

The Effect of High-Tech Clusters on the Productivity of Top Inventors: Comment

Michael Wiebe
Independent*

March 11, 2026

Moretti (2021a,b) (henceforth M21) studies agglomeration effects for innovation, testing whether the size of technology clusters causes patenting. Agglomeration effects are important for understanding technological progress (Kerr and Robert-Nicoud, 2020), and are affected by constraints on housing supply (Duranton and Puga, 2020). Using US data on patents filed between 1971 and 2007 (Zucker and Darby, 2014), M21 presents multiple lines of evidence supporting a causal effect of cluster size on patenting. The baseline results are from linear regressions of patents on cluster size, controlling for an extensive set of fixed effects (including inventor effects). The main finding is an elasticity of patenting with respect to cluster size of 0.0676.

In this comment, I identify two major problems in M21. First, M21 uses an event study to test for selection bias from promising inventors sorting into large clusters. The event is inventors moving across cities, but M21's specification does not use the variation in cluster size generated by moving. I construct an event study following the literature on 'mover' designs, where the change in average cluster size is interacted with event time indicators. I find a null effect, which is consistent with no causal effect of cluster size, but could also be explained by attenuation bias from moves being misclassified in the data.

Second, M21 uses an instrumental variables strategy to control for omitted variable bias from sources like local field-specific subsidies. The instrument is based on changes in the number of inventors in other cities. Due to a coding error, the data is not sorted by city, so the instrument is constructed incorrectly by taking first-differences across cities. When I calculate

*maswiebe@gmail.com. I thank Enrico Moretti for feedback that significantly improved this comment. I am grateful to Enrico Berkes, Abel Brodeur, Matthew Courchene, Rafael Latham-Proença, Matthew Lilley, David Macdonald, Aruni Mitra, David Molitor, Adlai Newson, Paul Novosad, Otis Reid, David Roodman, Jonathan Roth, Kurt Schmidheiny, Jacob Schwartz, Mengying Wei, Toban Wiebe, the editor, and an anonymous referee for helpful feedback. I thank Open Philanthropy for financial support.

the first-difference within cities, the estimates are nonsignificant.

I Event study

M21 Figure 6 uses an event study to address worries about sorting, where ‘rising star’ inventors (with increasing patent counts) in small clusters are systematically hired by employers in large clusters.¹ If these promising inventors select into larger clusters, then the positive elasticity between cluster size and patenting would represent selection bias rather than a causal effect of cluster size. Note that the baseline regressions (M21 Table 3) control for inventor fixed effects, so the threat to identification comes only from inventors whose inventiveness is increasing over time. If sorting does occur, we should observe patenting rising in the years before an inventor moves to a larger city.

To test for sorting, M21 performs an event study using variation in cluster size from inventors who move across cities exactly once.² That is, ‘stayers’ who never move are excluded, so the event study does not have a never-treated group. To generate a treatment-control comparison, M21 uses average cluster size before and after the move as a continuous treatment variable. Specifically, M21 interacts pre-move average cluster size with the pre-move event-time indicators, and post-move average cluster size with the post-move event-time indicators.³ M21 does not show a regression equation. My reconstruction from the text⁴ is

$$\begin{aligned} \ln y_{ijkct} = & \sum_{s=-5}^{-1} \beta_s \text{Size}_{-i}^{pre} \times \mathbb{1}\{t = s\} + \sum_{s=0}^5 \beta_s \text{Size}_{-i}^{post} \times \mathbb{1}\{t = s\} \\ & + d_{cf} + d_{ck} + d_{f\tau} + d_{k\tau} + d_{c\tau} + d_i + d_j + \varepsilon_{ijkct}. \end{aligned} \quad (1)$$

Here y is the number of patents by inventor i in firm j , research field f , technology class k , city c , and event year t . Event time is measured as the number of years since the inventor’s move. Fixed effects include city-field, city-class, field-year, class-year, city-year, inventor, and firm; note that the time fixed effects are based on calendar year τ . Size_{-i}^{pre} is average log cluster

¹A cluster is defined as a city \times research-field pair; there are 179 cities and 5 fields. Cluster size is the number of inventors excluding the focal inventor, and inventors are counted if they file a patent in that year. The main analysis focuses on the top 10% of inventors by lifetime patents.

²M21 identifies a mover as an inventor who patents in different cities in consecutive years. The inventor-year panel is otherwise unbalanced, with time gaps allowed between observations.

³M21 states: “the timing relative to an inventor’s move is interacted with the corresponding cluster size. As treatment variable, I use average cluster size before the move and after the move. [...] this specification is an event study, based on a “pure” movers design solely exploiting variation in cluster size induced by moves.” (p.3352)

⁴Note that this specification does not exactly match the M21 code, because the M21 code disagrees with the M21 text. See the discussion in footnote 8.

size calculated over the 5-year period before the move, while $\text{Size}_{-i}^{\text{post}}$ is calculated over the 5-year period after the move.⁵ Inventor i is excluded when calculating cluster size. Event time indicators are denoted by $\mathbb{1}\{s = t\}$. For a given event year t , the sums evaluate to a single term (with $s = t$), so one coefficient β_t is estimated per event year. Note that the coefficients are estimated relative to the omitted period, which is event years outside of $[-5, 5]$. I reproduce M21’s Figure 6 in Figure 1a.⁶

However, Equation 1 is not a proper event study specification.⁷ ‘Mover’ studies use the change in an individual’s average environment before and after the move as the treatment variable (Finkelstein et al., 2016; Molitor, 2018; Cantoni and Pons, 2022; Badinski et al., 2023). This constant treatment variable is then interacted with event time indicators. In contrast, M21 uses the pre- and post-move averages of cluster size ($\text{Size}_{-i}^{\text{pre}}$ and $\text{Size}_{-i}^{\text{post}}$), but interacts them separately with the pre- and post-move event time indicators. This does not exploit the variation in cluster size generated by moving across cities. In fact, this specification would be identified even if moving induced no change in cluster size (that is, $\text{Size}_{-i}^{\text{pre}} = \text{Size}_{-i}^{\text{post}}$).

Instead, Equation 1 effectively uses the same contemporaneous variation as the baseline regressions, but with cluster size measured as an average over the pre- and post-move periods. That is, β_t is estimated by regressing an inventor’s patents on their (averaged) cluster size in the same year. Since the treatment variation does not come from moving, M21 Figure 6 does not test for sorting. By interacting $\text{Size}_{-i}^{\text{pre}}$ and $\text{Size}_{-i}^{\text{post}}$ with event time indicators, this specification actually tests for heterogeneity in the cross-sectional relationship between (averaged) cluster size and patenting by event time. Hence, if the relationship between cluster size and patenting is constant over event time, the β_t coefficients should all be equal, as they estimate the same constant elasticity.⁸

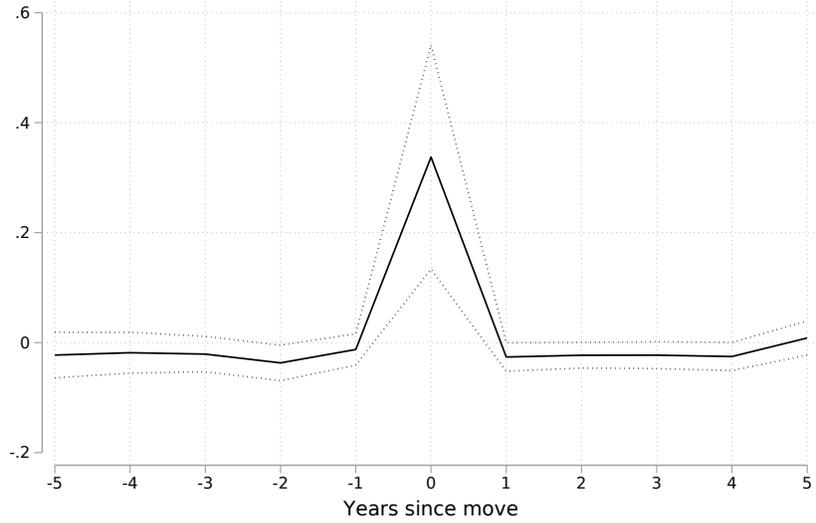
⁵M21 excludes $t = 0$ when calculating post-move average cluster size, despite being the inventor’s first year in the new city.

⁶There is a large immediate effect of cluster size in $t = 0$ (the year of the move), which dissipates entirely after one year. Note that the bottom panel of M21 Figure 6 plots the cumulative effect from the event study, which M21 interprets as showing that “most of the effect persists five years after the move.” (p.3355) However, the event study coefficients in the top panel of M21 Figure 6 already capture the cumulative treatment effect, so the bottom panel calculates the cumulative response function of a cumulative effect. M21 incorrectly treats the M21 Figure 6 top panel estimates as incremental estimates from a distributed lag model, which do need to be aggregated; see Schmidheiny and Sieglöcher (2023). Note that M21 Figure 5 uses a distributed lag model, while M21 Figure 6 uses an event study.

⁷I thank Enrico Moretti and an anonymous referee for discussion of this point.

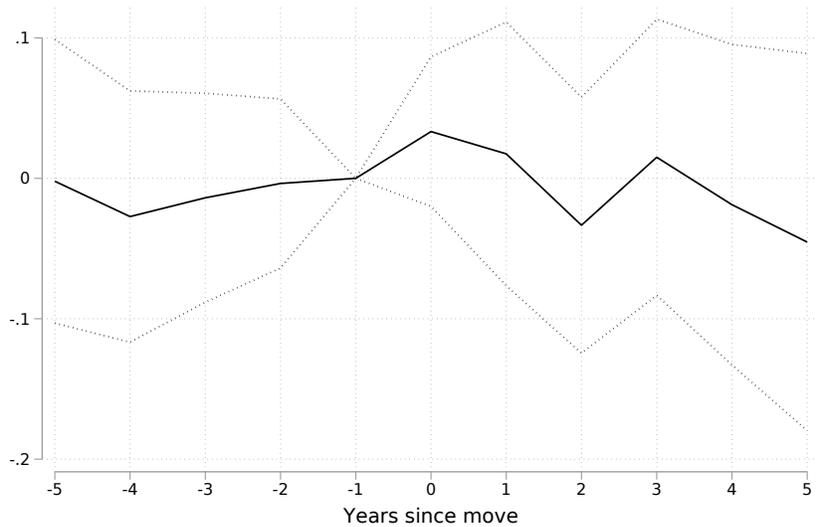
⁸This pattern is not what we see in M21 Figure 6, with a large effect only in $t = 0$. The M21 Figure 6 results are driven entirely by one coding error. The M21 code does not generate an interaction of post-move cluster size with a $t = 0$ indicator ($\text{Size}_{-i}^{\text{post}} \times \mathbb{1}\{s = 0\}$), but instead uses the time-varying cluster size variable that is used in the baseline regressions ($\ln \text{Size}_{-i\text{fct}}$). Hence, β_0 is estimated using data from all event-years, instead of capturing the effect in $t = 0$. I correct the code by interacting post-move cluster size with the $t = 0$ indicator. Figure A1 shows the original and corrected versions of Figure 6. The large effect in $t = 0$ disappears when using

Figure 1: Mover event study



(a) Reproduction of M21 Figure 6 (top panel)

Note: The solid line plots event study coefficients from Equation 1. Dotted lines are 95% confidence intervals. $N = 18,443$. Fixed effects: city \times field, city \times class, field \times year, class \times year, inventor, city \times year, and firm. Standard errors are clustered by city \times research field. My figure does not exactly match M21 Figure 6, because the M21 cleaning code is unreproducible; see Appendix A.



(b) Mover event study specification

Note: The solid line plots event study coefficients from Equation 2. Dotted lines are 95% confidence intervals. $N = 22,155$. Fixed effects: origin city \times field, destination city \times field, origin city \times class, destination city \times class, field \times year, class \times year, inventor, origin city \times year, destination city \times year, and firm. Standard errors are clustered by city \times research field.

I follow the literature on mover event studies and construct an event study specification using $\Delta\text{Size}_{-i} = \text{Size}_{-i}^{\text{post}} - \text{Size}_{-i}^{\text{pre}}$ as the treatment variable, which measures the change in average cluster size associated with moving across cities. $\text{Size}_{-i}^{\text{post}}$ is an inventor’s average log cluster size in the five years after the move (as well as the move year), and $\text{Size}_{-i}^{\text{pre}}$ is average log cluster size in the five years before the move. While cluster size varies over time as other inventors enter and exit a cluster, ΔSize_{-i} is a constant and isolates the variation in size generated specifically by moving. Figure A2 shows that moving across cities produces substantial variation in cluster size, with a roughly equal number of inventors moving to larger and smaller cities. To test how the change in size affects inventor output over time, I interact ΔSize_{-i} with event time indicators. I omit the interaction for event time $t = -1$ as the reference period. Since the sample does not include stayers, the dynamic treatment effects may not be separately identified from the time fixed effects (Schmidheiny and Siegloch, 2023). To achieve identification, I include observations corresponding to event years outside of $[-5, 5]$ and assume that the treatment effects are constant before event year -5 and after event year 5 .

Following Equation (1) in Molitor (2018), the mover event study specification is

$$\begin{aligned} \ln y_{ijkct} = & \sum_{s=-5, s \neq -1}^5 (\alpha_s D_{it}^s + \beta_s D_{it}^s \times \Delta\text{Size}_{-i}) \\ & + d_{cf} + d_{ck} + d_{f\tau} + d_{k\tau} + d_{c\tau} + d_i + d_j + \varepsilon_{ijkct}. \end{aligned} \quad (2)$$

where for event time t we have $D_{it}^{-5} = \mathbb{1}\{t \leq -5\}$ when $s = -5$; $D_{it}^s = \mathbb{1}\{s = t\}$ for $-5 < s < 5$; and $D_{it}^5 = \mathbb{1}\{t \geq 5\}$ when $s = 5$. Here D_{it}^s is the binned treatment adoption indicator (Schmidheiny and Siegloch, 2023), where the binned endpoints enforce constant treatment effects for event years outside of $[-5, 5]$.⁹ As before, for a given event year t , the sum evaluates to a single term (with $s = t$).¹⁰ Here α_t captures generic effects of moving, while β_t identifies the effect of a change in cluster size. The time fixed effects are based on calendar year τ .

I include the same set of fixed effects as in Equation 1, except for city fixed effects, which can generate collinearity in the context of inventors moving from an origin city to a destination

the corrected code.

⁹Without binning the endpoints, the equation simplifies to $\ln y_{it} = \sum_{s=-5}^5 (\alpha_s \mathbb{1}\{s = t\} + \beta_s \mathbb{1}\{s = t\} \times \Delta\text{Size}_{-i}) + \text{FEs} + \varepsilon_{it}$. When observations outside of event years $[-5, 5]$ are included, this specification assumes the treatment effect in those event years is zero.

¹⁰Note that this ‘expanded-sum’ specification is equivalent to alternative event study specifications used in the literature (for example, Equation (6) in Finkelstein et al. (2016), Equation (6) in Cantoni and Pons (2022), and Equation (2) in Badinski et al. (2023)). These versions are evaluated at a given event year, so the sum collapses to a single term with $\mathbb{1}\{s = t\} = 1$. Following their notation, Equation 2 is equivalent to $\ln y_{it} = \theta_{r(i,t)} \Delta\text{Size}_{-i} + \text{FEs} + \varepsilon_{it}$, where t is calendar time, $r(i, t)$ is relative event time, and the event time indicators are included in the fixed effects.

city. In particular, when a city is observed only as an origin or only as a destination, then the post-move event time indicators are collinear with the city effect.¹¹ To address this problem, I replace the city fixed effects with origin- and destination-city effects; note that city varies within inventor, while origin and destination are constant.

The estimates of β_t are shown in Figure 1b. There is no statistically significant pretrend indicating sorting (for $t < 0$), but also no statistically significant treatment effect (for $t \geq 0$). Consistent with the assumption of constant treatment effects outside of event years $[-5, 5]$, the estimates are stable around 0 at the endpoints. Note that the baseline elasticity (using linear regression with fixed effects) estimated on the mover event study sample is large and significant: 0.3145 (0.0953). For comparison, the M21 baseline elasticity (Table 3, Column 8) is 0.0676. Hence, while there is a large effect in the mover sample when using all observed variation in cluster size, we do not detect an effect when using only variation from moving.

To show how controlling for city matters, Figure A3 plots the event study results when controlling only for origin-city fixed effects, and when omitting city-level fixed effects altogether. The post-move coefficients are small and positive, but remain nonsignificant in all but one case. I also run specifications including stayers as a never-treated group, to provide additional variation for separately identifying the dynamic treatment effects from the year fixed effects. Figures A4 and A5 plot the results. The estimates are similar in magnitude, with at most one post-move coefficient attaining significance.

I test whether the null effect is driven by heterogeneity between inventors moving “up” to a larger city and inventors moving “down” to a smaller city. Specifically, I allow for separate effects by moving-up and moving-down by replacing the sum in Equation 2 with $\sum_{s=-5, s \neq -1}^5 (\beta_s^{up} D_{it}^s \times \Delta \text{Size}_{-i} \times \mathbb{1}\{\Delta \text{Size}_{-i} > 0\} + \beta_s^{down} D_{it}^s \times \Delta \text{Size}_{-i} \times \mathbb{1}\{\Delta \text{Size}_{-i} < 0\})$. This specification captures the absolute effect for each subgroup, instead of the differential effect between subgroups. I plot the event study coefficients in Figure A6, with alternative versions in Figure A7 (controlling only for origin-city fixed effects) and Figure A8 (omitting all city-level fixed effects). The effect is slightly larger when moving to a larger city, though the difference is nonsignificant. When fewer fixed effects are included, some preperiod coefficients become significant, suggesting a violation of the event study design.¹²

To understand the treatment variation generated by moving, I investigate how patents vary with event time in the raw data. Figure A12 plots average log patents by event time, separately by the move-up and move-down subsamples. If a treatment effect of cluster size is visible in the

¹¹To see the collinearity, consider the difference-in-differences specification with *Post* indicating time periods after the move. Here, origin-only cities have *Post*=0 and destination-only cities have *Post*=1 for all observations, so city fixed effects are collinear with *Post*.

¹²I find similar nonsignificant results when repeating this exercise using $\mathbb{1}\{\Delta \text{Size}_{-i} > 0\}$ and $\mathbb{1}\{\Delta \text{Size}_{-i} < 0\}$ as binary treatment variables (see Figures A9–A11); note that in this case, we expect a positive effect of moving-up and a negative effect of moving-down.

raw data, we should observe opposite patterns: an increase in patents for inventors moving-up and a decrease for inventors moving-down. This is not what we see, with both groups following the same pattern: an increase in patents in $t = 1$, followed by a decline until $t = 5$. While there is an opposite change of the expected sign from $t = -1$ to $t = 0$, the magnitude is small relative to pre-move fluctuations.

If the M21 elasticity was driven by selection bias, then the event study should show an upward-sloping pre-trend and positive post-move coefficients. Neither pattern appears in Figure 1b. This null effect is potentially explained by moves being misclassified in the data. The M21 code creates a unique identifier based on inventor names, but the patent data in Zucker and Darby (2014) does not distinguish between different inventors who share the same name. Hence, what appears to be an inventor moving across cities could simply be two inventors (with the same name) living in different cities.

For example, John P. Hansen is observed patenting in Austin, Texas from 1993-2002 and in Wadsworth, Ohio from 1997-2003.¹³ M21 assigns these observations the same inventor identifier, and transforms the patent data into an inventor-year panel by assigning the inventor’s modal city, which is Austin from 1993-2002 and Wadsworth in 2003. This is coded as a move from Austin to Wadsworth in 2003. However, if these are actually different inventors, each named John P. Hansen, then no move occurred.¹⁴ Without checking manually, it is unclear how many moves are misclassified. Since misclassified moves generate attenuation bias, the lack of a unique person-level identifier could explain the null result in Figure 1b.¹⁵¹⁶ Another explanation is that moving across clusters has no effect on patenting, and the M21 elasticity is driven by omitted variable bias, which I discuss below.

¹³The patents in Wadsworth are all assigned to The J.M. Smucker Company, a food and beverage manufacturer. The patents in Austin are mostly in computer science, and are assigned to Advanced Micro Devices, Motorola, and Freescale Semiconductor.

¹⁴Moreover, as I discuss in Appendix A.2, the Zucker and Darby (2014) dataset does not have a unique mapping from cities to numeric city identifiers. This also generates misclassified moves. For example, Abbas Rafii is observed in Los Altos in 1997-1998, but is coded as moving in 1998. This occurs because the city identifier changes from 097 to 146 (despite no change in the recorded city) and because moves are defined using the numeric identifier.

¹⁵A non-unique inventor identifier also undermines the use of inventor fixed effects in the main regressions.

¹⁶Since inventors move in different years, note that this is a staggered adoption design where a two-way fixed-effects estimator may be biased when treatment effects are heterogeneous over time (Goodman-Bacon, 2021). But since the treatment is continuous, unbiased estimators like Sun and Abraham (2021) and Callaway and Sant’Anna (2021) that are defined for binary treatments are not applicable.

II Instrumental variables estimates

M21 Table 5 uses an instrumental variables (IV) strategy to address worries about omitted variable bias, where cluster size is correlated with unobserved time-varying shocks at the city-field level. For example, a city subsidizing biotech firms could increase both biotech patents and the size of the biotech cluster.¹⁷ The idea for the instrument is to use variation in the number of inventors in other cities who are employed by firms that are also active in the focal inventor’s city.

To illustrate with an example, for the focal inventor i in the field of computer science at Google in San Francisco, we instrument for cluster size using variation in the number of computer science inventors at Microsoft in Seattle, where Microsoft also has inventors in San Francisco. Intuitively, the number of inventors working at Microsoft in Seattle will be correlated with the number of computer science inventors in San Francisco, since Microsoft also has a presence in San Francisco, allowing shocks to Microsoft to propagate to San Francisco. The exclusion restriction is that, controlling for fixed effects, the number of computer science inventors at Microsoft in Seattle is not directly correlated with computer science patents at Google in San Francisco. While this example uses two firms, the instrument uses all firms active in the focal inventor’s city, excluding the inventor’s own firm.

In contrast to the baseline regressions (M21 Table 3), M21 defines the instrument in first-differences. Let $\Delta N_{jfc(-c)t} = N_{jfc(-c)t} - N_{jfc(-c)(t-1)}$ be the first-difference over time in the number of inventors at firm j in research field f in all cities excluding c , and let ΔN_{ft} be the national change by field. With $D_{jfc(t-1)}$ as an indicator for firm j employing at least one inventor in city c and field f in year $t - 1$ (so that the first-difference is well-defined), the IV is

$$IV_{jfc} = \sum_{s \neq j} D_{sfc(t-1)} \frac{\Delta N_{sf(-c)t}}{\Delta N_{ft}}, \quad (3)$$

where the sum is taken across all firms excluding j . The regression equation is

$$\Delta \ln y_{ijfkct} = \alpha \Delta \ln(\text{Size})_{-ifct} + d_j + d_{ft} + d_{kt} + \varepsilon_{ijfkct}, \quad (4)$$

where y is the number of patents, Size is cluster size (excluding the focal inventor i), and k is research class. Fixed effects include firm, field-year, and class-year.¹⁸

The IV results are affected by a coding error. When manually calculating the first-differenced term $\Delta N_{jfc(-c)t}$, M21 sorts the data by firm, field, and year, but not by city. This means that

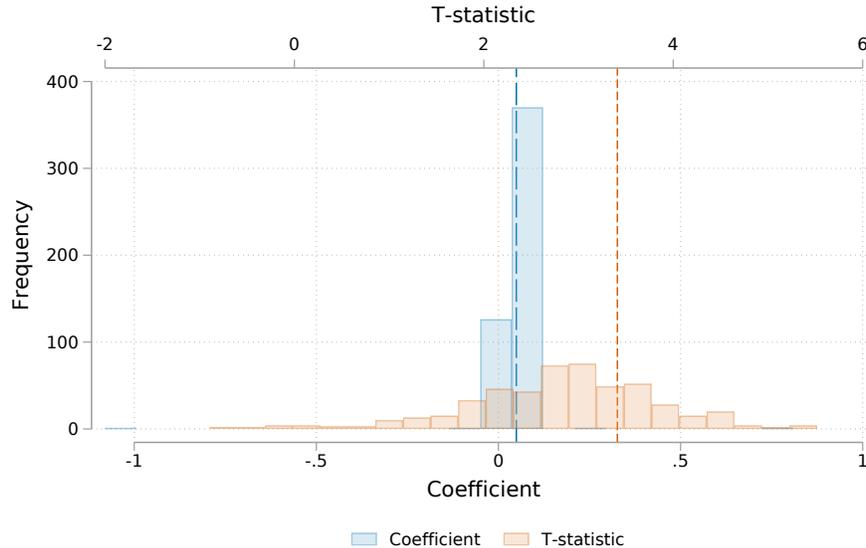
¹⁷Recall that cluster size for a focal inventor is measured as the number of other inventors filing patents that year in the same research field and city.

¹⁸M21’s Equation 3 includes redundant fixed effects, e.g., separate field and year effects when the field-year effect is also included. For clarity, I report only the nonredundant terms.

the first-difference is taken across different cities, instead of within the focal city. Hence, the code generates an incorrect instrument that does not match the definition in the text. I correct the code to sort by city and take first-differences within city.

Moreover, this coding error causes the M21 code to generate unreproducible results. This is because firm-field-year is not a unique sorting order, as there are observations in different cities with the same values for firm-field-year. Since Stata’s `sort` command pseudo-randomly orders the data before sorting, each run of the code results in a different sorting by city within tied values of firm-field-year. Hence, the first-difference is calculated differently each time. In Figure 2, I show the variance in the original results from 500 iterations of generating the instrument and running the same regression. The original M21 estimate is close to the estimate from an average run of the code, while the original t-statistic is slightly larger.

Figure 2: Distribution of coefficients and t-statistics from 2SLS regression



Note: Estimation results from the 2SLS regression in Table 5, Column 6. The distribution is constructed from 500 iterations of generating the instrument and running the regression. The original code sorts by firm-field-year but not city, which leads to a random ordering by city within tied values of firm-field-year. The first-difference in the numerator of the instrument is then (incorrectly) taken across cities, which causes the instrument to be different for each run of the code. The original estimate is 0.0491 (blue long-dashed line) with a t-statistic of 3.41 (orange short-dashed line).

The original results are reproduced in Panels A and B of Table 1, and the results using the corrected instrument are in Panels C and D. Since M21’s results are not reproducible, my reproductions are similar but not identical to the original estimates in M21 Table 5. The results using the corrected instrument are markedly different, however. The 2SLS estimates are now unstable and nonsignificant, possibly due to the first stage being much smaller with an F-statistic around 7. This is not evidence against agglomeration effects, but it does mean

that M21 fails to provide evidence against confounding from city-field-year shocks, such as local field-specific subsidies that affect both patenting and cluster size.

Table 1: Reproduction and correction of M21 Table 5 IV estimates

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: 2SLS (original)						
$\Delta \log \text{ size}$	0.0494 (0.0222)	0.0698 (0.0263)	0.0553 (0.0250)	0.0571 (0.0176)	0.0503 (0.0163)	0.0498 (0.0147)
Observations	419665	419665	419634	405305	405305	404148
Panel B: First stage (original)						
IV	0.8870 (0.1758)	0.8424 (0.1782)	0.8581 (0.1769)	1.1301 (0.2234)	1.1813 (0.1935)	1.1951 (0.1910)
Observations	419665	419665	419634	405305	405305	404148
F-statistic	25.45	22.33	23.53	25.59	37.27	39.16
Panel C: 2SLS (corrected)						
$\Delta \log \text{ size}$	-0.0303 (0.2205)	-0.0118 (0.2166)	0.0002 (0.2045)	-0.0458 (0.2763)	0.1181 (0.2667)	0.1376 (0.2451)
Observations	419665	419665	419634	405305	405305	404148
Panel D: First stage (corrected)						
IV	0.0272 (0.0126)	0.0275 (0.0131)	0.0274 (0.0127)	0.0226 (0.0105)	0.0278 (0.0105)	0.0289 (0.0109)
Observations	419665	419665	419634	405305	405305	404148
F-statistic	4.64	4.45	4.62	4.66	7.06	6.98
Year	Yes	Yes	Yes	Yes		
Field		Yes	Yes	Yes		
Class			Yes	Yes	Yes	
Firm				Yes	Yes	Yes
Field \times year					Yes	Yes
Class \times year						Yes

Standard errors in parentheses

Note: The dependent variable is the change in the log number of patents in a year (2SLS) and the change in log cluster size (first stage). Panels A and B calculate the instrument by incorrectly taking first-differences across cities, while Panels C and D take first-differences within cities. Standard errors are clustered by city. My estimates in Panels A and B do not exactly match the original M21 Table 5, because the M21 code to construct the instrument is unreproducible. Moreover, the M21 cleaning code is unreproducible; see Appendix A.

III Additional coding and interpretation issues

In the Appendix, I document seven other issues. First, the positive effect of cluster size on patent quality in M21 Table 6 is due to using $\log(y + 0.00001)$ to calculate log citations; the estimate when using $\log(y + 1)$ or Poisson regression is negative and significant. Second, the estimates of heterogeneous elasticities by cluster size in M21 Table 8 have a coding error; the

corrected estimates in fact have substantial heterogeneity. Third, a coding error in M21 Figure 5 means that leads and lags are defined incorrectly; the corrected estimate is nonsignificant. Fourth, the estimates by cluster quality in M21 Table A.8 use a definition of cluster size that differs from the text; the corrected estimates are about half as large. Fifth, M21 Table A.7 varies the time unit of the data and interprets the results as capturing an extensive margin effect; this interpretation is incorrect, since no zeros are observed and no extensive margin effect can be estimated. Sixth, M21 Table A.8 controls for the number of coauthors on a patent; this is conceptually flawed, since the dependent variable is already measured as patents per coauthor. Seventh, M21 Table A.6 has a coding error that understates the effect of interpolating missing observations.

IV Conclusion

I reanalyze Moretti (2021b) using the original data and code, and find two major problems that affect the elasticity from the baseline regressions. First, the event study does not use variation from inventors moving across cities; I use a mover event study design and find a null result. Second, the instrumental variable is calculated incorrectly by taking first-differences across cities. I correct the coding error in the instrument and find a null result. Hence, the positive elasticity of patenting with respect to cluster size could be biased upwards due to sorting and omitted variable bias.

References

- Badinski, I., A. Finkelstein, M. Gentzkow, and P. Hull (2023). Geographic variation in healthcare utilization: The role of physicians. Working paper.
- Callaway, B. and P. H. Sant’Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics* 225(2), 200–230.
- Cantoni, E. and V. Pons (2022, April). Does context outweigh individual characteristics in driving voting behavior? evidence from relocations within the united states. *American Economic Review* 112(4), 1226–72.
- Duranton, G. and D. Puga (2020). The economics of urban density. *Journal of Economic Perspectives* 34(3), 3–26.
- Finkelstein, A., M. Gentzkow, and H. Williams (2016, 07). Sources of geographic variation in health care: Evidence from patient migration. *The Quarterly Journal of Economics* 131(4), 1681–1726.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics* 225(2), 254–277. Themed Issue: Treatment Effect 1.
- Kerr, W. R. and F. Robert-Nicoud (2020). Tech clusters. *Journal of Economic Perspectives* 34(3), 50–76.
- Molitor, D. (2018, February). The evolution of physician practice styles: Evidence from cardiologist migration. *American Economic Journal: Economic Policy* 10(1), 326–56.
- Moretti, E. (2021a). Data and code for: The effect of high-tech clusters on the productivity of top inventors. <https://doi.org/10.3886/E140662V1>. American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor].
- Moretti, E. (2021b). The effect of high-tech clusters on the productivity of top inventors. *American Economic Review* 111(10), 3328–75.
- Schmidheiny, K. and S. Siegloch (2023). On event studies and distributed-lags in two-way fixed effects models: Identification, equivalence, and generalization. *Journal of Applied Econometrics* 38(5), 695–713.
- Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics* 225(2), 175–199.

Zucker, L. G. and M. R. Darby (2014). Connecting outcome measures in entrepreneurship technology and science (COMETS) database, public-use version comets 2.0. <https://kauffman.org/entrepreneurship/research/comets/>.

Supplemental Appendix

The Effect of High-Tech Clusters on the Productivity of Top Inventors: Comment

Michael Wiebe

A Computational reproducibility

The cleaning code in Moretti (2021a,b) (henceforth M21) is unreproducible and generates a slightly different dataset for each run of the code. In particular, it uses many-to-many merges with a nonunique sort order. For example, when merging inventors with patent assignees (firms), there are multiple firms per patent, so the patent identifier does not uniquely identify observations. Stata’s ‘merge’ command sorts the data, and ‘sort’ pseudo-randomly orders the data; hence, different runs of ‘merge’ produce a different sorting within tied values of the patent identifier. For coauthored patents with multiple firms, the many-to-many merge matches a different firm to different coauthors on separate runs of the code. The same problem arises when there are multiple research fields per patent. Hence, the code generates a different dataset each time it is run.¹⁹

I change the code to implement a many-to-one merge. I select the first-listed assignee as a patent’s unique assignee; the first-listed assignee is more likely to be a firm rather than an individual, which enables more inventors to share a firm fixed effect. The COMETS data includes numeric weights for each listed research field. I select the highest-weighted field; to break ties, I choose randomly. I follow the M21 code in selecting the first-listed class. With only one assignee, field, and class per patent, a reproducible many-to-one merge is achieved. Note that this produces a different sample than M21, so the results are slightly different. For example, the baseline elasticity in M21 is 0.0676; with my sample, it is 0.0588.

A.1 Coding issues

M21’s code uses the Stata package ‘reghdfe’ with the option ‘keepsingletons’ to not drop singleton groups. The package explicitly warns the user against including singletons, due to biased standard errors.

The cleaning code removes suffixes from inventor names. This creates cases where different inventors are assigned the same name and identifier. For example, patent 3986449 was filed by Paul H. Hamisch, Jr. and Paul H. Hamisch, Sr.; after string cleaning, both are assigned the name “Paul H. Hamisch”.

¹⁹This explains why M21 reports different baseline elasticities throughout the paper: 0.0676 (Table 3 Column 8 and Table A.7 Column 5), 0.0669 (Table 9 Column 3), 0.0662 (in text, p.3369).

M21 assigns multi-cluster inventors (who patent in different clusters in the same year) to their modal cluster. The code calculates the modal field and modal city independently, which can assign inventors to non-existent clusters. For example, for a cluster (city, field), if an inventor patents in clusters (A,1) and (B,2) in the same year, the code could assign them to (A,2) or (B,1), even if those clusters never occur in the data.

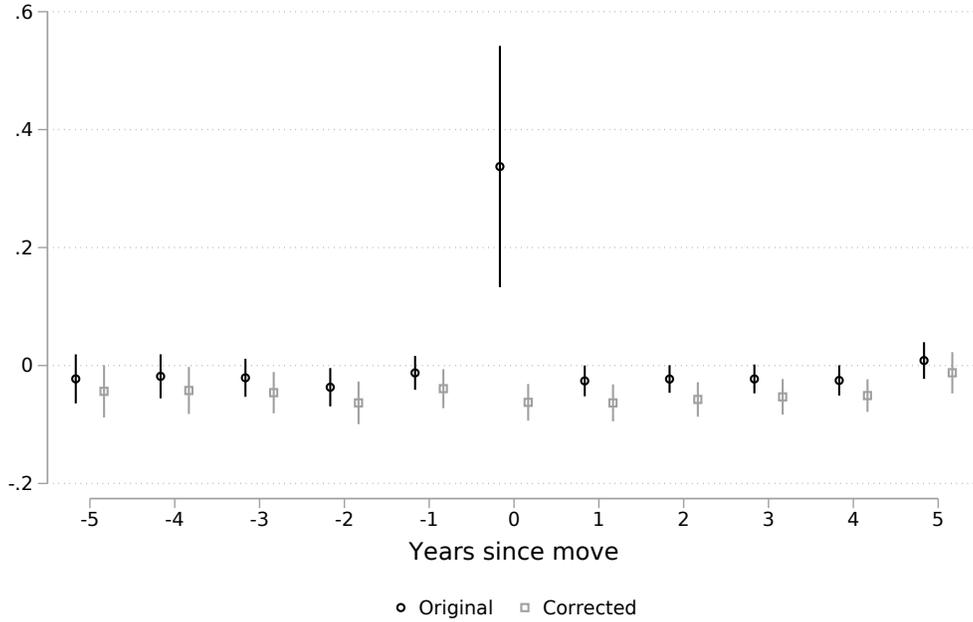
A.2 Data issues

The COMETS dataset appears to not have a unique mapping from cities to numeric BEA codes. For example, Fremont, California is associated with BEA codes 097, 140, and 146 with frequencies 23%, 34%, and 43%. This seems to contradict the description of the BEA variable: “Each county in the U.S. is assigned to a unique BEA area, with multiple counties contained within each BEA area.” (p.291, Zucker et al. (2014)).

On p.3334, M21’s description of the COMETS data is inaccurate. M21 mentions “the patent assignee/owner” and states that “patents are assigned to one of five main “research fields” and 579 “technology classes.””, implying there is a one-to-one mapping between patents and assignees, fields, and classes. This is incorrect: there can be multiple assignees, fields, and classes per patent. M21’s code uses many-to-many merges to handle multiple assignees and fields. It keeps the first-listed class, which enables a many-to-one merge; note that this modelling choice is not defended.

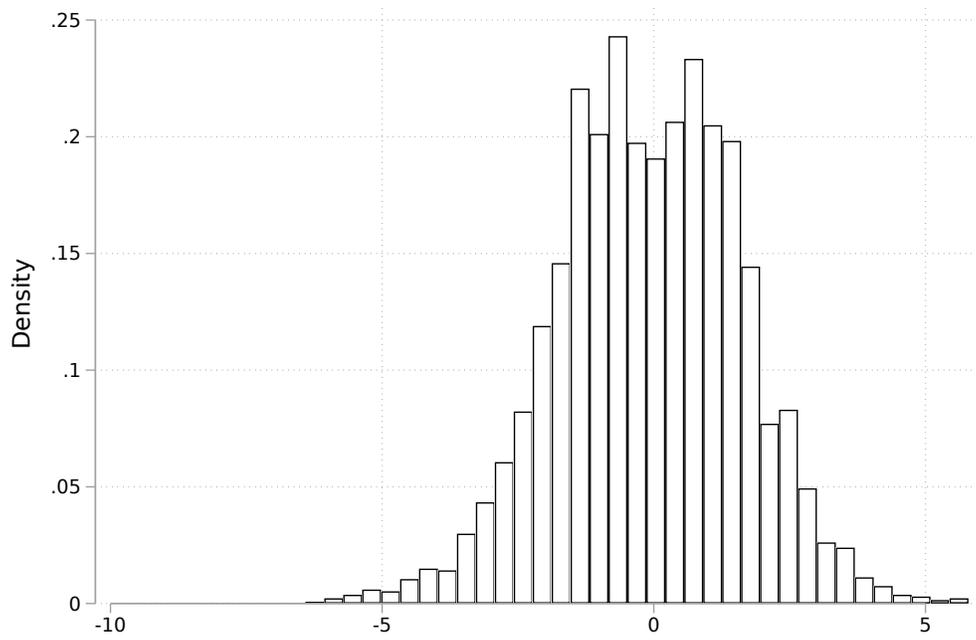
B Event study

Figure A1: Reproduction and correction of M21 Figure 6 (top panel)



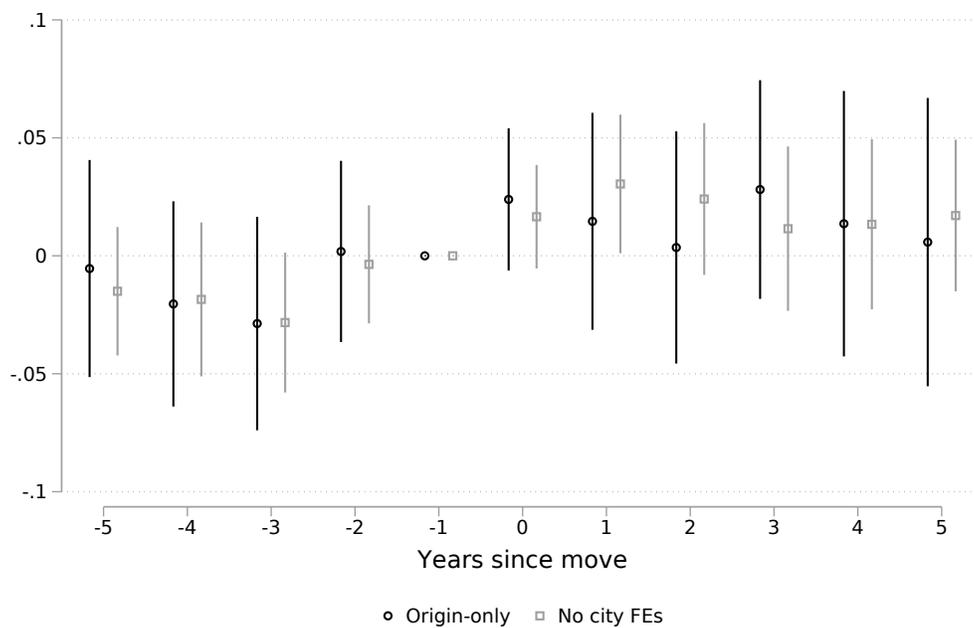
Note: Coefficients and 95% confidence intervals from Equation 1. *Original* follows the M21 code and estimates β_0 using time-varying $Size_{-ifc}$ and without interacting with $\mathbb{1}\{s = 0\}$. *Corrected* estimates β_0 using $Size_{-ifc}^{post} \times \mathbb{1}\{s = 0\}$; following the M21 code, $Size_{-ifc}^{post}$ is calculated over event years 1 to 5, excluding the year of the move ($t = 0$). $N = 18,443$ in *Original*, $N = 18,444$ in *Corrected*. Standard errors are clustered by city \times research field. M21 claims that the event study includes only “inventors who are in the data for 11 consecutive years” (p.3354). However, the code allows time gaps of any length between observations. My figure does not exactly match M21 Figure 6, because the M21 cleaning code is unreproducible; see Appendix A.

Figure A2: Distribution of changes in cluster size across move



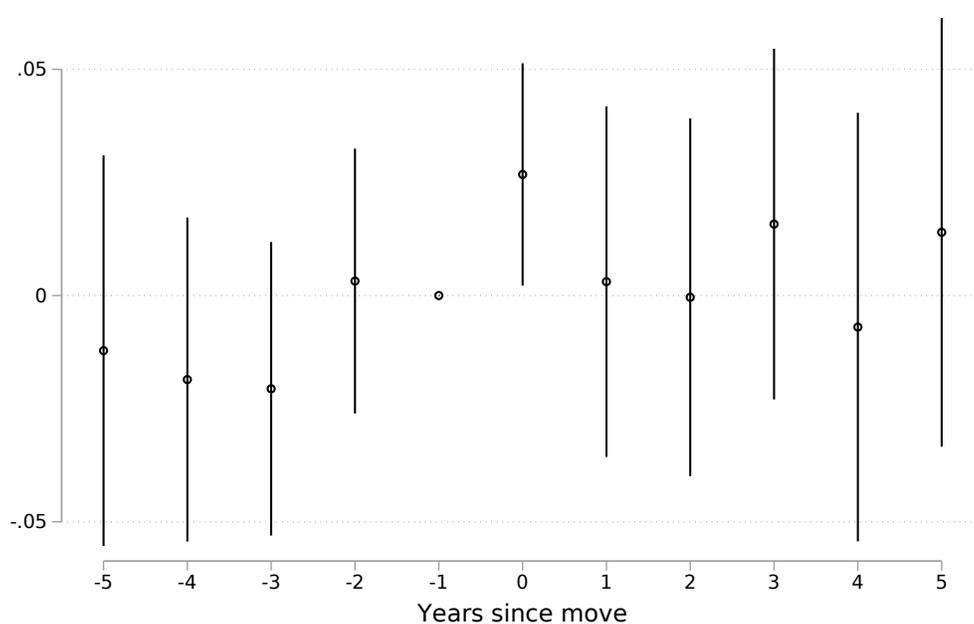
Note: Histogram of changes in cluster size among inventors who move across cities. The change for inventor i is $\Delta\text{Size}_{-i} = \text{Size}_{-i}^{\text{post}} - \text{Size}_{-i}^{\text{pre}}$, where $\text{Size}_{-i}^{\text{post}}$ is average log cluster size over the 5 years after the move and $\text{Size}_{-i}^{\text{pre}}$ is average log cluster size over the 5 years before the move. Inventor i is excluded when calculating cluster size, which is the number of inventors filing a patent. A cluster is a city-research field pair. $N = 3,834$ inventors. Average log cluster size in the mover sample is -3.40.

Figure A3: Mover event study: alternative city controls



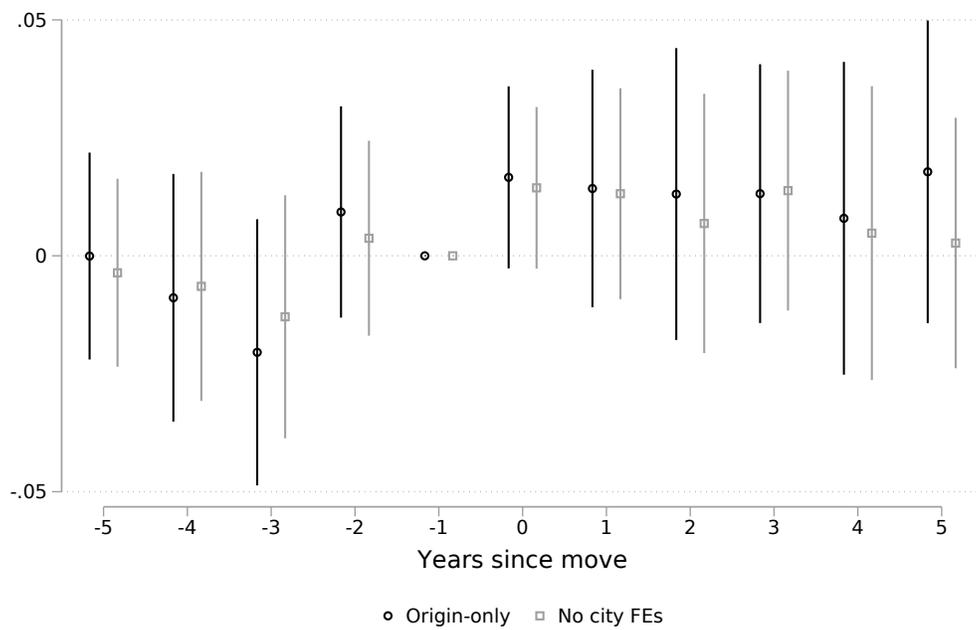
Note: Event study coefficients from Equation 2. Fixed effects: field \times year, class \times year, inventor, and firm. *Origin-only* also includes fixed effects for origin city \times field, origin city \times class, and origin city \times year; $N = 26,486$. *No city FEs* omits all city fixed effects; $N = 30,967$. Standard errors are clustered by city \times research field.

Figure A4: Mover event study: including stayers



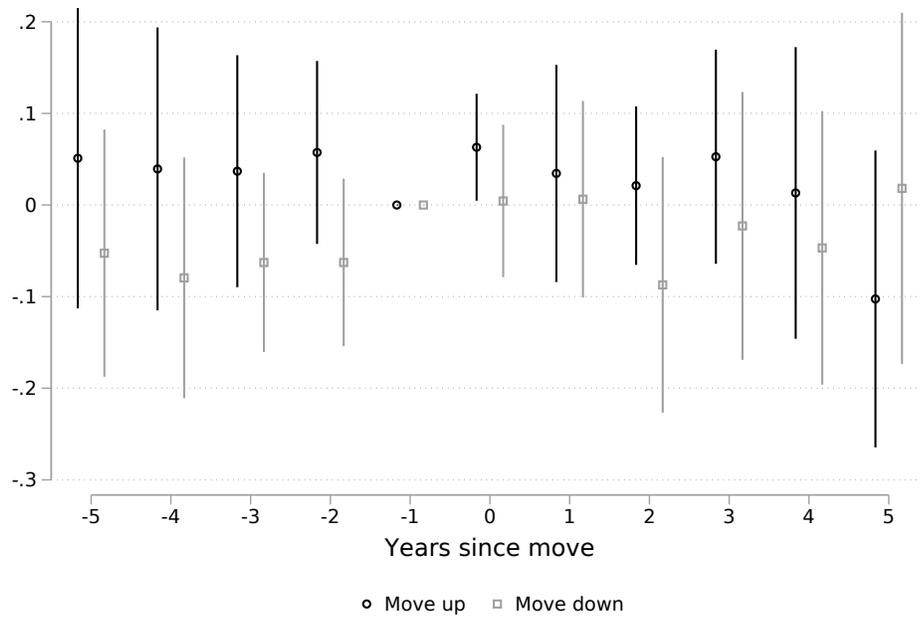
Note: Event study coefficients from Equation 2. The sample includes non-moving inventors as stayers, who are assigned event time -1 and $\Delta\text{Size} = 0$. $N = 474,575$. Fixed effects: origin city \times field, destination city \times field, origin city \times class, destination city \times class, field \times year, class \times year, inventor, origin city \times year, destination city \times year, and firm. Standard errors are clustered by city \times research field.

Figure A5: Mover event study: including stayers, alternative city controls



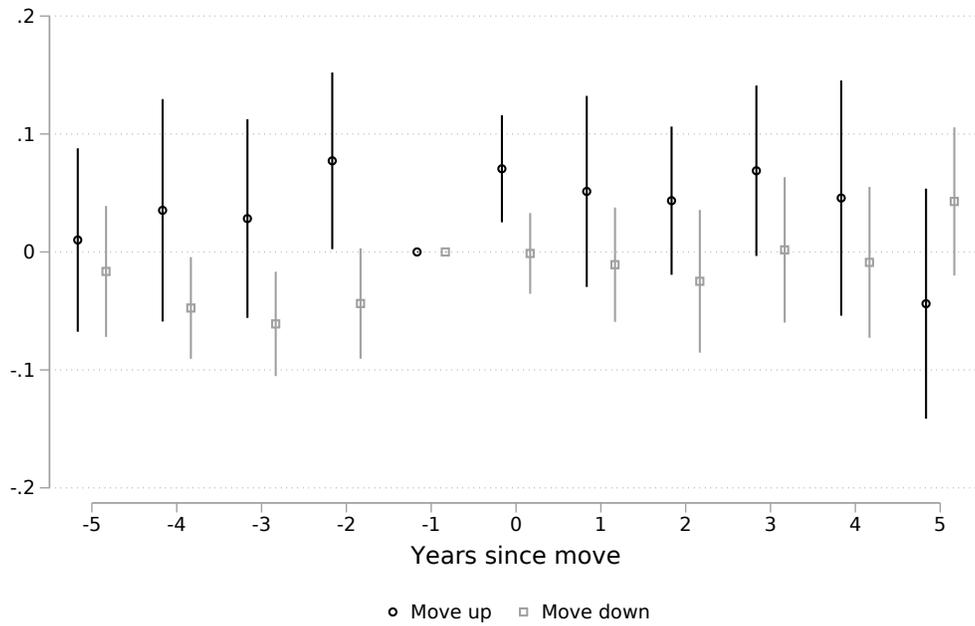
Note: Event study coefficients from Equation 2. The sample includes non-moving inventors as stayers, who are assigned event time -1 and $\Delta\text{Size} = 0$. Fixed effects: field \times year, class \times year, inventor, and firm. *Origin-only* includes fixed effects for origin city \times field, origin city \times class, and origin city \times year; $N = 476,610$. *No city FEs* omits all city fixed effects; $N = 483,324$. Standard errors are clustered by city \times research field.

Figure A6: Mover event study: heterogeneity by moving-up vs. moving-down



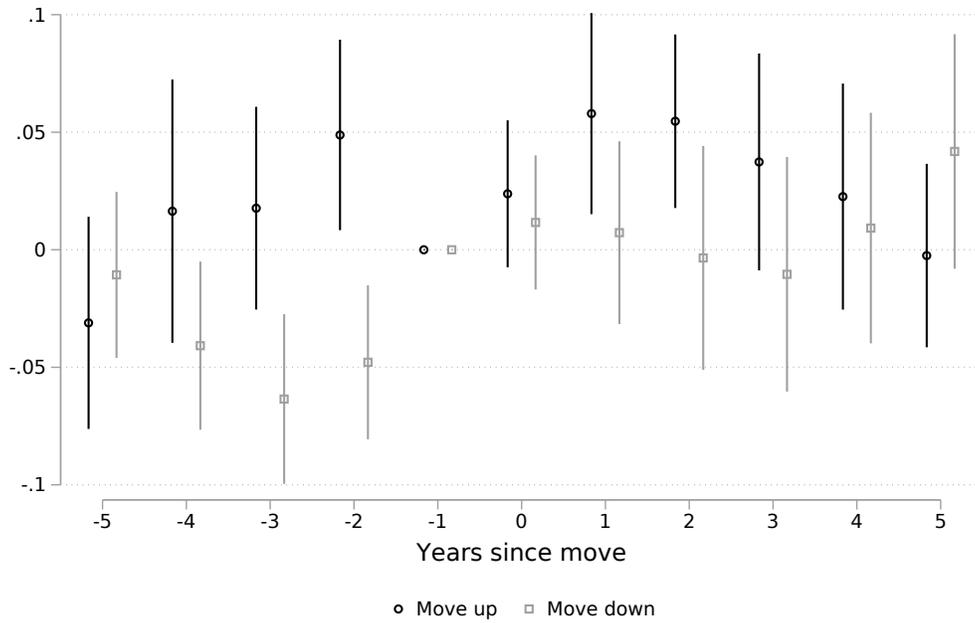
Note: Mover event study with heterogeneous effects by moving to a larger or smaller city. Event study coefficients β_t^{up} and β_t^{down} from $\ln y_{it} = \sum_{s=-5, s \neq -1}^5 (\beta_s^{up} D_{it}^s \times \Delta \text{Size}_{-i} \times \mathbb{1}\{\Delta \text{Size}_{-i} > 0\} + \beta_s^{down} D_{it}^s \times \Delta \text{Size}_{-i} \times \mathbb{1}\{\Delta \text{Size}_{-i} < 0\}) + \text{FEs} + \varepsilon_{it}$. $N = 22,155$. Fixed effects: origin city \times field, destination city \times field, origin city \times class, destination city \times class, field \times year, class \times year, inventor, origin city \times year, destination city \times year, and firm. Standard errors are clustered by city \times research field.

Figure A7: Mover event study: heterogeneity by moving-up vs. moving-down (origin-city fixed effects)



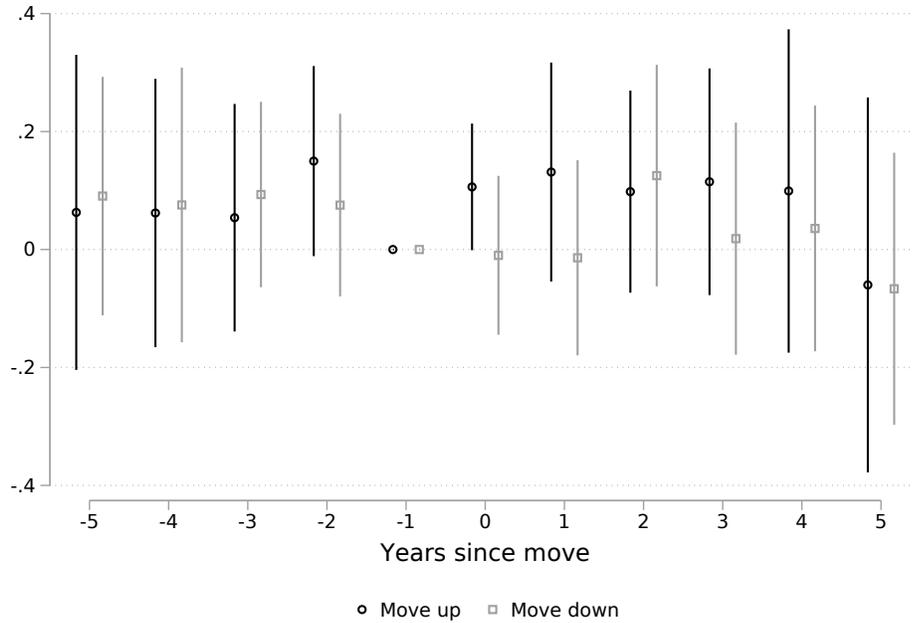
Note: Mover event study with heterogeneous effects by moving to a larger or smaller city. Event study coefficients β_t^{up} and β_t^{down} from $\ln y_{it} = \sum_{s=-5, s \neq -1}^5 (\beta_s^{up} D_{it}^s \times \Delta \text{Size}_{-i} \times \mathbb{1}\{\Delta \text{Size}_{-i} > 0\} + \beta_s^{down} D_{it}^s \times \Delta \text{Size}_{-i} \times \mathbb{1}\{\Delta \text{Size}_{-i} < 0\}) + \text{FEs} + \varepsilon_{it}$. $N = 26,486$. Fixed effects: origin city \times field, origin city \times class, field \times year, class \times year, inventor, origin city \times year, and firm. Standard errors are clustered by city \times research field.

Figure A8: Mover event study: heterogeneity by moving-up vs. moving-down (no city fixed effects)



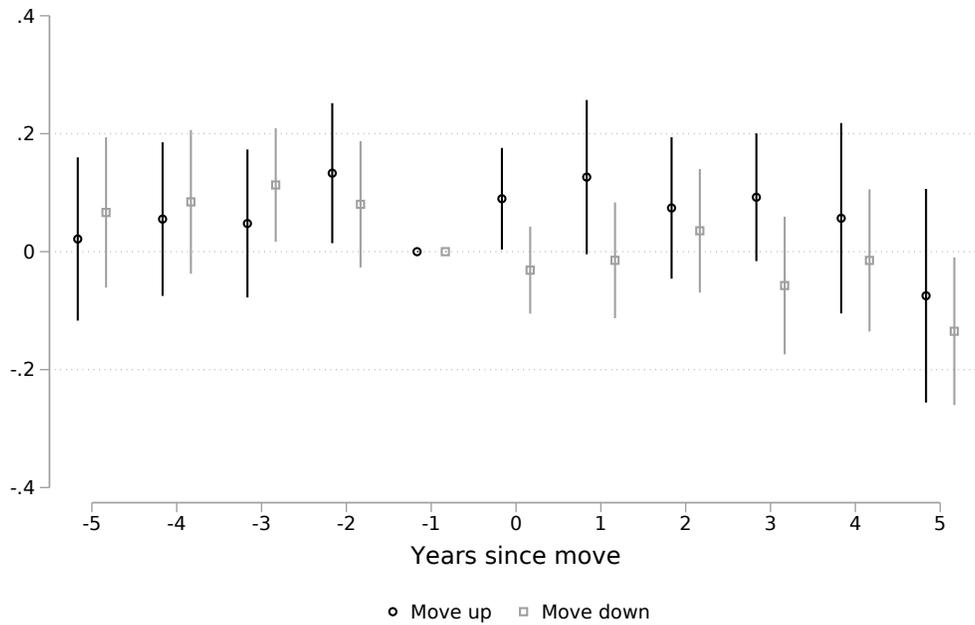
Note: Mover event study with heterogeneous effects by moving to a larger or smaller city. Event study coefficients β_t^{up} and β_t^{down} from $\ln y_{it} = \sum_{s=-5, s \neq -1}^5 (\beta_s^{up} D_{it}^s \times \Delta \text{Size}_{-i} \times \mathbb{1}\{\Delta \text{Size}_{-i} > 0\} + \beta_s^{down} D_{it}^s \times \Delta \text{Size}_{-i} \times \mathbb{1}\{\Delta \text{Size}_{-i} < 0\}) + \text{FEs} + \varepsilon_{it}$. $N = 30,967$. Fixed effects: field \times year, class \times year, inventor, and firm. Standard errors are clustered by city \times research field.

Figure A9: Binary mover event study: heterogeneity by moving-up vs. moving-down



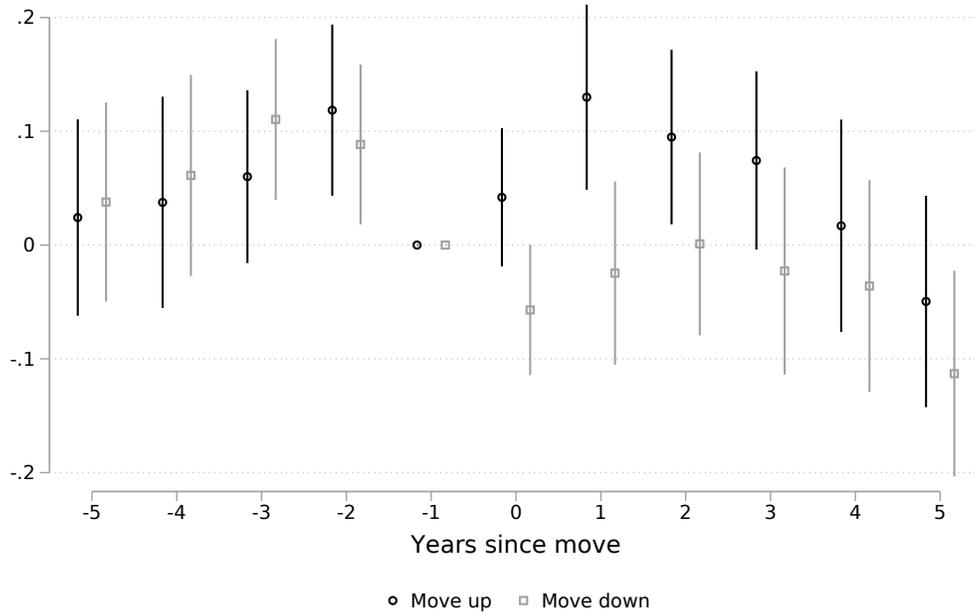
Note: Mover event study with binary treatment variable and heterogeneous effects by moving to a larger or smaller city. Event study coefficients β_t^{up} and β_t^{down} from $\ln y_{it} = \sum_{s=-5, s \neq -1}^5 (\beta_s^{up} D_{it}^s \times \mathbb{1}\{\Delta \text{Size}_{-i} > 0\} + \beta_s^{down} D_{it}^s \times \mathbb{1}\{\Delta \text{Size}_{-i} < 0\}) + \text{FEs} + \varepsilon_{it}$. N = 22,155. Fixed effects: origin city \times field, destination city \times field, origin city \times class, destination city \times class, field \times year, class \times year, inventor, origin city \times year, destination city \times year, and firm. Standard errors are clustered by city \times research field.

Figure A10: Binary mover event study: heterogeneity by moving-up vs. moving-down (origin-city fixed effects)



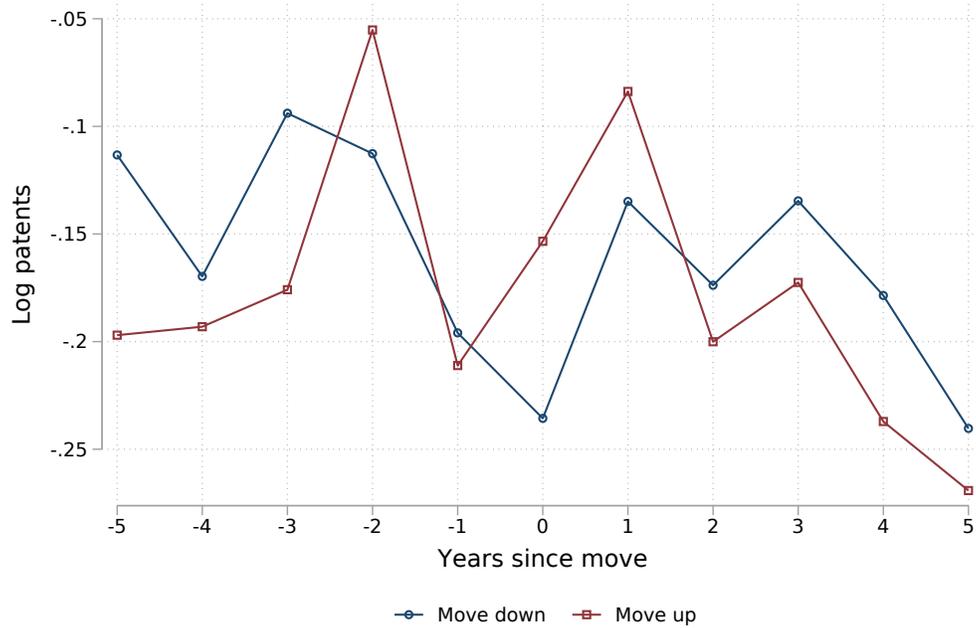
Note: Mover event study with binary treatment variable and heterogeneous effects by moving to a larger or smaller city. Event study coefficients β_t^{up} and β_t^{down} from $\ln y_{it} = \sum_{s=-5, s \neq -1}^5 (\beta_s^{up} D_{it}^s \times \mathbb{1}\{\Delta \text{Size}_{-i} > 0\}) + \beta_s^{down} D_{it}^s \times \mathbb{1}\{\Delta \text{Size}_{-i} < 0\}) + \text{FEs} + \varepsilon_{it}$. $N = 26,486$. Fixed effects: origin city \times field, origin city \times class, field \times year, class \times year, inventor, origin city \times year, and firm. Standard errors are clustered by city \times research field.

Figure A11: Binary mover event study: heterogeneity by moving-up vs. moving-down (no city fixed effects)



Note: Mover event study with binary treatment variable and heterogeneous effects by moving to a larger or smaller city. Event study coefficients β_t^{up} and β_t^{down} from $\ln y_{it} = \sum_{s=-5, s \neq -1}^5 (\beta_s^{up} D_{it}^s \times \mathbb{1}\{\Delta \text{Size}_{-i} > 0\} + \beta_s^{down} D_{it}^s \times \mathbb{1}\{\Delta \text{Size}_{-i} < 0\}) + \text{FEs} + \varepsilon_{it}$. $N = 30,967$. Fixed effects: field \times year, class \times year, inventor, and firm. Standard errors are clustered by city \times research field.

Figure A12: Average patents by event time



Note: Average log patents by event time, separately by inventors moving to larger ('move up') and smaller ('move down') clusters. The sample is restricted to inventors in the estimation sample from Figure 1b. N = 15,767.

C Citations and patent quality

M21 Table 6 Panel B uses citations per patent as the dependent variable, to test for the effect of cluster size on patent quality.²⁰ If patent quality is increasing in cluster size, then M21’s positive patenting-size elasticity is consistent with agglomeration effects, since larger clusters produce both a higher quantity and quality of patents. However, if patent quality decreases in cluster size, then the positive baseline elasticity could be interpreted as defensive patenting, where firms create low-quality patents to use in intellectual property lawsuits. But this suggests the elasticity could even be *negative* when calculated using quality-adjusted patents, contrary to the standard agglomeration model. Hence, the effect of cluster size on citations is key to understanding the sign and magnitude of agglomeration effects for quality-adjusted patents.

There are two issues affecting the M21 Table 6 results. First, M21 adds 0.00001 when calculating log citations, which assigns a large weight to the extensive margin effect. Second, M21 incorrectly calculates citations per patent using citations per coauthor and patents per coauthor; this nullifies the per-coauthor adjustment and inflates the number of citations for inventors with many coauthors.

According to the text, the dependent variable is the “log of number of subsequent patents that cite patents filed by inventor i in year t divided by the number of patents filed by inventor i in year t .” (p.3359) However, this is implemented in the code as $\log\left(\frac{\text{citations}+0.00001}{\text{patents}}\right) = \log\left(\frac{\text{citations}}{\text{patents}} + \frac{0.00001}{\text{patents}}\right)$. In the code, M21 notes that adding 0.00001 avoids dropping observations with zero citations.²¹ But Chen and Roth (2023) reports that the choice of constant c in a $\log(y + c)$ transformation can yield an estimate of any magnitude when there is a nonzero extensive margin. To demonstrate this, I repeat the analysis using $\log(\text{citations}+1)$ as the dependent variable (while noting that $\log(y + 1)$ and $\log(y + 0.00001)$ are both arbitrary). My preferred approach is Poisson regression, which treats the extensive and intensive margins the same way. That is, moving from creating patents with 0 citations to creating patents with 1 citation contributes to the estimate the same as moving from 1 to 2 citations.²²

To adjust for the number of coauthors on a patent, M21 uses fractional attribution when calculating the number of patents and citations. Coauthors of a patent are each assigned an equal share of the patent and the citations received.²³ However, M21 uses fractional patents

²⁰The dependent variable in M21 Table 6 Panel A is total citations across all patents created by an inventor in a given year. This does not measure patent quality on a per-patent basis, since an inventor can receive many citations from creating many patents.

²¹When an inventor does not patent in a given year, that observation is missing from the data. $Citations_{it} = 0$ refers to patents created by inventor i in year t that are never cited, as opposed to inventor-year observations with no patents.

²²Correia et al. (2020) explains that Poisson regression is valid for a non-integer (non-negative) dependent variable.

²³“If a patent has multiple inventors, I assign equally weighted fractions of the patent and citations to each of

in the denominator to calculate citations per patent: $\frac{\text{citations}}{\text{coauthors}} / \frac{\text{patents}}{\text{coauthors}}$. For example, consider an inventor i with a single patent that has one other coauthor and receives 10 citations. Since the total number of coauthors is 2, the number of fractional patents for i is $1/2 = 0.5$ and the number of fractional citations is $10/2 = 5$. In the code, fractional citations per patent is calculated incorrectly as $5/0.5 = 10$. But we should divide by the whole number of patents: $5/1 = 5$ fractional citations.

For inventors with a single patent in a year, this is equivalent to not adjusting for the number of coauthors, since the coauthor terms cancel out.²⁴ I correct the code to calculate citations per patent as fractional citations divided by the whole number of patents. Figure A13 shows a binned scatterplot of log fractional citations per patent against the log number of coauthors, using fractional patents in Panel (a) and whole patents in Panel (b). Since fractional citations is defined as $\frac{\text{citations}}{\text{coauthors}}$, we expect a negative correlation with the number of coauthors. However, Panel (a) shows that M21's incorrect citations-per-fractional-patent variable is positively related to the number of coauthors. When using the corrected citations-per-whole-patent variable in Panel (b), the correlation is negative, as expected. Hence, M21's code inflates the number of citations for inventors with many coauthors. Since the number of coauthors is positively correlated with cluster size, this overstates the number of citations for inventors in larger clusters, meaning that the estimated effect of cluster size on citations is biased upwards.

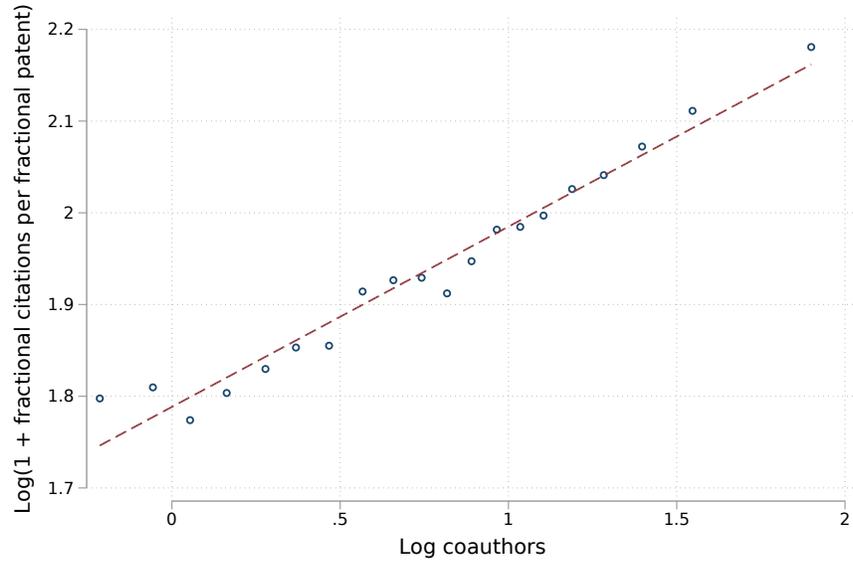
Table A1 presents the M21 Table 6 Panel B results correcting for both issues. First, Panel A reproduces the large positive effect in M21 using $\log\left(\frac{\text{citations}}{\text{patents}} + \frac{0.00001}{\text{patents}}\right)$, where the denominator is fractional patents. As noted above, when the extensive margin is nonzero, the value of the constant c in $\log(y + c)$ determines the magnitude of the estimate. Table A2 shows that the extensive margin effect is positive, implying that cluster size increases the likelihood of creating a cited patent. Hence, the large effect on citations in M21 could be due to the choice of $c = 0.00001$, which assigns a log-scale value of -11.5 to observations with zero citations.

Panel B presents the results using $\log\left(\frac{\text{citations}}{\text{patents}} + 1\right)$, where the denominator is whole patents. Now the effect is negative and statistically significant. Given the positive extensive margin effect, this implies a negative intensive margin effect, where as cluster size increases, inventors shift from creating patents with some citations to creating patents with fewer citations. Is the negative effect from using $c=1$ or from correctly defining citations per patent (using whole patents)? Table A3 shows that the Column 6 estimate is -0.0231 (0.0130) when using $c=1$ and

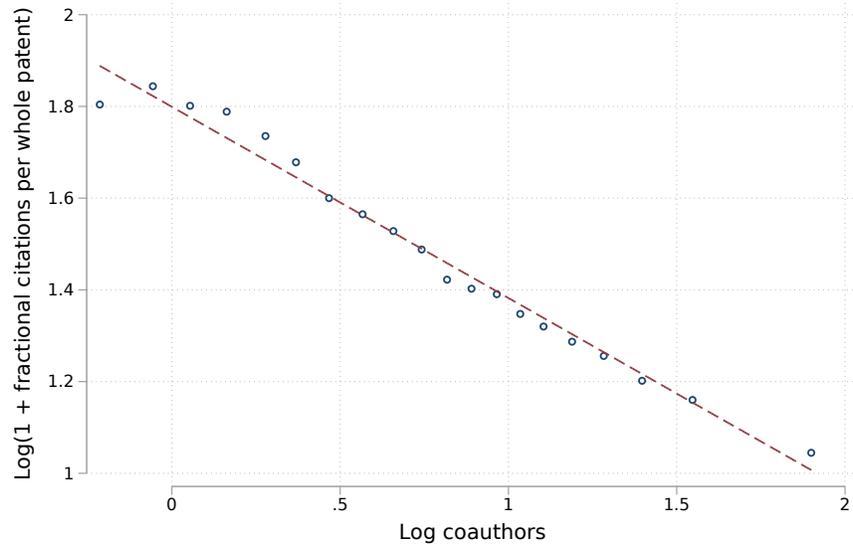
its inventors. For example, if a patent has four inventors, each inventor is credited with one-quarter of a patent and a quarter of the subsequent citations.” (p.3334)

²⁴For inventors with multiple patents in a year, the equivalence is not exact, since fractional patents and citations are calculated and then aggregated to the inventor-year level, before calculating $\frac{\text{citations}}{\text{coauthors}} / \frac{\text{patents}}{\text{coauthors}}$ using the average number of coauthors per patent.

Figure A13: Citations per coauthor is positively correlated with the number of coauthors



(a)



(b)

Note: Panel (a) shows citations per patent calculated incorrectly as the log of fractional citations per fractional patent, so the per-coauthor adjustment cancels out. Panel (b) shows the correct definition as the log of fractional citations per whole patent. Fractional values are calculated by dividing by the number of coauthors. Fixed effects for city, field, and year.

dividing by fractional patents, and is 0.0548 (0.0380) when using $c=0.00001$ and dividing by whole patents. Hence, the positive effect reported in M21 is mostly due to $c=0.00001$ putting more weight on the extensive margin.²⁵ To further explore the impact of the constant c in $\log(y + c)$, Table A5 varies the constant from $c=0.00001$ to $c=1$ while following M21 in dividing by fractional patents. The coefficient decreases monotonically as the constant grows larger, from 0.1174 to -0.0231.²⁶

Panel C uses Poisson regression with fractional citations per whole patent as the dependent variable. Since Poisson regression treats the extensive and intensive margins the same way, it avoids weighting the extensive margin effect in an unsystematic way. Moreover, in this context the extensive margin does not correspond to a discrete economic choice (such as participating in the labor market), but would seem to represent idiosyncrasies in patent quality, as some patents happen to go uncited. That is, a change between 0 and 1 citations is similar to a change between 1 and 2 citations. Hence, Poisson regression is appropriate. The estimates in Panel C are negative and statistically significant, which implies that the intensive margin effect is negative, and outweighs the positive extensive margin.²⁷ The overall effect of cluster size on citations is negative.

As a robustness check, I repeat the regressions using the measure of patent importance from Kelly et al. (2021) as the dependent variable. Figure A14 shows that importance is positively correlated with citations per patent from M21's data. Consistent with the results based on citations, Table A6 shows that cluster size also has a negative effect on patent importance, though the estimate loses significance with city-year fixed effects.

So the original M21 finding is reversed: patent quality is in fact decreasing in cluster size. While cluster size increases patent quality along the extensive margin (by raising the probability of creating a cited patent), it decreases quality along the intensive margin (by reducing the number of citations per patent for inventors with cited patents), and the overall effect is negative. This potentially changes the interpretation of M21's main findings. That is, larger clusters produce more patents than smaller clusters, but those patents are of lower quality, possibly because they are used strategically in intellectual property lawsuits rather than

²⁵Following Chen and Roth (2023), in Table A4 I calibrate the value of the extensive margin using the transformed outcome $m(y) = \log(y)$ for $y > 0$ and $m(0) = -x$. In this specification, the extensive margin is valued as a $100x$ percent change in the intensive margin. For example, when $x = 0.1$, the extensive margin is valued the same as a 10 log point change along the intensive margin. I match the M21 coefficient using $x = 6$. Hence, the M21 specification using $\log(y + 0.00001)$ can be interpreted as valuing the extensive margin as the equivalent of a 600 log point change (or $100(e^6 - 1)\% \approx 40,000\%$ change) in the intensive margin.

²⁶The elasticity for $c=0$ is -0.0325, representing the intensive margin effect when dropping observations with zero citations.

²⁷Since the independent variable is log cluster size and Poisson regression uses a log link function, the Poisson coefficient is an elasticity.

Table A1: Reanalysis of M21 Table 6, Panel B (citations per patent)

	(1)	(2)	(3)	(4)	(5)	(6)	(7) Same city	(8) Different city
Panel A: Original: $\log(\text{citations per fractional patent} + 0.00001/\text{patents})$								
Log size	0.1002 (0.0835)	0.1086 (0.0427)	0.1565 (0.0362)	0.1165 (0.0260)	0.1131 (0.0368)	0.1119 (0.0389)	0.2497 (0.0766)	-0.0340 (0.0296)
Observations	923934	923934	923226	922390	921843	786680	691890	691890
Adjusted R^2	0.449	0.453	0.468	0.514	0.515	0.514	0.287	0.346
Panel B: $\text{Log}(\text{citations per whole patent} + 1)$								
Log size	-0.1231 (0.0173)	-0.1292 (0.0133)	-0.0844 (0.0088)	-0.0850 (0.0105)	-0.0517 (0.0097)	-0.0611 (0.0107)	-0.0533 (0.0108)	-0.0621 (0.0111)
Observations	923934	923934	923226	922390	921843	786680	691890	691890
Adjusted R^2	0.461	0.464	0.487	0.584	0.586	0.605	0.402	0.544
Panel C: Poisson regression (citations per whole patent)								
Log size	-0.1274 (0.0317)	-0.1107 (0.0152)	-0.0893 (0.0144)	-0.0849 (0.0140)	-0.0653 (0.0153)	-0.0774 (0.0160)	-0.0938 (0.0266)	-0.0771 (0.0165)
Observations	923500	923500	921878	920045	919251	782919	654526	691374
Year	Yes							
City \times field	Yes							
City \times class	Yes							
Field \times year		Yes						
Class \times year			Yes	Yes	Yes	Yes	Yes	Yes
Inventor				Yes	Yes	Yes	Yes	Yes
City \times year					Yes	Yes	Yes	Yes
Firm						Yes	Yes	Yes

Standard errors in parentheses

Notes: Citations are measured as the number of patents that cite any patent filed by inventor i in year t ; citations are always calculated per coauthor, dividing by the number of inventors listed on the patent. Citations per patent is measured as citations divided by the number of patents filed by inventor i in year t . Fractional patents are calculated per coauthor, dividing by the number of inventors listed on the patent. Panel A follows the M21 code in calculating citations per patent using fractional patents in the denominator; in this case, the per-coauthor adjustment cancels out. My results in panel A do not exactly match the original Table 6, because the M21 cleaning code is unreproducible; see Appendix A. Panels B and C calculate citations per patent using whole patents in the denominator. The dependent variable in panel C is citations per patent in levels. Column 7 is for citations by inventors located in the same city as the focal inventor. Column 8 is for citations by inventors located in a city different from the focal inventor city. Standard errors are clustered by city \times research field.

representing genuine innovation.

C.1 Coding issues

In Columns 7 and 8 of Table 6, M21 tests whether citations come from patents in the same city. The code to join citing and cited patents uses a many-to-many merge, which makes the results unreproducible by producing a different sample each time the code is run. Moreover, in cases where citing patents have multiple coauthors from different cities, this approach assigns different values of ‘citing-city’ to different cited patents in an unsystematic way. For example, suppose patent p cites patents $q1$, $q2$, and $q3$; the coauthors of p are from cities X and Y . The many-to-many merge will produce the following matching of cited-patent and citing-city: $\{q1:X, q2:Y, q3:Y\}$. In unreported results, I correct the code by assigning each patent the modal city of its inventors and using a many-to-one merge to join citing and cited patents. The estimates are similar.

C.2 Supplementary tables and figures

Table A2: Reanalysis of M21 Table 6, Panel B: extensive margin effect

	(1)	(2)	(3)	(4)	(5)	(6)	(7) Same city	(8) Different city
Log size	0.0196 (0.0071)	0.0210 (0.0034)	0.0203 (0.0029)	0.0183 (0.0023)	0.0120 (0.0028)	0.0133 (0.0031)	0.0126 (0.0056)	-0.0015 (0.0016)
Observations	923934	923934	923226	922390	921843	786680	786680	786680
Adjusted R^2	0.348	0.353	0.369	0.401	0.402	0.391	0.180	0.043
Year	Yes							
City \times field	Yes	Yes						
City \times class	Yes	Yes						
Field \times year		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Class \times year			Yes	Yes	Yes	Yes	Yes	Yes
Inventor				Yes	Yes	Yes	Yes	Yes
City \times year					Yes	Yes	Yes	Yes
Firm						Yes	Yes	Yes

Standard errors in parentheses

Notes: The dependent variable is an indicator for a strictly positive value of total citations received by an inventor for patents created in a given year. Citations are measured as the number of patents that cite any patent filed by inventor i in year t . Column 7 is for citations by inventors located in the same city as the focal inventor. Column 8 is for citations by inventors located in a city different from the focal inventor city. Standard errors are clustered by city \times research field.

Table A3: Reanalysis of M21 Table 6, Panel B (citations per patent): alternative specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7) Same city	(8) Different city
Panel A: $\text{Log}(\text{citations per fractional patent} + 1)$								
Log size	-0.0928 (0.0233)	-0.0993 (0.0139)	-0.0502 (0.0098)	-0.0654 (0.0118)	-0.0126 (0.0123)	-0.0231 (0.0130)	-0.0412 (0.0141)	-0.0250 (0.0131)
Observations	923934	923934	923226	922390	921843	786680	691890	691890
Adjusted R^2	0.483	0.487	0.511	0.601	0.602	0.623	0.421	0.545
Panel B: $\text{Log}(\text{citations per whole patent} + 0.00001/\text{patents})$								
Log size	0.0590 (0.0894)	0.0662 (0.0414)	0.1082 (0.0350)	0.0849 (0.0257)	0.0516 (0.0357)	0.0548 (0.0380)	0.1929 (0.0777)	-0.0908 (0.0293)
Observations	923934	923934	923226	922390	921843	786680	691890	691890
Adjusted R^2	0.466	0.470	0.484	0.527	0.528	0.526	0.280	0.373
Year	Yes							
City \times field	Yes							
City \times class	Yes							
Field \times year		Yes						
Class \times year			Yes	Yes	Yes	Yes	Yes	Yes
Inventor				Yes	Yes	Yes	Yes	Yes
City \times year					Yes	Yes	Yes	Yes
Firm						Yes	Yes	Yes

Standard errors in parentheses

Notes: Citations are measured as the number of patents that cite any patent filed by inventor i in year t ; citations are always calculated per coauthor, dividing by the number of inventors listed on the patent. Citations per patent is measured as citations divided by the number of patents filed by inventor i in year t . Fractional patents are calculated per coauthor, dividing by the number of inventors listed on the patent. Panel A follows the M21 code in calculating citations per patent using fractional patents in the denominator; in this case, the per-coauthor adjustment cancels out. Departing from the M21 code, I add 1 instead of $\frac{0.00001}{\text{patents}}$. Panel B calculates citations per patent using whole patents in the denominator, and follows the M21 code in adding $\frac{0.00001}{\text{patents}}$. Column 7 is for citations by inventors located in the same city as the focal inventor. Column 8 is for citations by inventors located in a city different from the focal inventor city. Standard errors are clustered by city \times research field.

Table A4: Reanalysis of M21 Table 6, Panel B: calibrated extensive margin value ($x = 6$)

	(1)	(2)	(3)	(4)	(5)	(6)	(7) Same city	(8) Different city
Log size	0.0932 (0.0858)	0.1004 (0.0414)	0.1466 (0.0354)	0.1094 (0.0253)	0.1034 (0.0356)	0.1039 (0.0377)	0.1155 (0.0572)	0.0984 (0.0369)
Observations	923934	923934	923226	922390	921843	786680	524513	766583
Adjusted R^2	0.454	0.457	0.472	0.518	0.519	0.517	0.585	0.536
Year	Yes							
City \times field	Yes	Yes						
City \times class	Yes	Yes						
Field \times year		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Class \times year			Yes	Yes	Yes	Yes	Yes	Yes
Inventor				Yes	Yes	Yes	Yes	Yes
City \times year					Yes	Yes	Yes	Yes
Firm						Yes	Yes	Yes

Standard errors in parentheses

Notes: The dependent variable is citations per patent, which is measured as the number of subsequent patents that cite patents filed by inventor i in year t divided by the number of patents filed by inventor i in year t . Following the M21 code, I calculate citations per patent is using fractional citations and fractional patents, so that the per-coauthor adjustment cancels out. Following Chen and Roth (2023), I use the transformed outcome $m(y) = \log(y)$ for $y > 0$ and $m(0) = -x$. I first normalize the outcome by dividing by the minimum non-zero number of citations, so that the extensive margin is a move from 0 to 1. I choose x to match the estimates in M21 Table 6 Panel B. Column 7 is for citations by inventors located in the same city as the focal inventor. Column 8 is for citations by inventors located in a city different from the focal inventor city. Standard errors are clustered by city \times research field.

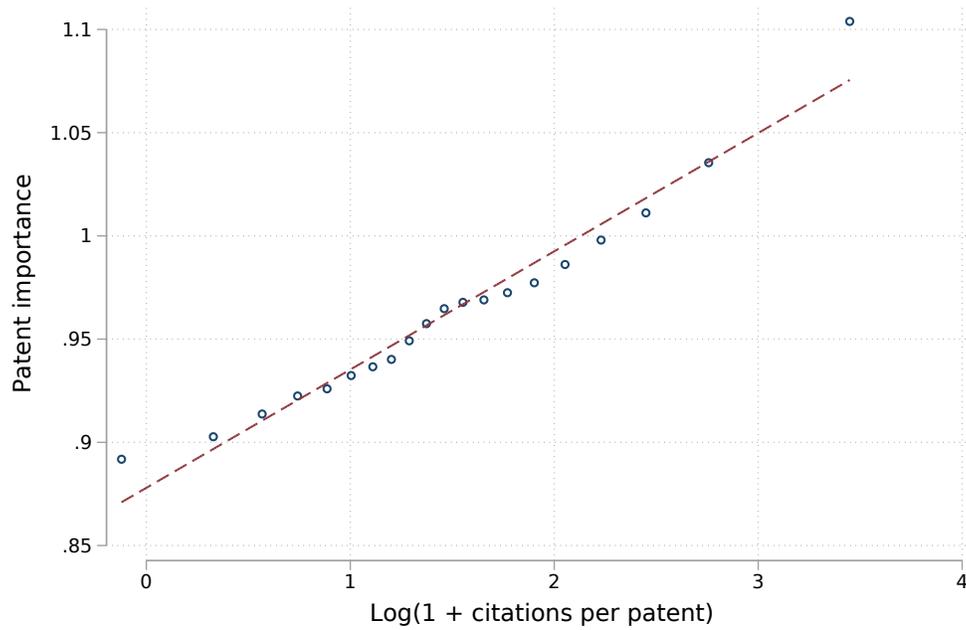
Table A5: Reanalysis of M21 Table 6, Panel B (citations per patent): varying c in $\log(y + c)$

	(1) c=0.00001	(2) c=0.0001	(3) c=0.001	(4) c=0.01	(5) c=0.1	(6) c=1	(7) c=0
Log size	0.1174 (0.0406)	0.0869 (0.0341)	0.0563 (0.0278)	0.0260 (0.0221)	-0.0025 (0.0172)	-0.0231 (0.0130)	-0.0325 (0.0154)
Observations	786680	786680	786680	786680	786680	786680	691890
Adjusted R^2	0.509	0.530	0.557	0.592	0.626	0.623	0.555

Standard errors in parentheses

Notes: Each column follows the Table 6, Panel B, column 6 specification, with fixed effects for city \times field, city \times class, field \times year, class \times year, inventor, city \times year, and firm. The dependent variable is citations per patent, which is measured as the number of subsequent patents that cite patents filed by inventor i in year t divided by the number of patents filed by inventor i in year t . Following M21, citations per patent is calculated using fractional patents in the denominator, so that the per-coauthor adjustment cancels out. Column 1 does not match Table A1 Column 6, since it uses $\log(\frac{\text{citations}}{\text{patents}} + 0.00001)$ rather than $\log(\frac{\text{citations}}{\text{patents}} + \frac{0.00001}{\text{patents}})$. Column 7 drops observations with zero citations. Standard errors are clustered by city \times research field.

Figure A14: Patent importance and citations per patent



Note: The dependent variable is average patent importance for patents filed by inventor i in year t . Importance is measured as the log ratio of forward 10-year and backward 5-year aggregate patent similarity; see Kelly et al. (2021). Citations per patent is measured using whole patents in the denominator. Fixed effects for city, field, and year.

Table A6: Reanalysis of M21 Table 6, Panel B: alternative measure of patent quality

	(1)	(2)	(3)	(4)	(5)	(6)
Log size	-0.0268 (0.0141)	-0.0330 (0.0081)	-0.0126 (0.0042)	-0.0176 (0.0048)	-0.0048 (0.0034)	-0.0048 (0.0034)
Observations	923655	923655	922948	922111	921564	786448
Adjusted R^2	0.598	0.625	0.702	0.773	0.774	0.795
Year	Yes					
City \times field	Yes	Yes	Yes	Yes	Yes	Yes
City \times class	Yes	Yes	Yes	Yes	Yes	Yes
Field \times year		Yes	Yes	Yes	Yes	Yes
Class \times year			Yes	Yes	Yes	Yes
Inventor				Yes	Yes	Yes
City \times year					Yes	Yes
Firm						Yes

Standard errors in parentheses

Notes: The dependent variable is average patent importance for patents filed by inventor i in year t . Importance is measured as the log ratio of forward 10-year and backward 5-year aggregate patent similarity; see Kelly et al. (2021). Standard errors are clustered by city \times research field.

D Aggregate effects of agglomeration

M21 Section V quantifies the aggregate benefits of agglomeration in the United States, using the estimated elasticity of patenting to cluster size. M21 performs an exercise of equalizing cluster size within each field and year, to calculate the decrease in the total number of patents. When agglomeration effects are positive, reallocating inventors from high- to low-productivity places will reduce aggregate output.

To illustrate, consider an example with two cities in the same research field. Inventor i in city c creates patents as a function of cluster size S_c . Given the elasticity α , each inventor produces S_c^α patents; this functional form is derived from the original regression equation with $\ln(y) = \alpha \ln(S)$.²⁸ Since there are S_c inventors, total patents created in city c is then $S_c \cdot S_c^\alpha$. For cities A and B, initial aggregate output is $Y^1 = S_A \cdot S_A^\alpha + S_B \cdot S_B^\alpha$.

In the counterfactual, cluster sizes are equalized: $S'_A = S'_B = \frac{S_A + S_B}{2} := \frac{S}{2}$, where S is the total number of inventors. Aggregate output is $Y^2 = S(S/2)^\alpha$. The change in output in the case of two cities in a single field and year is $Y^2 - Y^1 = S_A \left[\left(\frac{S}{2}\right)^\alpha - S_A^\alpha \right] + S_B \left[\left(\frac{S}{2}\right)^\alpha - S_B^\alpha \right]$. In the general case with many cities, the difference in aggregate output in field f and year t is $Y_{ft}^2 - Y_{ft}^1 = \sum_c S_{cft} [\bar{S}_{ft}^\alpha - S_{cft}^\alpha]$, where \bar{S}_{ft} is the resulting cluster size after the population is equalized across cities, and S_{cft} is the actual cluster size. The percentage difference is $(Y_{ft}^2 - Y_{ft}^1)/Y_{ft}^1$. I reproduce the M21 Table 12 results in Columns 1 and 2 of Table A8 below. Similar to the original, aggregate output falls by around 10% when using a constant elasticity, and by roughly 11% when using heterogeneous estimates that allow the elasticity to vary by cluster size quartiles.

There are several problems with the analysis in M21. First, M21 estimates the elasticity using the subsample of the top 10% of inventors (defined using lifetime patents). Hence, the elasticity should be α_{T10} (estimated on the top 10% subsample) rather than α (estimated on the full sample). So per-inventor output is $S_c^{\alpha_{T10}}$ rather than S_c^α ; note that per-inventor output is based on the total number S_c of inventors in the same cluster, even when restricting to top 10% inventors. More importantly, total patents should be calculated using $S_{T10,c}$, the number of inventors in the top 10% subsample, rather than the total number S_c . Hence, total patents created should be $S_{T10,c} \cdot S_c^{\alpha_{T10}}$ rather than $S_c \cdot S_c^{\alpha_{T10}}$.

Second, M21 describes the exercise as using the number of inventors, but the code uses cluster density, defined as the number of inventors in a city-field-year divided by the field-year total. But when $0 < S < 1$, we have the undesirable behavior that $S^\alpha < S^\beta$ when $\alpha > \beta$: a larger elasticity is associated with a smaller output. This affects the results in Column 2 of M21 Table 12, which uses heterogeneous elasticities that vary by cluster size.

²⁸There is a typo on M21 p.3369: “ $\ln g(S) = \alpha S$, so that $g(S) = S^\alpha$ ”. The first equation should read “ $\ln g(S) = \alpha \ln S$ ”.

Third, the M21 estimates of heterogeneity by cluster size are affected by a coding error. M21 Table 8 tests whether the elasticity of patenting to cluster size varies with cluster size, to understand whether agglomeration effects are different in small versus large clusters. Specifically, Panel A of M21 Table 8 interacts log cluster size with indicators for cluster size quartiles, to estimate the effect of cluster size on patenting separately for each quartile. There is little evidence that the elasticity varies with cluster size. However, the M21 code incorrectly omits the quartile indicators from the regression, which forces the baseline level of patents to be the same across quartiles, resulting in biased estimates of heterogeneity.²⁹ I reproduce the original results in Column 1 of Table A7. In Column 2, I show the corrected regression, which includes the quartile indicators. Now there is substantial heterogeneity across cluster size quartiles, with the elasticity in the fourth quartile being seven times larger than the elasticity in the first quartile (0.0192 versus 0.1432).

Table A7: Reanalysis of M21 Table 8: heterogeneity by cluster size quartiles

	(1)	(2)	(3)
	Original	Corrected: field-year	Corrected: global
Log size \times Q1	0.0601 (0.0136)	0.0192 (0.0126)	-0.0074 (0.0139)
Log size \times Q2	0.0641 (0.0139)	0.0890 (0.0179)	0.0728 (0.0235)
Log size \times Q3	0.0623 (0.0143)	0.1483 (0.0213)	0.1753 (0.0368)
Log size \times Q4	0.0690 (0.0159)	0.1432 (0.0320)	0.2190 (0.0339)
Observations	786680	786680	786680
Adjusted R^2	0.242	0.242	0.242

Standard errors in parentheses

Notes: The dependent variable is log patents. Regression equation: $\ln y_{it} = \sum_{q=1}^4 \beta_q Q_q \times \ln S_{-it} + \sum_{q=1}^3 \lambda_q Q_q + \text{FEs} + \varepsilon_{it}$. Column 1 omits the term $\sum_{q=1}^3 \lambda_q Q_q$, following the original code. My results in Column 1 do not exactly match the original estimates in Table 8, Panel A, Column 2, because the M21 cleaning code is unreproducible; see Appendix A. Cluster size quartiles are calculated within field-year in columns 1 and 2, and globally (across fields and years) in column 3. Coefficients for the quartile indicators Q_q are not reported. Fixed effects: city \times field, city \times class, field \times year, class \times year, inventor, city \times year, and firm. Standard errors are clustered by city \times research field.

Moreover, while unspecified in the text, the M21 code calculates cluster size quartiles within

²⁹M21 Table 8 Panel B tests for heterogeneity by firm productivity instead of cluster size. The code for this regression correctly includes the quartile indicators.

each research field and year, presumably to match the exercise in M21 Table 12, which equalizes cluster size within field-year. But since cluster size varies across fields and years, calculating quartiles within field-year limits the amount of variation across quartiles, and hence understates how much the elasticity varies by cluster size. In Column 3, I repeat the analysis when calculating cluster size quartiles globally, across both fields and years. The heterogeneity is even larger, ranging from -0.0074 in the first quartile to 0.2190 in the fourth quartile.

I repeat the aggregation exercise correcting for these three issues. I multiply per-inventor output by the number of top 10% inventors, instead of the number of all inventors. To avoid the issue of larger elasticities resulting in lower output, I measure cluster size as the number of inventors (so $S \geq 1$) rather than as a density. Finally, I use the corrected heterogeneous estimates from Table A7. The results are shown in Columns 3-5 of Table A8. The reduction in aggregate output is slightly smaller using the constant elasticity (-9.15%, Column 3). When using the corrected heterogeneous estimates based on within-field-year size quartiles from Table A7 Column 2, the reduction in output is twice as large (-23.75%, Column 4).³⁰ Column 5 uses the elasticities from Table A7 Column 3, where size quartiles are calculated across fields and years to capture the full variation in cluster size. Since the heterogeneity across quartiles is more pronounced, the output losses from equalization are even larger, at -35.11%.

Table A8: Reproduction and correction of M21 Table 12

	(1) Original: constant α	(2) Original: het. α	(3) Corrected: constant α	(4) Corrected: het. α (field-year)	(5) Corrected: het. α (global)
Computer science	-11.80	-12.69	-12.28	-29.93	-43.94
Biology and chemistry	-8.94	-9.98	-9.07	-23.67	-35.74
Semiconductors	-13.24	-14.21	-13.40	-30.28	-44.04
Other engineering	-6.75	-7.45	-7.39	-20.71	-28.74
Other science	-8.57	-9.61	-9.01	-23.03	-34.26
All fields	-9.91	-10.84	-9.15	-23.75	-35.11

Notes: Entries are percentage changes in aggregate output from equalizing cluster size within field-year, for year=2007. Columns 1 and 2 follow the original code: per-inventor output is multiplied by the total number of inventors; cluster size is measured as a density; and the heterogeneous elasticities are from Table A7, Column 1. I use the constant elasticity as estimated on my sample: 0.0588. Columns 3-5 multiply per-inventor output by the number of top 10% inventors and measure cluster size as the number of inventors. Column 4 uses the corrected heterogeneous elasticities from Table A7, Column 2, based on cluster size quartiles calculated within field-year. Column 5 uses the heterogeneous elasticities from Table A7, Column 3, based on cluster size quartiles calculated across fields and years. My results in Columns 1 and 2 do not exactly match the original estimates in M21 Table 12, because the M21 cleaning code is unreproducible; see Appendix A.

³⁰When I correct only one issue at a time, the reduction in aggregate output is: -11.66% (using the top 10% population); -9.43% (using the number of inventors); -18.03% (using the corrected heterogeneous estimates). So most of the difference between Column 2 and Column 4 is from using the corrected heterogeneous elasticities.

The above analysis uses elasticities estimated on the subsample of top 10% inventors. M21 does not report estimates using the subsample of bottom 90% inventors. However, M21 Table 9 reports a small elasticity (0.0182) when using the full sample, implying a negative elasticity for bottom 90% inventors. In fact, the baseline elasticity for this subsample is -0.0346 (0.0096). The negative elasticity suggests that cluster size has a crowding-out effect for low-productivity inventors; agglomeration has costs as well as benefits. Table A9 estimates the heterogeneity by cluster size quartiles for the bottom 90% subsample. The negative effect is concentrated in the first quartile, with an elasticity of -0.0681. Hence, it appears that the crowding-out effect is strongest in the smallest clusters.

I repeat the aggregation exercise using the elasticities and inventor populations from the bottom 90% subsample. The negative first-quartile elasticity suggests that equalizing cluster size would amplify the decrease in aggregate output, by reallocating inventors to clusters with agglomeration *diseconomies*. However, the fourth quartile elasticity is not as large, so the overall reduction in output is smaller at -5.83% (see Table A10 Column 1). I also equalize cluster size using the full sample of inventors, with separate elasticities for top 10% and bottom 90% inventors.³¹ This approach effectively averages the results of the top 10% and bottom 90% exercises, yielding a reduction in aggregate output of -22.08% (see Table A10 Column 2).

Hence, the aggregate productivity gains from agglomeration appear to be several times larger than those reported in M21. However, this conclusion must be tempered in light of my earlier results. First, the null results in the event study and IV regressions leave open the possibility that the elasticity is not causal. Second, the negative effect of cluster size on citations implies that patents created in larger clusters tend to be of lower quality. While reallocating inventors from large to small clusters reduces the total number of patents, the corresponding decrease in quality-adjusted patents is less pronounced.³²

³¹I formalize the full sample exercise in Appendix D.1.

³²I repeat the aggregation exercise using citation-weighted patents to adjust for quality. Specifically, I weight each patent by the number of citations received, so the dependent variable is the total number of citations for patents created in a year. To address observations with zero citations, I use Poisson regression, which treats the extensive and intensive margins the same way; the log link function also produces the same functional form in the aggregation exercise. For the top 10% sample, the Poisson estimate of the baseline elasticity (using unweighted patents) is 0.12, and the estimate using citation-weighted patents is 0.08. The corresponding reductions in aggregate patents are -17.75% and -13.07%. In the other specifications, the unweighted and weighted reductions in aggregate patents are -44.94% and -24.08% (using the top 10% sample with heterogeneous elasticities), and -31.95% and -12.16% (using the full sample). Hence, the aggregate productivity gains are 26%, 46%, and 62% smaller when accounting for patent quality.

D.1 Aggregation exercise

The main analysis uses the subsample of top 10% inventors. Here I repeat the aggregation exercise using all inventors.

The size of city c is the sum of its Top10 and Bottom90 inventors:

$$S_c = S_{T10,c} + S_{B90,c}$$

We estimate the elasticities separately by subsample: α_{T10} and α_{B90} . Initial aggregate output is:

$$Y^1 = (S_{T10,A} \cdot S_A^{\alpha_{T10}} + S_{B90,A} \cdot S_A^{\alpha_{B90}}) + (S_{T10,B} \cdot S_B^{\alpha_{T10}} + S_{B90,B} \cdot S_B^{\alpha_{B90}})$$

$S_{T10,A}$ is the number of Top10 inventors in city A. The per-inventor output for Top10 inventors in city A is $S_A^{\alpha_{T10}}$, who have elasticity α_{T10} . Note that the elasticity is applied to the full population of city A: an inventor's output is determined by the total number of other inventors in the same cluster, regardless of type. To obtain patents created, we multiply the per-inventor output ($S_c^{\alpha_i}$) by the number of inventors in each subgroup ($S_{i,c}$) for $c = A, B$ and $i = T10, B90$. In the homogeneous case ($\alpha_{T10} = \alpha_{B90} = \alpha$), we have $Y^1 = S_A \cdot S_A^\alpha + S_B \cdot S_B^\alpha$.

With inventor types, we equalize across clusters by equalizing the number of inventors in each subgroup. As a result, the clusters have the same size.

$$S'_A = S'_B = \frac{S_{T10} + S_{B90}}{2} := \frac{S}{2}$$

where S_{T10} is the total number of Top10 inventors, S_{B90} is the total number of Bottom90 inventors, and S is the total number of inventors (in both cities).

Aggregate output after equalizing is:

$$Y^2 = S_{T10} \left(\frac{S}{2}\right)^{\alpha_{T10}} + S_{B90} \left(\frac{S}{2}\right)^{\alpha_{B90}}$$

For homogeneous α , this simplifies to $Y^2 = S(S/2)^\alpha$.

The difference in aggregate output is:

$$\begin{aligned} Y^2 - Y^1 = & S_{T10,A} \left[\left(\frac{S}{2}\right)^{\alpha_{T10}} - S_A^{\alpha_{T10}} \right] + S_{T10,B} \left[\left(\frac{S}{2}\right)^{\alpha_{T10}} - S_B^{\alpha_{T10}} \right] \\ & + S_{B90,A} \left[\left(\frac{S}{2}\right)^{\alpha_{B90}} - S_A^{\alpha_{B90}} \right] + S_{B90,B} \left[\left(\frac{S}{2}\right)^{\alpha_{B90}} - S_B^{\alpha_{B90}} \right] \end{aligned}$$

This equation is for two cities within a single field and year. In the general case with many

cities, the difference in aggregate output for a field-year is:

$$Y_{ft}^2 - Y_{ft}^1 = \sum_c \{ S_{T10,c} [\bar{S}_{ft}^{\alpha_{T10}} - S_{cft}^{\alpha_{T10}}] + S_{B90,c} [\bar{S}_{ft}^{\alpha_{B90}} - S_{cft}^{\alpha_{B90}}] \},$$

where \bar{S}_{ft} is the resulting cluster size in field f and year t after the T10 and B90 populations are equalized across cities, and S_{cft} is the actual cluster size. The percentage difference is $(Y_{ft}^2 - Y_{ft}^1)/Y_{ft}^1$.

D.2 Supplementary tables and figures

Table A9: Reanalysis of M21 Table 8: bottom 90% inventors

	(1) Bottom 90%
Log size × Q1	-0.0681 (0.0112)
Log size × Q2	-0.0156 (0.0189)
Log size × Q3	-0.0215 (0.0241)
Log size × Q4	0.0545 (0.0336)
Observations	925201
Adjusted R^2	0.359

Standard errors in parentheses

Notes: The dependent variable is log patents. The sample is bottom 90% inventors, defined using lifetime patents. Regression equation: $\ln y_{it} = \sum_{q=1}^4 \beta_q Q_q \times \ln S_{-it} + \sum_{q=1}^3 \lambda_q Q_q + \text{FEs} + \varepsilon_{it}$. Cluster size quartiles are calculated globally (across fields and years). Coefficients for the quartile indicators Q_q are not reported. Fixed effects: city × field, city × class, field × year, class × year, inventor, city × year, and firm. Standard errors are clustered by city × research field.

Table A10: Reproduction and correction of M21 Table 12: additional results

	(1)	(2)
	Bottom 90%:	Full sample:
	het. α	het. α
Computer science	-10.08	-29.29
Biology and chemistry	-7.72	-22.99
Semiconductors	-9.14	-33.54
Other engineering	-2.71	-16.69
Other science	-5.94	-21.14
All fields	-5.83	-22.08

Notes: Entries are percentage changes in aggregate output from equalizing cluster size within field-year, for year=2007. Column 1 uses the sample of bottom 90% inventors (defined using lifetime patents), with heterogeneous elasticities based on global size quartiles from Table A9, Column 1. Column 2 uses the full sample of inventors, with separate elasticities for top 10% and bottom 90% inventors; elasticities are heterogeneous based on global size quartiles (calculated separately by subgroup). Top 10% elasticities are from Table A7 Column 3, and bottom 90% elasticities are from Table A9, Column 1. Aggregate output is calculated by multiplying per-inventor output by the number of inventors in the corresponding subsample, and cluster size is measured as the number of inventors. In Column 2, the counterfactual difference in aggregate output for a field-year is $Y_{ft}^2 - Y_{ft}^1 = \sum_c \left\{ S_{T10,c} \left[\bar{S}_{ft}^{\alpha_{T10}} - S_{cft}^{\alpha_{T10}} \right] + S_{B90,c} \left[\bar{S}_{ft}^{\alpha_{B90}} - S_{cft}^{\alpha_{B90}} \right] \right\}$, where $S_{j,cft}$ is the number of inventors in subgroup j in city c , field f , and year t ; α_j is the elasticity for subgroup j ; \bar{S}_{ft} is the resulting cluster size after the population is equalized across cities; and S_{cft} is the actual cluster size.

E Distributed lag model

M21 Figure 5 estimates the distributed lag model in Equation A1 to test for selection bias from rising star inventors moving into larger clusters. The regression includes five leads and five lags of cluster size. The lead terms (β_1 to β_5) capture how current patenting is affected by future changes in cluster size, while the lag terms (β_{-5} to β_{-1}) capture lagged effects of past changes in cluster size.

$$\ln y_{ijkct} = \sum_{s=-5}^5 \beta_s \ln(\text{Size})_{-ifc(t+s)} + d_{cf} + d_{ck} + d_{ft} + d_{kt} + d_{ct} + d_i + d_j + \varepsilon_{ijkct} \quad (\text{A1})$$

Since cluster size is a continuous variable and changes every year, this specification represents the most demanding form of a distributed lag model, with multiple treatments of different signs and varying intensities (Schmidheiny and Siegloch, 2023). This is a generalization of an event study, where an individual can experience multiple events and event time indicators are scaled by the treatment magnitude. In this case, a different event occurs in every year: a change in cluster size. As a result, this is a staggered treatment adoption design where two-way fixed effects estimates may be biased. M21 does not discuss the assumptions for identifying a treatment effect of cluster size in this setting.³³

The M21 code for Figure 5 defines leads and lags incorrectly. To illustrate, it does not implement a one-year lag as $t-1$, referring to the year prior to t . Instead, it uses the *observation* prior to t . Since the inventor-year panel is unbalanced, a one-year lag can hence be defined as any number of years before t . In other words, the code defines a one-year lag as $t-j$ for $j \geq 1$. For example, β_{-5} is supposed to be the effect of cluster size five years in the past on current patents, but actually corresponds to cluster sizes from five to twenty-nine years in the past.³⁴ M21 states that “for an inventor to be in this sample, the five leads and lags need to be nonmissing, which implies that only inventors observed in 11 consecutive years are included.” (p.3352) However, the code does not enforce the requirement that the observations span consecutive calendar years, so the estimation sample also includes inventors with 11 observations in non-consecutive years.

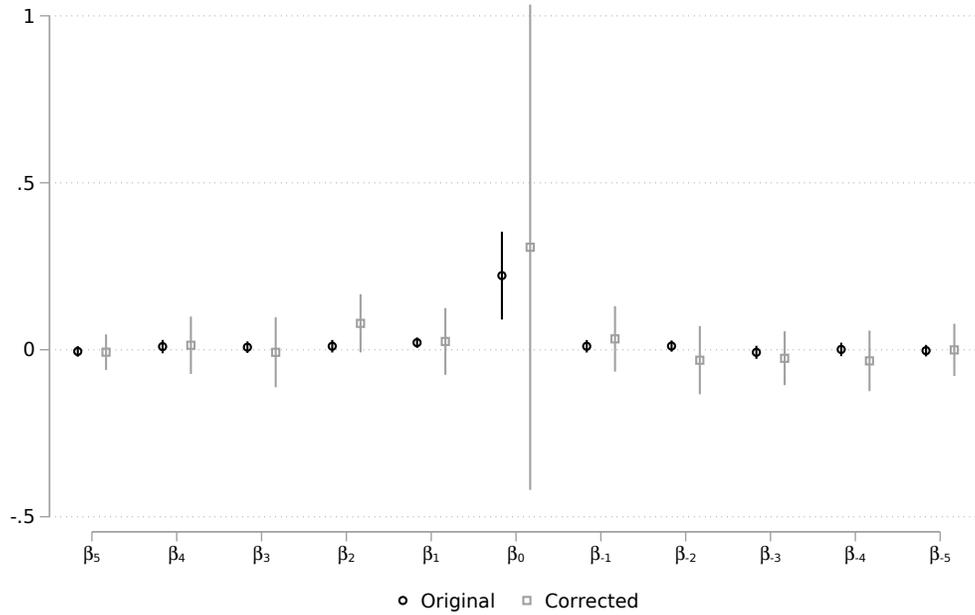
I correct the code to use inventors with observations in 11 consecutive years, and plot the original and corrected versions of M21 Figure 5 in Figure A15. The point estimates are similar, but restricting to observations in consecutive years reduces the sample size, so the confidence

³³Note that the distributed-lag parameters β_s represent the incremental effect of a treatment, whereas the parameters in an event study specification capture cumulative effects. Hence, we aggregate the β_s terms to present the dynamic treatment effects of a change in cluster size.

³⁴The M21 code does not use Stata’s panel operators L. or F., which require a lag to have a one-year gap. Instead, it defines a lag using Stata’s `_n` operator, meaning that a lag is the previous observation, with no restrictions on the length of the gap.

intervals are much wider.³⁵ As a result, the estimate of β_0 becomes nonsignificant.

Figure A15: Reproduction and correction of M21 Figure 5 (top panel)



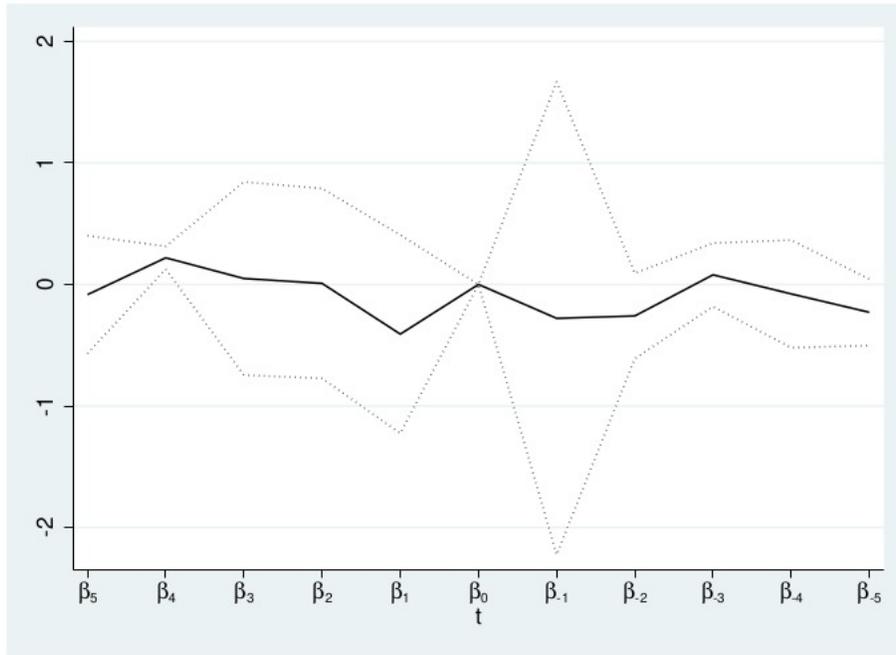
Note: Distributed lag coefficients and 95% confidence intervals from Equation A1, corresponding to the top panel of M21 Figure 5. *Original* includes inventors observed in 11 possibly non-consecutive years. *Corrected* restricts the sample to inventors observed in 11 consecutive years, as stated in the text. $N=56,887$ in *Original*, $N=6,326$ in *Corrected*. Fixed effects: city \times field, city \times class, field \times year, class \times year, inventor, city \times year, and firm. Standard errors are clustered by city \times research field.

As a placebo test, M21 repeats the M21 Figure 5 analysis using the subset of pharmaceutical patents, where R&D takes longer and cluster size should not have an immediate effect on patenting. I show a screenshot of the published M21 Figure A.1 in Figure A16a. We can see that the output does not match M21 Figure 5, since the β_0 term is incorrectly omitted and has no confidence interval. I run the M21 code and display the output in Figure A16b. In contrast to the published figure, the β_0 term is reported; hence, the M21 code does not produce the same figure that is shown in the published article. As with M21 Figure 5, the code for M21 Figure A.1 does not enforce the requirement that inventors are observed in 11 consecutive years. When I correct the code, the sample size is $N=58$ and the β_0 term is dropped as collinear. Since there are very few inventors who create only pharmaceutical patents in 11 consecutive years, the placebo test is not feasible with this sample.

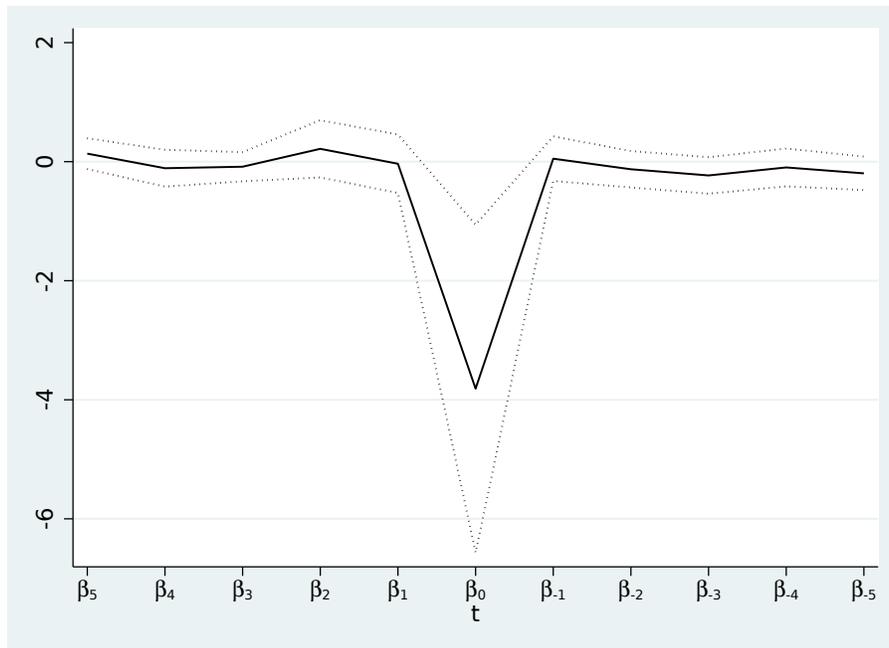
³⁵M21 reports a sample size of 21,787. When I run the M21 code, I obtain a sample of 56,887.

Figure A16: Distributed lag model: pharmaceutical patents

APPENDIX FIGURE 1. DYNAMIC RESPONSE FOLLOWING A CHANGE IN CLUSTER SIZE – PHARMACEUTICAL ONLY



(a) M21 Figure A.1



(b) Reproduction of Figure A.1

Note: Panel (a) is a screenshot of M21 Figure A.1 (top panel), which replicates M21 Figure 5 (top panel) using pharmaceutical patents (class code 424). Panel (b) plots the output from running the original code, corresponding to Equation A1. The solid line represents the coefficients, and the dotted lines are 95% confidence intervals. $N = 585$.

F Cluster quality

In Columns 1-4 of Table A.8, M21 tests whether the elasticity is larger for higher-quality clusters. Cluster quality is defined either as a weighted sum of lifetime patents or citations, or as cluster size measured using inventors with lifetime patents or citations above a threshold. M21 finds that agglomeration effects created by high quality scientists are twice as large.³⁶ There are two coding issues that affect the results.

First, M21's code calculates cluster size differently than described in the text. For the main results, when inventors patent in different cities or fields in the same year, M21 assigns each inventor the cluster size of their *modal* cluster.³⁷ For example, when an inventor patents twice in city A and once in city B, that inventor-year observation is assigned city=A and the cluster size of A. In contrast, to calculate cluster size based on inventor quality in Table A.8, M21's code takes the *average* size across clusters for inventors with patents in different cities. That is, M21 aggregates to the inventor-year level, and then assigns the averaged cluster size to each inventor's modal cluster. Continuing the example, the inventor-year observation is assigned city=A and the unweighted average cluster size of A and B. I correct the code by assigning the cluster size of the modal cluster.

Second, in contrast to the main results which measure fractional patents and citations by dividing by the number of coauthors on a patent, M21 uses unadjusted patents and citations to calculate cluster quality. This treats inventors on large teams as higher quality, since they are assigned the full value of a patent or citation, regardless of the number of coauthors making up the team; moreover, team size is positively correlated with cluster size. I repeat the cluster quality analysis using fractional patents and citations.

Table A11, Panel A reproduces the original results, where an inventor's cluster size is averaged across clusters, and cluster size is calculated using unadjusted (not fractional) values of patents and citations. To show how using fractional values affects the results, Panel B uses average cluster size and fractional patents and citations. The estimates tend to be slightly larger. Panel C assigns each inventor the cluster size of their modal cluster, while using unadjusted patents and citations. The elasticity is much smaller for the weighted cluster sizes (Columns 1 and 3), and is close to the baseline estimate when using inventors above a threshold (Columns 2 and 4). The estimates in Panel D are slightly larger compared to Panel C, when using modal cluster size and fractional patents and citations.

Hence, correcting the definition of cluster size produces elasticities about half as large as

³⁶Note that this implies a smaller (or negative) effect created by low quality scientists.

³⁷"If an inventor has multiple patents with multiple addresses in a single year, I use the modal city for that inventor-year pair. If an inventor has patents in more than one research field or technology class in an year, I use the modal research field or technology class for that inventor-year pair." (p.3334)

those in M21. My results provide mixed evidence for M21’s claim that higher quality inventors matter more for agglomeration effects. The estimates using weighted measures of cluster size are smaller than the baseline elasticity estimated on my sample (0.0588), while the estimates measuring cluster size with inventors above a threshold are slightly larger.³⁸

Table A11: Reanalysis of M21 Table A.8: cluster quality

	(1)	(2)	(3)	(4)
	Weighted: patents	Patents $\geq P$	Weighted: citations	Citations $\geq C$
Panel A: Original: average size, unadjusted				
Log size	0.1200 (0.0430)	0.1301 (0.0513)	0.1097 (0.0351)	0.1301 (0.0529)
Observations	786802	786494	786767	786717
Adjusted R^2	0.244	0.244	0.244	0.244
Panel B: Average size, fractional				
Log size	0.1261 (0.0446)	0.1334 (0.0520)	0.1149 (0.0362)	0.1278 (0.0494)
Observations	786802	786437	786767	786299
Adjusted R^2	0.244	0.244	0.244	0.244
Panel C: Modal size, unadjusted				
Log size	0.0417 (0.0107)	0.0648 (0.0131)	0.0153 (0.0080)	0.0557 (0.0139)
Observations	782036	781636	781997	781934
Adjusted R^2	0.243	0.243	0.243	0.243
Panel D: Modal size, fractional				
Log size	0.0500 (0.0106)	0.0747 (0.0135)	0.0232 (0.0080)	0.0595 (0.0133)
Observations	782036	781563	781997	781396
Adjusted R^2	0.243	0.243	0.243	0.243

Standard errors in parentheses

Notes: Reproduction and reanalysis of Columns 1-4 in M21 Table A.8. The dependent variable is log patents. Cluster size is measured as: the weighted sum of inventors in a given city-field-year cell, with weights reflecting the lifetime number of patents of each inventor (column 1); the number of inventors with a lifetime patent count above a threshold P (column 2); the weighted sum of inventors in a given city-field-year cell, with weights reflecting the lifetime number of citations received (column 3); the number of inventors with a lifetime patent citation count above a threshold C (column 4). Panels A and B follow the original code for Table A.8 and calculate cluster size by averaging across clusters, for inventors with patents in different clusters in the same year. Panels C and D correct the code by assigning each inventor the cluster size from their modal cluster, following the text. Panels A and C follow the original code for Table A.8 and use unadjusted patents and citations, with thresholds in unadjusted values of $P = 3$ and $C = 5$; these correspond to the 75th percentile of patents and the 40th percentile of citations. Panels B and D use fractional patents and citations, with thresholds in fractional values set to the 75th percentiles: $P = 1.5$ and $C = 14$. Fractional values are calculated by dividing by the number of coauthors on a patent. My results in Panel A do not exactly match the original Table A.8 (Columns 1-4), because the M21 cleaning code is unreproducible; see Appendix A. Fixed effects: city \times field, city \times class, field \times year, class \times year, inventor, city \times year, and firm. Standard errors are clustered by city \times research field.

³⁸Given my finding in Table A1 that cluster size reduces citations per patent, citations may not be a good proxy for cluster quality.

G Varying the time unit

In Table A.7, M21 varies the time unit of the data, from 1-month to 3-year periods. As with interpolation in M21 Table A.6, the motivation is to address missing observations when inventors do not patent in a given time period. M21 claims that these missing observations create a downward bias, because the extensive margin effect is not estimated when zeros are not observed:

I expect that when the temporal unit of analysis is short (months), the problem of sample selection and the downward bias are more pronounced. In the extreme, if I were to measure productivity second by second, very few inventor-second pairs would be nonmissing and the selection bias would be large. By contrast, when the temporal unit of analysis is long (two or three years), I expect the problem of sample selection and the downward bias to be less pronounced. In the extreme, if I were to have just one observation per inventor with productivity defined as the number of patents created in all the years in the sample, there would be no selection and both the intensive margin and extensive margin would be reflected in my estimates. (p.3365)

The results in M21 Table A.7 show that the elasticity is negative for 1- and 2-month time units, and grows with the length of the time unit, with the 2- and 3-year elasticities being larger than the baseline.

There are two possible arguments being made. First, that as the time unit shortens, more inventor-time pairs have no patents and are unobserved, so the downward bias is larger; conversely, as the time unit lengthens, fewer inventor-time pairs are unobserved, so the downward bias is smaller. The downward bias here is the difference between the intensive margin effect β_{IM} and the total effect from both the intensive and extensive margins: $\beta_{TOT} = \beta_{IM} + \beta_{EM}$. In other words, the magnitude of the bias is the extensive margin effect $\beta_{EM} > 0$. M21 claims that the bias is increasing as the time unit shortens; that is, β_{EM} is larger when the time unit is shorter. However, β_{EM} can be estimated only if we add zeros to the data by interpolating the unobserved inventor-year pairs with zero patents. Table A.7 does not interpolate, but estimates β_{IM} for different time units. Hence, while it may be true that β_{EM} differs as the time unit changes, this argument does not explain the results in Table A.7, which show estimates of β_{IM} .

A second interpretation is that, as the time unit lengthens, the more the extensive margin is ‘reflected’ in the estimated elasticity; conversely, as the time unit shortens, the estimate ‘reflects’ less of the extensive margin. In other words, with a longer time unit, more inventor-time pairs with zero patents (defined under the shorter time unit) are aggregated into a nonzero inventor-time pair (defined under the longer time unit). It is not clear what it means for the

estimate to ‘reflect’ more of the extensive margin in this case. Suppose an inventor patents in 2000 and 2002, but not in 2001 or 2003. Using a 1-year time unit, the observed data is {2000: 1, 2002: 1}. Using a 2-year time unit, we would observe {2000/2001: 1, 2002/2003: 1}. That is, in each case we observe only the inventor-time pairs with nonzero patents. While there is an extensive margin effect from, for example, not patenting in 2001 and patenting in 2002, this is not captured in either case. Since no zeros are observed, no extensive margin can be estimated.

The intensive margin effect does change with the length of the time unit, because the identifying variation is different. For example, with shorter time units, fewer inventors patent at the same time, so measured cluster sizes are smaller, and fewer patents are summed within each inventor-time bin. Hence, the independent and dependent variables differ. Moreover, with shorter time units, the time fixed effects more finely control for variation over time. These changes explain the pattern of growing elasticities in Table A.7.

Furthermore, there is a coding issue with Table A.7. M21 does not recalculate cluster size at the level of the new time unit, but instead uses the baseline 1-year cluster size; M21 assigns the 1-year value to months within a year, and the first 1-year value in a j -year period to that period. For example, with a 1-month time unit, M21 assigns the cluster size for 2000 to each month in that year; with a 2-year time unit, M21 assigns the cluster size for 2000 to the 2-year period covering 2000-2001.

I reproduce the original results in Panel A of Table A12. Panel B recalculates cluster size at the appropriate time unit. While M21 found negative effects for 1- and 2-month time units, I find large negative and significant effects for 1- to 3-month time units. Similar to the original, the effect size continues to grow as the time unit lengthens. In contrast to the original, the coefficients are larger in magnitude: the 1-month effect is -0.0903 (vs. -0.0258) and the 3-year effect is 0.1972 (vs 0.1613).

Why is the elasticity negative when using short time units? If the intensive margin effect is positive, then the estimated elasticity should still be positive when the extensive margin is omitted. The explanation is based on how cluster size changes with the time unit. Because patents are measured on a per-coauthor basis, the negative elasticity is driven by the mechanical negative relationship between patents per coauthor and the number of coauthors in the same cluster. By definition, coauthors patent at the same time as the focal inventor. As the time-unit shortens, coauthors make up a larger share of an inventor’s cluster; in the extreme case, with patents measured per second, the only other inventors in a cluster would be the focal inventor’s coauthors. So cluster size becomes a proxy for the number of coauthors, which generates a negative correlation between patents per coauthor and cluster size. Table A12 Panel C uses unadjusted patents instead of patents per coauthor as the dependent variable; since coauthors are not included in the denominator, there should not be a negative correlation between total patents and cluster size. As expected, the elasticity is positive in each time unit. Panel D shows

the elasticity between the number of coauthors (team size) and cluster size. The elasticity is about twice as large when the time unit is 1-month compared to the baseline of 1-year. Hence, the negative elasticities are explained by the share of coauthors in the cluster.

Table A12: Reanalysis of M21 Table A.7: recalculating cluster size at new time unit

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	1-Month	2-Month	3-Month	6-Month	1-Year	2-Year	3-Year
<i>Panel A: Original</i>							
Log size	-0.0258 (0.0084)	-0.0147 (0.0091)	-0.0028 (0.0096)	0.0229 (0.0107)	0.0588 (0.0139)	0.1290 (0.0144)	0.1613 (0.0188)
Observations	1286847	1199567	1130305	977635	786680	572459	460922
Adjusted R^2	0.313	0.293	0.282	0.260	0.242	0.211	0.198
<i>Panel B: Corrected</i>							
Log size	-0.0903 (0.0045)	-0.0681 (0.0061)	-0.0422 (0.0073)	0.0045 (0.0094)	0.0588 (0.0139)	0.1512 (0.0179)	0.1972 (0.0255)
Observations	1242696	1181503	1120045	974633	786680	573500	461906
Adjusted R^2	0.338	0.306	0.290	0.264	0.242	0.212	0.198
<i>Panel C: Log unadjusted patents</i>							
Log size	0.0332 (0.0024)	0.0392 (0.0035)	0.0537 (0.0044)	0.0783 (0.0064)	0.1158 (0.0108)	0.1826 (0.0167)	0.2309 (0.0240)
Observations	1242696	1181503	1120045	974633	786680	573500	461906
Adjusted R^2	0.142	0.151	0.167	0.202	0.238	0.255	0.254
<i>Panel D: Log team size</i>							
Log size	0.1279 (0.0032)	0.1124 (0.0038)	0.1025 (0.0045)	0.0826 (0.0054)	0.0684 (0.0071)	0.0475 (0.0092)	0.0541 (0.0105)
Observations	1242696	1181503	1120045	974633	786680	573500	461906
Adjusted R^2	0.475	0.467	0.466	0.472	0.483	0.495	0.499

Standard errors in parentheses

Notes: The dependent variable is log patents per coauthor in panels A and B; log unadjusted patents (not adjusting for coauthors) in panel C; and log team size (average number of coauthors) in panel D. Panel B corrects the original code by calculating cluster size at the level of the corresponding time unit; the M21 code uses the 1-year cluster size for all time units. My results in Panel A do not exactly match the original Table A.7 because the M21 cleaning code is unreproducible; see Appendix A. Fixed effects: city \times field, city \times class, field \times year, class \times year, inventor, city \times year, and firm. Standard errors are clustered by city \times research field.

H Team size

In Table A.8, M21 investigates how team size, measured as the number of coauthors on a patent, affects the relationship between cluster size and patenting. M21 states: “If larger teams are both more productive and more likely to be in larger clusters, team size could be an important omitted variable.” (p.3366) The effect size is larger when controlling for team size or excluding coauthors when measuring cluster size. These results are conceptually flawed because patents are already measured on a per-coauthor basis, so adjusting for team size again means controlling for the same variable twice.

When controlling for a quadratic in team size (column 5), the effect of cluster size on patenting is nearly double the baseline magnitude. However, team size is not a confounder for the effect of cluster size on patenting, but a mediator, and should not be controlled for. With agglomeration effects, larger clusters enable collaboration and lead to larger teams, which in turn increase inventors’ productivity. The total effect of cluster size on patenting includes the direct effect as well as the indirect effect through team size. Hence, controlling for team size should reduce the estimated elasticity by controlling for the indirect effect. But M21 finds a larger elasticity when controlling for team size, implying that the indirect effect is negative.

The explanation for this puzzle is that M21 measures coauthored patents using fractional attribution, assigning each coauthor an equal share of a patent. Hence, controlling for team size means controlling for the denominator of a rate variable (patents per coauthor). As shown in Table A13, Column 1, there is a mechanical negative relationship between patents per coauthor and the number of coauthors. Moreover, team size is positively related to cluster size, since a patent having more coauthors implies there are more inventors contributing to cluster size (see Table A13, Column 2). Hence, using the formula for omitted variable bias, the bias from omitting team size is negative, which explains the larger elasticity when controlling for team size.

A more fundamental problem is that team size is observed only for patents that are created, which complicates estimation of the effect of team size on patenting. We do not observe the full distribution of team size, since team size is unobserved for teams that did not complete a patent. In this dataset, team size is a characteristic of created patents rather than a factor influencing the creation of patents.

In Column 6 of Table A.8, M21 recalculates cluster size to exclude all members of the focal inventor’s team. The M21 code subtracts the inventor’s average team size in a given year.³⁹

³⁹Note that this undercounts the number of coauthors, when the coauthors on different patents are distinct. For example, if inventor A coauthors with B and C on one patent, and separately coauthors with D and E on another, the average number of coauthors is 2, but there are 4 coauthors contributing to A’s cluster size. This approach also overcounts by subtracting coauthors from different cities (who do not contribute to A’s cluster size).

Table A13: Mechanical correlations of team size

	(1)	(2)
	Log patents	Log team size
Log team size	-0.9049 (0.0033)	
Log cluster size		0.0684 (0.0071)
Observations	787061	786680
Adjusted R^2	0.450	0.483

Standard errors in parentheses

Notes: *Patents* is the total number of patents per coauthor created by an inventor in a year. *Team size* is the number of coauthors listed on a patent, averaged by inventor-year. Fixed effects: city \times field, city \times class, field \times year, class \times year, inventor, city \times year, and firm. Standard errors are clustered by city \times research field.

The estimated elasticity is now over three times larger than the baseline estimate (0.237 vs. 0.0676). As above, adjusting the independent variable for team size is inappropriate when the dependent variable is already normalized by team size.⁴⁰

⁴⁰Furthermore, M21's code for the Column 6 result deviates from the text in the same way as the cluster quality results, as discussed in Section F. To calculate cluster size excluding team members, M21 takes the average size across clusters for inventors with patents in different clusters in the same year. This contrasts with the main results, where M21 assigns each inventor the cluster size of their modal cluster.

I Interpolating missing observations

When inventors do not patent in a year, they have zero patents, but their city and field are unobserved. Since these inventor-year observations are missing, the main estimates do not capture the extensive margin effect of cluster size on patenting.⁴¹ In Table A.6, M21 interpolates missing inventor-year observations when inventors have gaps of length 1 or 2 years, and have the same city, field, class, and firm in the years immediately before and after the gap. Specifically, M21 assigns patent=0 and the values of city, field, class, and firm from the adjacent years. M21 then estimates the effect of cluster size on patenting with the extended sample, using the inverse hyperbolic sine and $\log(\text{patents} + 1)$ to account for zero values in the dependent variable.

M21's code contains an error that biases the estimates downwards. The code to interpolate two-year gaps fills in only the second year of the gap, leaving the first year as missing. I correct the code to fill in both years of two-year gaps. Table A14 presents the original and corrected interpolation. Since the extensive margin effect is positive, adding more zero observations further increases the elasticity.

Table A14: Reanalysis of M21 Table A.6: interpolation

	(1)	(2)	(3)	(4)
	Baseline	1 year	One year (2-year gap)	Two years (2-year gap)
<i>Panel A: Inverse hyperbolic sine</i>				
Log size	0.0470 (0.0088)	0.0610 (0.0084)	0.0637 (0.0083)	0.0672 (0.0080)
Observations	767636	806612	819379	850036
Adjusted R^2	0.251	0.221	0.217	0.221
<i>Panel B: Log(patents + 1)</i>				
Log size	0.0354 (0.0066)	0.0467 (0.0063)	0.0489 (0.0063)	0.0517 (0.0060)
Observations	767636	806612	819379	850036
Adjusted R^2	0.254	0.221	0.216	0.221

Standard errors in parentheses

Notes: The dependent variable is patents, as specified in the panel titles. Inventor-year observations are interpolated by assigning patent=0 and the city, field, class, and firm from the year preceding the gap. Column 2 interpolates one-year gaps. Column 3 interpolates one-year gaps and the second year of two-year gaps, following the original code. Column 4 corrects the code by interpolating one-year gaps and both years of two-year gaps. My results in Columns 1-3 do not exactly match the original Table A.6 because the M21 cleaning code is unreproducible; see Appendix A. Fixed effects: city \times field, city \times class, field \times year, class \times year, inventor, city \times year, and firm. Standard errors are clustered by city \times research field.

⁴¹For this sample of star inventors, the extensive margin should be interpreted as an inventor joining a research project in a given year, rather than a non-inventor becoming an inventor.

References

- Chen, J. and J. Roth (2023). Logs with Zeros? Some Problems and Solutions. *The Quarterly Journal of Economics* 139(2), 891–936.
- Correia, S., P. Guimarães, and T. Zylkin (2020, March). Fast poisson estimation with high-dimensional fixed effects. *The Stata Journal: Promoting communications on statistics and Stata* 20(1), 95–115.
- Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy (2021, September). Measuring technological innovation over the long run. *American Economic Review: Insights* 3(3), 303–20.
- Moretti, E. (2021a). Data and code for: The effect of high-tech clusters on the productivity of top inventors. <https://doi.org/10.3886/E140662V1>. American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor].
- Moretti, E. (2021b). The effect of high-tech clusters on the productivity of top inventors. *American Economic Review* 111(10), 3328–75.
- Schmidheiny, K. and S. Siegloch (2023). On event studies and distributed-lags in two-way fixed effects models: Identification, equivalence, and generalization. *Journal of Applied Econometrics* 38(5), 695–713.
- Zucker, L. G., M. R. Darby, and J. Fong (2014). Communitywide database designs for tracking innovation impact: Comets, stars and nanobank. *Annals of Economics and Statistics* (115/116), 277–311.