster size is interacted with the number of plant inventors (in logs). In columns (1) to (4), number of inventors refers to the focal plant. In columns (5) and (6), number of inventors refers to the firm's innovating plants in connected clusters. Observations are weighted by plant employment. Standard errors are double clustered at the city and year level. The sample period is from 1976 to 2018. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Within-Cluster			Cross-Cluster		
	Log(Pat/Inv)	TFP	Log(Pat/Inv)	TFP	Log(Pat/Inv)	TFP
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Cluster Size)	0.054*** (0.013)	0.015*** (0.006)	0.017*** (0.007)	0.009*** (0.004)	0.018*** (0.008)	0.009*** (0.004)
Log(Cluster Size) × Log(Inventors)	0.025*** (0.002)	0.012*** (0.002)	0.004*** (0.002)	0.003*** (0.002)	0.002*** (0.001)	0.002*** (0.001)
Log(Inventors)	0.101*** (0.049)	0.010*** (0.002)	0.005*** (0.002)	0.008*** (0.003)	0.003*** (0.002)	0.004* (0.002)
Field × Year FE	Yes	Yes	-	-		
City × Year FE	Yes	Yes	-	-		
City × Field FE	Yes	Yes	-	-		
City × Field × Year FE	-	-	Yes	Yes	Yes	Yes
Plant FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	134,000	134,000	57,000	57,000	57,000	57,000
R-squared	0.44	0.70	0.59	0.79	0.59	0.79

Table A.3: Aggregation to the Cluster Level

This table presents variants of the regressions in columns (3) and (4) of Table 3 in which observations are aggregated at the city \times field \times year level. The dependent variable is either the ratio (in logs) of the total number of patents to the total number of inventors based on all plants in the cluster or the employment-weighted average TFP across all plants in the cluster. (Connected) cluster size is either the sum (columns (1) and (2)) or average (columns (3) and (4)) of the plant-level (connected) cluster sizes across all plants in the cluster. Observations are weighted by city \times field employment. Standard errors are double clustered at the city and year level. The sample period is from 1976 to 2018. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Log(Pat/Inv)	TFP	Log(Pat/Inv)	TFP
	(1)	(2)	(3)	(4)
Log(Cluster Size)	0.016** (0.008)	0.010** (0.004)	0.017** (0.008)	0.011** (0.005)
Field \times Year FE	Yes	Yes	Yes	Yes
City \times Year FE	Yes	Yes	Yes	Yes
City \times Field FE	Yes	Yes	Yes	Yes
Observations	18,000	18,000	18,000	18,000
R-squared	0.50	0.48	0.50	0.48

Table A.4: Patent Output

This table presents variants of the regressions in Table 3 and columns (1) and (2) of Table 4 in which the dependent variable is the log number of patents (in the plant's field) at the plant level. In column (4), (connected) cluster size is the total number of other firms' inventors (in the plant's field) across all other cities in which the firm has plants but no inventors. Observations are weighted by plant employment. Standard errors are double clustered at the city and year level. The sample period is from 1976 to 2018. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Log(Patents)			
	(1)	(2)	(3)	(4)
Log(Cluster Size)	0.036*** (0.006)	0.031*** (0.006)	0.030*** (0.007)	0.004 (0.009)
Field FE	Yes	-	-	-
Year FE	Yes	-	-	-
Field × Year FE	No	Yes	-	-
City × Year FE	No	Yes	-	-
City × Field FE	No	Yes	-	-
City × Field × Year FE	No	No	Yes	Yes
Plant FE	Yes	Yes	Yes	Yes
Observations	57,000	57,000	57,000	80,000
R-squared	0.63	0.69	0.74	0.61

B Proofs

B.1 Proof of Lemma 1

Let $x \equiv \{h_{iJ}, K_i\}_{i \in J}$ and $\hat{x} \equiv \{\hat{h}_{iJ}, \hat{K}_i\}_{i \in J}$ denote two equilibria, and let $\tilde{y} \equiv \ln \hat{y} - \ln y$ for $y \in \{h_{iJ}, K_i\}_{i \in J}$. Let $\zeta \equiv \|\ln x - \ln \hat{x}\|_\infty$. In an equilibrium, $\{h_{iJ}, K_i\}$ must satisfy:

$$c_q h_{qJ} = \alpha \sum_n b_n \prod_{i \in J} \left(h_{iJ}^\alpha K_i^\delta \right)^{\omega_{ni}} \omega_{nq},$$

$$b_n = \sum_k \frac{1}{\sigma} \frac{\tau_{kn}^{1-\sigma}}{\sum_m \sum_{J' \ni m} N_{J'} \prod_{i \in J'} \left(h_{iJ'}^\alpha K_i^\delta \right)^{\omega_{mi}} \tau_{km}^{1-\sigma}},$$

and

$$K_n \equiv \sum_{J \ni n} N_J \prod_{i \in J} \left(h_{iJ}^\alpha K_i^\delta \right)^{\omega_{ni}}.$$

Define f_{qJ} and g_n as the equilibrium correspondence with:

$$h_{qJ} \equiv f_{qJ}(x)$$

and

$$K_n \equiv g_n(x).$$

By the intermediate value theorem, there must exist $x' \equiv \exp(t \ln x + (1-t) \ln \hat{x})$ for $t \in [0, 1]$ such that:

$$\tilde{h}_{qJ} = \sum_{iJ'} \frac{\partial \ln f_{qJ}(x')}{\partial \ln h_{iJ'}} \tilde{h}_{iJ'} + \sum_i \frac{\partial \ln f_{qJ}(x')}{\partial \ln K_i} \tilde{K}_i$$

and

$$\tilde{K} = \sum_{iJ'} \frac{\partial \ln g_n(x')}{\partial \ln h_{iJ'}} \tilde{h}_{iJ'} + \sum_i \frac{\partial \ln g_n(x')}{\partial \ln K_i} \tilde{K}_i.$$

Hence:

$$\begin{aligned}
|\widetilde{h}_{qJ}| &\leq \sum_{ij'} \left| \frac{\partial \ln f_{qJ}(x')}{\partial \ln h_{ij'}} \right| |\widetilde{h}_{ij'}| + \sum_i \left| \frac{\partial \ln f_{qJ}(x')}{\partial \ln K_i} \right| |\widetilde{K}_i| \\
&\leq \left(\sum_{ij'} \left| \frac{\partial \ln f_{qJ}(x')}{\partial \ln h_{ij'}} \right| + \sum_i \left| \frac{\partial \ln f_{qJ}(x')}{\partial \ln K_i} \right| \right) \zeta \\
&\leq (\alpha + \delta) (1 + (\mathcal{J} - 1) \omega) \zeta
\end{aligned}$$

and

$$\begin{aligned}
|\widetilde{K}| &\leq \sum_{ij'} \left| \frac{\partial \ln g_n(x')}{\partial \ln h_{ij'}} \right| |\widetilde{h}_{ij'}| + \sum_i \left| \frac{\partial \ln g_n(x')}{\partial \ln K_i} \right| |\widetilde{K}_i| \\
&\leq \left(\sum_{ij'} \left| \frac{\partial \ln g_n(x')}{\partial \ln h_{ij'}} \right| + \sum_i \left| \frac{\partial \ln g_n(x')}{\partial \ln K_i} \right| \right) \zeta \\
&\leq (\alpha + \delta) (1 + (\mathcal{J} - 1) \omega) \zeta.
\end{aligned}$$

We thus have $\zeta \leq (\alpha + \delta) (1 + (\mathcal{J} - 1) \omega) \zeta$. A sufficient condition for $\zeta = 0$ is that $(\alpha + \delta) (1 + (\mathcal{J} - 1) \omega) < 1$, establishing uniqueness.

B.2 Proof of Lemma 2

Equations (5) and (7) imply:

$$y_{ni}(v) = \frac{\sigma - 1}{\sigma} \frac{(\tau_{ni} \kappa_i(v))^{-\sigma}}{\sum_i \int_v (\tau_{ni} \kappa_i(v))^{1-\sigma} dv}.$$

Substituting into the definition of Y_n in (3) and using the fact that $\kappa_{nJ}^{1-\sigma} = k_{nJ}$, we obtain:

$$\begin{aligned}
\ln Y_n &= \ln \left[\sum_i \int y_{ni}(\nu)^{\frac{\sigma-1}{\sigma}} d\nu \right]^{\frac{\sigma}{\sigma-1}} \\
&= \frac{\sigma}{\sigma-1} \ln \left[\sum_i \int_{\nu} (\tau_{ni} \kappa_i(\nu))^{1-\sigma} d\nu \right] - \ln \left(\sum_i \int_{\nu} (\tau_{ni} \kappa_i(\nu))^{1-\sigma} d\nu \right) + \ln \frac{\sigma-1}{\sigma} \\
&= \frac{1}{\sigma-1} \ln \left[\sum_i \int_{\nu} (\tau_{ni} \kappa_i(\nu))^{1-\sigma} d\nu \right] + \ln \frac{\sigma-1}{\sigma} \\
&= \frac{1}{\sigma-1} \ln \left[\sum_i \sum_{J \ni i} \tau_{ni}^{1-\sigma} N_J k_{iJ} \right] + \ln \frac{\sigma-1}{\sigma}.
\end{aligned}$$

Taking the derivative with respect to k_{iJ} , summing across n , and substituting using the definition of b_i in (13), we have:

$$\begin{aligned}
\sum_n \frac{\partial \ln Y_n}{\partial k_{iJ}} &= \sum_n \frac{1}{\sigma-1} \frac{N_J \tau_{ni}^{1-\sigma}}{\sum_m \sum_{J \ni m} \tau_{nm}^{1-\sigma} N_J k_{mJ}} \\
&= \frac{\sigma}{\sigma-1} N_J b_i.
\end{aligned}$$

Lemma 2 then follows immediately from the definitions of \mathcal{W} and $\hat{\mathcal{W}}$.

B.3 Proof of Proposition 2

Note that:

$$\beta_i \equiv \sum_{J \ni i} N_J \beta_{iJ}, \quad \gamma_i \equiv \sum_{J \ni i} N_J \gamma_{iJ}.$$

Equation (20) implies:

$$\begin{aligned}
\beta_i &= \sum_{J \ni i} N_J \beta_{iJ} \\
&= N_J \alpha \sum_n \omega_{ni} \sum_{J \ni n} \pi_{nJ} \mathbf{1}_{ni}^J \\
&= \frac{\alpha}{\sigma} \sum_n \omega_{ni} \left(\sum_{J \ni n} N_J r_{nJ} \right) \frac{\sum_{J \ni n} N_J r_{nJ} \mathbf{1}_{ni}^J}{\sum_{J \ni n} N_J r_{nJ}} \\
&= \frac{\alpha}{\sigma} \sum_n \omega_{ni} r_n s_{ni},
\end{aligned}$$

where $r_n \equiv \sum_{J' \ni n} N_{J'} r_{nJ'}$ and $s_{ni} \equiv \frac{\sum_J N_J r_{nJ} \mathbf{1}_{ni}^J}{\sum_J N_J r_{nJ}}$. In vector form:

$$\beta' = \frac{\alpha}{\sigma} \mathbf{r}' (\Omega \circ \mathbf{S}).$$

Equation (24) implies:

$$\sum_{J \ni i} \frac{d \ln K_n}{d \ln h_{iJ}} = \alpha s_{ni} \omega_{ni} + \delta s_{nq} \omega_{nq} \sum_{J \ni i} \frac{d \ln K_q}{d \ln h_{iJ}}.$$

In vector form:

$$\left[\sum_{J \ni i} \frac{d \ln K_n}{d \ln h_{iJ}} \right] = \alpha (I - \delta \Omega \circ \mathbf{S})^{-1} \Omega \circ \mathbf{S}. \quad (\text{B.1})$$

We thus have:

$$\begin{aligned} \gamma_i &\equiv \sum_{J \ni i} N_J \gamma_{iJ} \\ &= \sum_{J \ni i} \frac{d \hat{\mathcal{W}}}{d \ln h_{iJ}} \\ &= \frac{\sigma}{\sigma - 1} \left(\frac{\alpha}{\sigma} \sum_n r_n s_{ni} \omega_{ni} + \frac{\delta}{\sigma} \sum_n r_n \sum_m s_{nm} \omega_{nm} \sum_{J \ni i} \frac{d \ln K_m}{d \ln h_{iJ}} \right). \end{aligned}$$

In vector form:

$$\begin{aligned} \gamma' &= \frac{\sigma}{\sigma - 1} \frac{\alpha}{\sigma} \mathbf{r}' (\Omega \circ \mathbf{S}) (I + \delta \Omega \circ \mathbf{S} (I - \delta \Omega \circ \mathbf{S})^{-1}) \\ &= \frac{\sigma}{\sigma - 1} \frac{\alpha}{\sigma} \mathbf{r}' (\Omega \circ \mathbf{S}) (I - \delta \Omega \circ \mathbf{S})^{-1} \\ &= \frac{\sigma}{\sigma - 1} \beta' (I - \delta \Omega \circ \mathbf{S})^{-1}. \end{aligned}$$

C Extensions

C.1 Heterogeneous Location Size

In our baseline model, the population in each location is normalized to one. We now extend the model to allow for population heterogeneity; precisely, we specify that location i has population μ_i . In this case, the total expenditure in location i is μ_i . The variable profit of variety v produced in location i that is sold to location n is:

$$\pi_{ni}(v) = \mu_n \frac{p_{ni}(v)^{1-\sigma}}{\sum_i \int_v p_{ni}(v)^{1-\sigma} dv},$$

and the total variable profit of a plant is:

$$\pi_i(v) = b_i \kappa_i(v)^{1-\sigma},$$

where b_i is the market access for a plant in location i :

$$b_i \equiv \sum_n \frac{\mu_n}{\sigma} \frac{\tau_{ni}^{1-\sigma}}{\sum_m \int_v (\kappa_m(v) \tau_{nm})^{1-\sigma} dv}. \quad (\text{C.1})$$

Note that, intuitively, the market access b_i for a plant in location i weighs each market n by its population μ_n . In the baseline model, $\mu_n = 1$ (as in (13)).

The utilitarian social welfare is $\mathcal{W} \equiv \sum_n \mu_n U_n$. It is easy to show that Lemma 2 continues to hold in this extended setting, and the marginal impact of a plant's innovation investment on social welfare is equal to the marginal impact on \mathcal{W} as defined in (19), with population-adjusted market access terms as in (C.1). Since our subsequent analysis is based on analyzing the impact of innovation investments on \mathcal{W} , our analysis (including the sufficient statistics) continues to hold in this setting without modification.

C.2 Innovation Decisions along the Extensive Margin

In our baseline model, all J plants of each firm are innovating plants. In this extension, we enrich the model by modeling firms with both innovating and non-innovating plants. Specifically, firms may have heterogeneous efficiencies and make endogenous plant-level innovation decisions along the extensive margin. The set of innovating plants for each firm

is thus endogenous. While solving for the equilibrium involves a difficult combinatorial discrete choice problem (cf. Arkolakis, Eckert, and Shi, 2023) and is outside the scope of this paper, we extend our main sufficient statistic result, Proposition 2, to this extended setting. That is, we derive the sufficient statistic for each location's marginal private and social value of innovation in equilibrium.

Formally, we now define a firm (\hat{J}, φ) by its exogenous set of production plants \hat{J} and efficiency φ . All plants produce, but only a subset of plants endogenously innovate. We assume φ is a scalar drawn from a continuous CDF $F_{\hat{J}}$, and that the firm-wide efficiency φ applies to all plants of the firm.

We specify that the marginal cost of a plant in location n is $\kappa_n(\hat{J}, \varphi) = k_n(\hat{J}, \varphi)^{\frac{1}{1-\sigma}}$, where:

$$k_n(\hat{J}, \varphi) = \varphi \prod_{i \in \hat{J}} \max \left\{ 1, \left(h_i(\hat{J}, \varphi)^\alpha K_i^\delta \right)^{\omega_{ni}} \right\}. \quad (\text{C.2})$$

Let $J(\hat{J}, \varphi) \subseteq \hat{J}$ denote the set of innovating plants of firm (\hat{J}, φ) . The productivity of a plant in location n depends on three components:

1. firm-wide efficiency φ ;
2. the firm's innovation investments $h_i(\hat{J}, \varphi)$ in all plant locations $i \in J(\hat{J}, \varphi)$, including locations $i \neq n$;
3. the location-wide knowledge K_i across all innovating plants in each location $i \in J(\hat{J}, \varphi)$,

$$K_i \equiv \sum_{\hat{J} \ni i} N_{\hat{J}} \int_{\varphi} k_i(\hat{J}, \varphi) \mathbf{1}_{i \in J(\hat{J}, \varphi)} dF_{\hat{J}}(\varphi). \quad (\text{C.3})$$

Note that, in particular, our specification implies that (i) non-innovating plants do not benefit from local spillovers, nor do they contribute to local knowledge: $k_n(\hat{J}, \varphi)$ does not benefit from higher K_i , and $k_i(\hat{J}, \varphi)$ does not contribute to higher K_i , unless $h_i(\hat{J}, \varphi) > 0$; (ii) non-innovating plants benefit from cross-location spillovers: $k_n(\hat{J}, \varphi)$ benefits from higher K_i even if $h_n(\hat{J}, \varphi) = 0$. Both implications are motivated by our empirical evidence (cf. Table 2 and Table 4).

The variable profit of a plant in location n is $\pi_n(\hat{J}, \varphi) = b_n k_n(\hat{J}, \varphi)$, where b_n is the market access in location n defined analogously to (13). Given the innovation subsidy

$t_{nj}(\varphi)$ to plant n of firm (\hat{J}, φ) , the firm solves:

$$\max_{h_n(\hat{J}, \varphi)} \sum_n \left[b_n \varphi \prod_{i \in \hat{J}} \max \left\{ 1, \left(h_i(\hat{J}, \varphi)^\alpha K_i^\delta \right)^{\omega_{ni}} \right\} - c_n h_n(\hat{J}, \varphi) / t_n(\hat{J}, \varphi) \right].$$

In an equilibrium, the private value of a marginal innovation investment in an innovating plant is:

$$\beta_i(\hat{J}, \varphi) \equiv \frac{\partial \sum_{n \in \hat{J}} \pi_n(\hat{J}, \varphi)}{\partial \ln h_i(\hat{J}, \varphi)} = \alpha \sum_{n \in \hat{J}} \pi_n(\hat{J}, \varphi) \omega_{ni} \mathbf{1}_{ni}^{\hat{J}}.$$

Following the same steps as in Lemma 2, we can show that the impact of a marginal innovation investment in an innovating plant $n \in J(\hat{J}, \varphi)$ on social welfare is the same as $\frac{\partial \hat{\mathcal{W}}}{\partial h_n(\hat{J}, \varphi)} |_{\text{holding } \{b_n\} \text{ constant}}$, with $\hat{\mathcal{W}}$ defined as:

$$\begin{aligned} \hat{\mathcal{W}} &\equiv \sum_n \left(\frac{\sigma}{\sigma - 1} \sum_{\hat{J} \ni n} N_{\hat{J}} \int_{\varphi} \left\{ b_n \varphi \prod_{m \in J(\hat{J}, \varphi)} \left(h_m(\hat{J}, \varphi)^\alpha K_m^\delta \right)^{\omega_{nm}} - \sum_{\hat{J} \ni n} N_{\hat{J}} c_n h_n(\hat{J}, \varphi) \right\} dF_{\hat{J}}(\varphi) \right) \\ &= \sum_n \left(\frac{\sigma}{\sigma - 1} \sum_{\hat{J} \ni n} N_{\hat{J}} \int_{\varphi} \left\{ \pi_n(\hat{J}, \varphi) - \sum_{\hat{J} \ni n} N_{\hat{J}} c_n h_n(\hat{J}, \varphi) \right\} dF_{\hat{J}}(\varphi) \right). \end{aligned}$$

Analogous to (22), the social value of a marginal innovation investment in an innovating plant is:

$$\begin{aligned}
& \gamma_i(\hat{J}, \varphi) \\
& \equiv N_j^{-1} \frac{d\hat{\mathcal{W}}}{d \ln h_i(\hat{J}, \varphi)} \\
& = \frac{\sigma}{\sigma-1} \frac{\partial \sum_{n \in \hat{J}} \pi_n(\hat{J}, \varphi)}{\partial \ln h_i(\hat{J}, \varphi)} f_{\hat{J}}(\varphi) \\
& \quad + \frac{\sigma}{\sigma-1} N_j^{-1} \sum_m \frac{\partial \sum_{\hat{J}'} N_{\hat{J}'} \sum_{n \in \hat{J}'} \int_{\varphi'} \pi_n(\hat{J}', \varphi') dF_{\hat{J}'}(\varphi')}{\partial \ln K_m} \frac{d \ln K_m}{d \ln h_i(\hat{J}, \varphi)} \\
& = \frac{\sigma}{\sigma-1} \alpha \sum_{n \in \hat{J}} \pi_n(\hat{J}, \varphi) \omega_{ni} \mathbf{1}_{ni}^{\hat{J}} f_{\hat{J}}(\varphi) \\
& \quad + \frac{\sigma}{\sigma-1} N_j^{-1} \delta \sum_{\hat{J}'} N_{\hat{J}'} \sum_{n \in \hat{J}'} \int_{\varphi'} \sum_{m \in J(\hat{J}', \varphi')} \pi_n(\hat{J}', \varphi') \omega_{nm} \mathbf{1}_{nm}^{\hat{J}'} dF_{\hat{J}'}(\varphi') \frac{d \ln K_m}{d \ln h_i(\hat{J}, \varphi)}. \quad (\text{C.4})
\end{aligned}$$

We now derive the private and social value of innovation at the location level. First, we substitute (C.2) into (C.3):

$$K_n \equiv \sum_{\hat{J} \ni n} N_{\hat{J}} \int_{\varphi} \prod_{i \in \hat{J}} \max \left\{ 1, \left(h_i(\hat{J}, \varphi)^\alpha K_i^\delta \right)^{\omega_{ni}} \right\} \mathbf{1}_{n \in J(\hat{J}, \varphi)} dF_{\hat{J}}(\varphi). \quad (\text{C.5})$$

Let

$$s_n(\hat{J}, \varphi) \equiv \frac{N_{\hat{J}} \varphi \prod_{i \in \hat{J}} \max \left\{ 1, \left(h_i(\hat{J}, \varphi)^\alpha K_i^\delta \right)^{\omega_{ni}} \right\} \mathbf{1}_{n \in J(\hat{J}, \varphi)} f_{\hat{J}}(\varphi)}{\sum_{\hat{J}' \ni n} N_{\hat{J}'} \int_{\varphi'} \prod_{i \in \hat{J}'} \max \left\{ 1, \left(h_i(\hat{J}', \varphi')^\alpha K_i^\delta \right)^{\omega_{ni}} \right\} \mathbf{1}_{n \in J(\hat{J}', \varphi')} dF_{\hat{J}'}(\varphi')}.$$

Differentiating (C.5) with respect to $h_i(\hat{J}, \varphi)$ for a plant with interior innovation investments, we have:

$$\begin{aligned}
\frac{d \ln K_n}{d \ln h_i(\hat{J}, \varphi)} & = s_n(\hat{J}, \varphi) \alpha \omega_{ni} \\
& \quad + \int_{\varphi'} \sum_{\hat{J}'} s_n(\hat{J}', \varphi') \sum_{q \in J(\hat{J}', \varphi')} \delta \omega_{nq} \frac{d \ln K_q}{d \ln h_i(\hat{J}, \varphi)} F_{\hat{J}'}(\varphi').
\end{aligned}$$

Integrating across all innovating plants in location i , we obtain:

$$\begin{aligned}
\chi_{ni} &\equiv \sum_{\hat{j} \ni i} \int_{\varphi} \mathbf{1}_{i \in J(\hat{j}, \varphi)} \frac{d \ln K_n}{d \ln h_i(\hat{j}, \varphi)} dF_{\hat{j}}(\varphi) \\
&= \left(\sum_{\hat{j}} \int_{\varphi} s_n(\hat{j}, \varphi) \mathbf{1}_{i \in J(\hat{j}, \varphi)} dF_{\hat{j}}(\varphi) \right) \times \alpha \omega_{ni} \\
&\quad + \sum_q \left(\sum_{\hat{j}'} \int_{\varphi'} s_n(\hat{j}', \varphi') \mathbf{1}_{q \in J(\hat{j}', \varphi')} dF_{\hat{j}'}(\varphi') \right) \times \delta \omega_{nq} \chi_{qi}.
\end{aligned}$$

Note that $s_{ni} \equiv \sum_{\hat{j}} \int_{\varphi} s_n(\hat{j}, \varphi) \mathbf{1}_{i \in J(\hat{j}, \varphi)} dF_{\hat{j}}(\varphi)$ is the revenue-weighted share of innovating plants in location n that are connected to an innovating plant in location i . Hence:

$$\chi_{ni} = \alpha s_{ni} \omega_{ni} + \delta \sum_q s_{nq} \omega_{nq} \chi_{nq}.$$

In matrix form:

$$\mathbf{X} \equiv [\chi_{ni}] = \alpha (\mathbf{I} - \delta \mathbf{\Omega} \circ \mathbf{S})^{-1} \mathbf{\Omega} \circ \mathbf{S},$$

which is analogous to equation (B.1) in our baseline model.

The location-specific private value of innovation is:

$$\begin{aligned}
\beta_i &\equiv \sum_{\hat{j}} N_{\hat{j}} \int_{\varphi} \beta_i(\hat{j}, \varphi) \mathbf{1}_{i \in J(\hat{j}, \varphi)} dF_{\hat{j}}(\varphi) \\
&= \sum_{\hat{j}} N_{\hat{j}} \int_{\varphi} \alpha \sum_{n \in \hat{j}} \pi_n(\hat{j}, \varphi) \omega_{ni} \mathbf{1}_{i \in J(\hat{j}, \varphi)} dF_{\hat{j}}(\varphi) \\
&= \frac{\alpha}{\sigma} \sum_n \omega_{ni} \left(\sum_{\hat{j} \ni n} N_{\hat{j}} \int_{\varphi} r_n(\hat{j}, \varphi) \mathbf{1}_{i \in J(\hat{j}, \varphi)} dF_{\hat{j}}(\varphi) \right) \\
&= \frac{\alpha}{\sigma} \sum_n \omega_{ni} \hat{s}_{ni} r_n,
\end{aligned}$$

where r_n is the total revenue of all plants in location n , and \hat{s}_{ni} is the revenue-weighted fraction of all plants in location n that are connected to an innovating plant in location i . Note the distinction between \hat{s}_{ni} and s_{ni} : the former refers to all plants in location n ; the

latter refers to all innovating plants in location n . In vector form:

$$\beta' = \frac{\alpha}{\sigma} \mathbf{r}' \left(\Omega \circ \hat{\mathbf{S}} \right).$$

Finally, to derive the location-specific social value of innovation, we integrate (C.4) across all innovating plants in location i :

$$\begin{aligned} \gamma_i &\equiv \sum_{\hat{j}} N_{\hat{j}} \int_{\varphi} \gamma_i \left(\hat{j}, \varphi \right) \mathbf{1}_{i \in J(\hat{j}, \varphi)} dF_{\hat{j}}(\varphi) \\ &= \frac{\sigma}{\sigma-1} \alpha \sum_{\hat{j}} N_{\hat{j}} \int_{\varphi} \left(\sum_{n \in \hat{j}} \pi_n \left(\hat{j}, \varphi \right) \omega_{ni} \mathbf{1}_{ni}^{\hat{j}} f_{\hat{j}}(\varphi) \right) \mathbf{1}_{i \in J(\hat{j}, \varphi)} dF_{\hat{j}}(\varphi) \\ &\quad + \frac{\sigma}{\sigma-1} \delta \left(\sum_{\hat{j}'} N_{\hat{j}'} \sum_{n \in \hat{j}'} \int_{\varphi'} \sum_m \pi_n \left(\hat{j}', \varphi' \right) \omega_{nm} \mathbf{1}_{nm}^{\hat{j}'} dF_{\hat{j}'}(\varphi') \chi_{mi} \right) \\ &= \frac{1}{\sigma-1} \left(\alpha \sum_n r_n \hat{s}_{ni} \omega_{ni} + \delta \sum_n r_n \sum_m s_{nm} \omega_{nm} \chi_{mi} \right). \end{aligned}$$

In vector form:

$$\begin{aligned} \gamma' &= \frac{\alpha}{\sigma-1} \left(\mathbf{r}' \left(\Omega \circ \hat{\mathbf{S}} \right) + \delta \mathbf{r}' \left(\Omega \circ \hat{\mathbf{S}} \right) (\mathbf{I} - \delta \Omega \circ \mathbf{S})^{-1} \Omega \circ \mathbf{S} \right) \\ &= \frac{\alpha}{\sigma-1} \mathbf{r}' \left(\Omega \circ \hat{\mathbf{S}} \right) (\mathbf{I} - \delta \Omega \circ \mathbf{S})^{-1} \\ &= \frac{\sigma}{\sigma-1} \beta' (\mathbf{I} - \delta \Omega \circ \mathbf{S})^{-1}. \end{aligned}$$

We summarize these results in the following proposition.

Proposition 3. *With productivity heterogeneity and endogenous innovation investments along the extensive margin, the location-specific private value of innovation investments is $\beta' = \frac{\alpha}{\sigma} \mathbf{r}' \left(\Omega \circ \hat{\mathbf{S}} \right)$, where $\hat{\mathbf{S}}$ is the matrix whose ni -th entry captures the revenue-weighted share of all plants in location n that are connected to an innovating plant in location i . The corresponding social value of innovation investments is $\gamma' = \frac{\sigma}{\sigma-1} \beta' (\mathbf{I} - \delta \Omega \circ \mathbf{S})^{-1}$.*