The Creation and Diffusion of Knowledge: Evidence from the Jet Age*

Stefan Pauly[†] Fernando Stipanicic[‡]

December 23, 2024

This paper provides new causal evidence of the impact of air travel time on the creation and diffusion of knowledge. We construct a novel dataset of the U.S. flight network between 1951 and 1966 and exploit the beginning of the *Jet Age* as a quasi-natural experiment. Travel time between locations over 2,000 km apart decreased by 41%, explaining 30% of the observed increase knowledge diffusion as measured by patent citations. This diffusion spurred new knowledge creation, driving innovation convergence across locations and shifting innovation activity towards the South and West of the United States. JEL Classification: O31, O33, R12, R41, N72, F60

^{*}First version: February 2021. We are grateful for the continuous support of Christian Hellwig since the beginning of this project. We are indebted to Björn Larsson who has made this project possible by sharing with us his collection of historical flight schedules. We thank Taylor Jaworski and Carl Kitchens for sharing highway data with us. This paper has benefited from comments of Antonin Bergeaud, Davide Cantoni, Thomas Chaney, Fabrice Collard, Ben Faber, Ruben Gaetani, Victor Gay, Ulrich Hege, Enrico Moretti, Luigi Pascali, Mohamed Saleh, Mark Schankerman, Claudia Steinwender, Nicolas Werquin, Alex Whalley, Miguel Zerecero, and numerous participants in seminars and conferences in Berkeley, FREIT, IMF, INSEAD/Collège de France, RIDGE Growth, RIEF Paris, SED, Toulouse, Urban Economic Association, YES Princeton, and others.

⁺Email: paulystefan@gmail.com

[‡]University of Oslo. Corresponding author, email: *fernando.stipanicic@econ.uio.no*

1 Introduction

The geography of innovation in the United States underwent a radical shift after World War II. In 1950, the number of patents per capita in the South and the West was less than half as high as in the Northeast and the Midwest. By 2010, this difference had disappeared completely with the emergence of prominent technology clusters in the South and the West. Developing technology clusters is of great interest to policy makers as this may have strong effects on the local economy (Moretti (2010), Moretti and Thulin (2013)). While the literature has emphasized knowledge spillovers as one of the drivers of innovation (Storper and Venables (2004), Furman and Stern (2011), Acemoglu et al. (2016)), highlighting the role of physical proximity for facilitating face to face interactions (Glaeser (2011), Carlino and Kerr (2015), Atkin et al. (2022)), there is scarce evidence on the effect of infrastructure on knowledge spillovers and the geography of innovation (Agrawal et al. (2017), Chatterji et al. (2014)).

This paper provides new evidence on this question by exploiting the nationwide rollout of jet airplanes in the United States during the 1950s and 1960s, which led to a large reduction in air travel time. For instance, travel time between New York City and San Francisco decreased from 11 hours in 1951 to 5 hours and 35 minutes in 1966. This substantial reduction in travel time facilitated long-distance face-to-face interactions.

We find that the decrease in travel time led to an increase in the diffusion of knowledge, which we convert into an increase in access to knowledge. Next, we find that the increase in access to knowledge spurred an increase in the creation of new knowledge. The results provide evidence that jet airplanes drove innovation convergence across locations and contributed to the shift in innovation activity towards the South and the West of the United States.

Our findings align with historical accounts from this period, underscoring the importance of face-to-face interactions for knowledge diffusion and creation. Researchers frequently traveled across the country to meet in person with researchers of other companies to share their advances and learn from each other.¹

¹Gertner (2013), pages 251–252, illustrates this phenomenon based on interviews with inventors from this era. For example, researchers involved in the development of transistors and integrated circuits

We begin by constructing a new dataset of the flight network in the United States during the 1950s and 1960s. We digitize historical flight schedules of the major interstate airlines operating during this period and obtain the fastest routes between every two airports in the network.² We document that between 1951 and 1966, travel time decreased on average by 29%, with a 41% average decrease for Metropolitan Statistical Areas (MSAs) located more than 2,000 km apart.³

This nationwide shock was arguably exogenous as it happened in a strictly regulated environment. We decompose the change in travel time and find that 90% of the reduction was due to improvements in aircraft speed, while 10% resulted from changes in flight routes. This is consistent with the fact that the Civil Aeronautics Board (CAB) imposed strong regulations on the interstate airline market during this period. To promote a *stable* airline industry, the CAB set ticket prices and restricted airlines' entry into new or existing routes. Hence, the reduction in travel time occurred mostly within pre-existing routes rather than through the opening of new ones.

The large decrease in travel time was accompanied by a substantial increase in passenger transport. Reports from the Interstate Commerce Commission show that aggregate passenger miles increased five-fold during the 1950s and 1960s, surpassing those of all other modes of transportation combined (I.C.C. (1965), I.C.C. (1967)). At the same time, a travel survey indicates that 60% of air passenger travel was for business, with the average trip lasting 4.8 days (U.S. Department of Commerce (1958)). In contrast, air transport played a minimal role in goods shipment during this period (I.C.C. (1965)). Thus, the 11-hour reduction in coast-to-coast travel time primarily facilitated face-to-face interactions for business travelers rather than improving goods shipment.⁴

at Fairchild Semiconductor in Palo Alto, California, Texas Instruments in Dallas, Texas, and Bell Labs in Murray Hill, New Jersey, would frequently meet in person to share advances. We expand on these accounts in Section 2.

²The six domestic airlines in our data accounted for 75% of total air passenger transport.

³New York and Boston are about 300 km apart, while New York and San Francisco are located about 4,130 km apart. Between 1951 and 1966, we observe a 23% reduction (13 minutes) in travel time between New York and Boston, and a 49% reduction (5 hours and 25 minutes) between New York and San Francisco.

⁴Firm-employed inventors are likely to be highly sensitive to travel time (Perlow (1999)). For a review of elasticities of business travel with respect to travel time, see Wardman (2012).

To study knowledge creation and diffusion we use patent data. We follow Jaffe et al. (1993) and use patent citations as our observable measure of knowledge flow.⁵ We assemble one dataset with all corporate patents granted by the United States Patent and Trademark Office (USPTO) with filing year between 1949 and 1968, which includes for each patent: filing year, technology classification, location (MSA) of the inventors when they applied for the patent, owner of the patent and citations to other patents which were granted by the USPTO.

We document three facts of patenting activity during our sample period. First, patent growth was stronger in initially less innovative MSAs. Second, it was also stronger in the South and the West of the US. Third, the mass of citations shifted towards longer distances. Our results show that the decrease in travel time contributed to all three facts.

We do our analysis in two steps. In the first step, we estimate a gravity equation to obtain the elasticity of citations to travel time. We identify the elasticity exploiting only within establishment-pair across-time variation in citations and travel time. The estimated elasticity implies that citations increased on average 2.4% due to the decrease in travel time between 1951 and 1966. We find that the absolute value of the elasticity is increasing with the distance between the citing and cited establishments. At a distance of more than 2,000km, the change in travel time implies an increase in citations of 6.8%, contributing to the shift of citations towards longer distance.

In order to rule out the possibility that the opening of new routes or the timing of adoption of jets at the route level was driven by variables that also affected knowledge flows, we perform an instrumental variables estimation. We instrument the observed travel time with a fictitious travel time computed by holding flight paths constant at their initial configuration and assuming in each year all routes are operated with the year's average airplane. Hence, changes in instrumental travel time are only due to the nationwide roll out of jets and is thus independent of decisions at the route level. The

⁵A vast literature uses patent citations as a proxy for knowledge diffusion. Recent examples include Atkin et al. (2022), Cai et al. (2022), and Bahar et al. (2023). Inventor surveys indicate that citations are a noisy but still informative signal for knowledge diffusion between inventors (Jaffe et al. (2000a), Roach and Cohen (2013)).

key source of variation in the instrument is the increase of in-flight speed relative to takeoff and landing time. As consequence, longer paths or paths with fewer stops –where in-flight speed plays a larger role– experience greater reductions in the instrumental travel time, consistent with patterns observed in the data. The instrumental variable results confirm the baseline results, reflecting the reduced scope for endogeneity of travel time. In addition, the results are robust to controlling for potential confounding factors such as changes in highway travel time, telephone connectivity, flight ticket prices and for a time-varying effect of distance. The results also remain after restricting the sample to contain only establishments that existed in the initial time period and when restricting the sample to MSA-pairs that always require at least one connecting flight.

In the second step, using the estimated elasticity of diffusion of knowledge, we compute a measure of knowledge access that is specific to each location-technology. The measure captures changes in knowledge access that are only consequence of the change in travel time. We use this measure to estimate the elasticity of new patents to knowledge access. We identify the elasticity at the establishment level comparing only across time variation in patents and knowledge access across establishments within a location, conditional on aggregate technological trends. Thus, the identification is independent of location specific changes in local population or R&D subsidies. We find that the increase in access to knowledge led to a 3.5% yearly growth rate of new patents filed.

Given the reduction in travel time was larger for longer distances, the increase in knowledge access was stronger in locations geographically far from the initial innovation centers located in the Midwest and the Northeast. Hence, by increasing access to knowledge, the reduction in travel time contributed to the shift in the distribution of innovative activity towards the South and the West of the US that we observe in the data. The change in travel time predicts that the South and the West would have an average yearly growth rate of patenting 0.75 percentage points higher than the Northeast and the Midwest during our sample period.

We find that the value of the elasticity of patents to knowledge access is bigger in

magnitude for establishments located in initially less innovative locations. Within each technology class, we rank locations according to the amount of patents in the initial time period and split them into four quartiles. We find that the increase in knowledge access predicts a 4.5% yearly growth rate of patenting in locations in the lowest quartile of initial innovativeness, while it predicts a 3.4% yearly growth rate in the highest quartile. The difference in growth rates indicates that the increase in knowledge access acted as a convergence force between locations, consistent with the data. Results go in the same direction if we rank locations in terms of patents per capita.

We test whether improved face-to-face interactions led to an increase in patenting through mechanisms other than knowledge spillovers. We find that neither changes in market access nor financial access can explain the results. Two thirds of the effect remain after controlling for changes in market access. Especially, the convergence effect remains unchanged. Financial regulation restricted inter-state banking activity, limiting the scope of increase financial access. We present suggestive evidence that the results are not driven by a decrease in financial frictions.

Our results are robust to controlling for changes in market access by highway, time changing telephone connectivity and to computing knowledge access using only knowledge located at long distances.

We also estimate the elasticity by instrumental variables, constructing an instrumental knowledge access with the instrumental travel time. Using the instrumental knowledge access and a recentered version following Borusyak and Hull (2023), we obtain results that go in the same direction as in the baseline analysis.

Literature. This paper contributes to multiple branches of literature. First, it contributes to the literature on agglomeration and knowledge spillovers. Agglomeration forces are usually understood as happening in a geographically localized manner (Glaeser (2011), Arzaghi and Henderson (2008)). The literature on technology clusters also documents this fact (Duranton et al. (2009), Kerr and Robert-Nicoud (2020), Moretti (2021)). The seminal paper Jaffe et al. (1993) finds that patent citations decay rapidly with distance. Our results show that jet airplanes allowed long distance knowledge spillovers, facilitating the development of technology clusters in other regions. We contribute to the literature on transportation by constructing a novel data set and studying a large change in transportation technology that isolates a shock to the mobility of people. To the best of our knowledge, this is the first quantitative analysis of the change in air travel time due to the roll out of jet airplanes in commercial aviation. Other papers have studied the impact of transportation improvements on innovation. Agrawal et al. (2017) study the impact on innovation of a region's stock of highways, while Perlman (2016) uses 19th century data on locations' density of railroads. Berger and Prawitz (2024) and Tsiachtsiras (2021) do so using the historical railroad expansion in Sweden and France. Relative to them, we contribute by exploiting a quasi-natural experiment that allows us to isolate a channel of face to face interactions, with little scope for a trade channel.

In contemporaneous work Bai et al. (2023) study the impact of the opening of airline routes on patent citations in a more recent set up. Bahar et al. (2023) shows non-stop flights lead to more citations and collaborative patents in an international context. Both papers represent supporting evidence that face to face interactions are relevant for the diffusion of knowledge and hence it is a plausible and reasonable mechanism in our context. We contribute to this literature by exploiting a large, nationwide shock with limited risk of endogeneity. Importantly, we show that decreased travel time not only leads to an increase in knowledge diffusion, but also convert such knowledge diffusion into a measure of potential spillovers and show that it affects the creation of new knowledge. Additionally, we highlight that knowledge spillovers shape the geography of innovation and can act as a convergence force.

The impact of transportation improvements in economic outcomes has long been a subject of study (Fogel (1963), Baum-Snow (2007), Michaels (2008), Donaldson and Hornbeck (2016), Campante and Yanagizawa-Drott (2017), Jaworski and Kitchens (2019) and Herzog (2021)). Our convergence result contrasts with previous studies on improvements in other means of transport. Pascali (2017) finds that the introduction of steam engine vessels in the second half of the 19th century led to an increase in international trade which contributed to economic divergence between countries. Faber (2014) finds that the expansion of the highway system in China led to a reduction of GDP growth of peripheral counties, with evidence suggesting a trade channel. While both papers emphasize a trade channel, in our set up the trade channel would not be of first order. Hence, we uncover a new effect of improved connectivity.

Feyrer (2019) studies the time-varying effect of air and sea distance on international trade and economic growth, showing little to no effect of air distance for trade during the 1950s and early 1960s. Additionally, our results are robust to controlling for a time-varying effect of air distance, revealing that flight schedules contain relevant information. Our paper is related to Campante and Yanagizawa-Drott (2017) who study changes in international airplane connectivity, finding that it affects capital flows and the spatial distribution of economic activity. Other literature has also found that business travel affects innovation (Hovhannisyan and Keller (2015)), trade (Söderlund (2023)) and industrial activity (Coscia et al. (2020)).

Finally, we contribute to the contemporaneous literature on innovation on the post WW2 period. Gross and Sampat (2023) study the long-term effects of the public R&D funding by the Office of Scientific Research and Development (OSRD) during WW2. They find that this R&D shock enlarged pre-existing patenting gaps across locations. Kantor and Whalley (2023) study the effects of NASA spending and the race to the moon during the 1960s, finding stronger growth in county-industries that were more space-relevant before the Space Race. We contribute to them by studying a shock that is different in its nature: improved connectivity, rather than increased expenditure. We provide evidence that this shock contributed to the post WW2 shift in innovative activity towards the South and the West. Additionally, we show that improved connectivity had a differential effect on innovation which contributed to closing the patenting gap between locations. This is in line with the decline in innovation concentration during our period of analysis documented in Andrews and Whalley (2021).

2 Historical context

This section describes elements of the historical context that are relevant for identification and interpretation of the results. A more detailed description is presented in Appendix B.

Face to face interactions. Historical accounts from this time period refer to face to face interactions as being an important driver to diffuse and access to knowledge. For example, Bell Labs would organize in-person conferences to explain their new technologies, including the conferral of considerable informal and tacit knowledge (Nagler et al. (2022), Holbrook et al. (2000)). Gertner (2013) on pages 251 and 252, quoting interviews with key inventors of the period, states that during this period, "information was freely exchanged" and inventors of different firms would visit each other in their respective laboratories. Ian Ross, president of Bell Labs during more than 10 years, would state that in order to learn about semiconductor devices inventors would travel to Bell Labs subsidiaries in Murray Hill, NJ, and in Building 2, NY, clearly showing the relevance of being able to travel in person to the places in which innovation takes place in order to learn.⁶

Large increase in air passenger travel and little air transport of goods. Figure 1 in Appendix B extracted from a report of the Interstate Commerce Commission (I.C.C. (1965) and I.C.C. (1967)) shows that air transport accounted for less than 0.1% of the total ton-miles. On the other hand, Appendix Figure 2 shows that air passenger-miles increased five-fold during our period of analysis, reaching a level equal to three times the one of rail travel. Thus, improvements in air travel mainly affected the mobility of people and not the shipment of goods.

Business travelers. The report of the first-ever nationwide survey of travelers (U.S. Department of Commerce (1958)) shows two important pieces. First, presented tables imply that more than 60% of air passenger travel was due to business travel. Second, the average business trip took 4.8 days. Given this context, an 11-hour reduction in travel time for a return coast-to-coast travel represents a substantial improvement for business travelers. Hence, the observed increase in air passenger travel in Appendix

⁶"If you wanted to know about semiconductor devices, you went to Murray Hill [New Jersey] and Building 2 [Manhattan, NY]". Quotes are from Gertner (2013) pages 251 and 252, referring to an interview with Ian Ross president of Bell Labs in 1979-1991. Inventors who would visit each other include Morrey Tanenbaum, inventor of the silicon transistor who was based at Bell Labs in Murray Hill, New Jersey; Jack Kilby, Nobel laureate in Physics inventor of the hybrid integrated circuit at Texas Instruments in Dallas, Texas and; Robert Noyce co-founder of Fairchild Semiconductor and later on Intel Corporation, inventor of the monolithic integrated circuit.

Figure 2 likely reflects a considerable increase in business travel.

Regulation. As explained in Borenstein and Rose (2014), the Civil Aeronautics Board (CAB) regulated the market by deciding which airlines could fly, in which routes they could operate, the price that they charged in each route, the structure of subsidies and merger decisions. When the CAB was created in 1938, it conceived special rights to the existing airlines over the connections they were operating. The CAB did not permit entry of new airlines on interstate routes and gradually allowed current airlines to expand their routes. Importantly, Borenstein and Rose (2014) in pages 68-69 explain that *"the regulatory route award process largely prevented airlines from reoptimizing their networks to reduce operation costs or improve service as technology and travel patterns changed."* As a consequence, any technological improvement such as increases in aircraft speed, capacity or range would not affect each airline's flight network in the short term.

3 Air travel data

We construct a new data set of the flight network in the United States during the 1950s and 1960s. We collected and digitized information of all the flights operated by the main airlines and obtained the fastest route and travel time between every two airports in the network. Additionally, we constructed an instrumental travel time that is not affected by the opening of new routes neither the allocation of jets across routes. Details are provided in Appendix C.

To construct the flight network we digitize historical flight schedules for the years 1951, 1956, 1961 and 1966 of six domestic airlines: American Airlines (AA), Eastern Airlines (EA), United Airlines (UA), Trans World Airlines (TWA), Braniff International Airways (BN), Northwest Airlines (NW).⁷ The six domestic airlines together accounted for between 77% and 81% of interstate air revenue passenger miles C.A.B. (1966).⁸

Appendix Figures 5 and 6 display the flight network in continental United States.

⁷Appendix Figure 3 is a fragment from a page of the 1961 flight schedule of American Airlines. The selection of years was done based on data availability and with a criteria to be equally spaced.

⁸AA, EA, UA and TWA were referred as the *Big 4* which accounted for between 69% and 74% of interstate air revenue passenger miles in the US in the years collected. BN and NW were digitized in order to increase the geographical coverage.

In total we have digitized 5,910 flights (unique combinations of flight number-year). However, flights often have multiple stops. If we count each non-stop part (*leg*) of these flights separately, our sample contains 17,469 legs. Our data connects 275 US airports creating 2,541 unique origin-destination (directional) airport links.

Using departure and arrival time of each flight at each airport, we obtain the fastest route and corresponding travel time between every two airports in our data. We modify the Dijkstra algorithm to account for layover time in case the fastest route includes connecting flights. We then match every airport to 1950 Metropolitan Statistical Areas (MSA) in contiguous United States using the shape file from Manson et al. (2020). We match each airport to all MSAs for which it lies inside the MSA or is at most 15km away from its boundary.⁹ We use the sample of 108 MSAs that are matched to at least one airport in the four years as our baseline travel time data.¹⁰

3.1 Descriptive statistics: Air travel

Appendix Figure 8 shows that from 1951 to 1966, the adoption of jets was stronger for longer routes. In 1961, all MSA-pairs more than 3,000km apart connected with a non-stop flight operated at least one jet flight, and this expanded to all those more than 2,000km apart in 1966. Additionally, conditional on flying non-stop and having a jet, longer routes got a larger reduction in travel time.¹¹ Appendix Figure 10 shows the large decrease in travel time for MSAs connected with a non-stop flight. The figure also exposes the emergence of +8 hours flights in 1956 due to a change in regulation.¹² Given the concern that such change in regulation may be endogenous and that it affects long distance routes, in the construction of our instrument we keep only non-stop routes existing in 1951.

⁹The 15km distance was chosen after inspecting airports outside MSAs that are near the border and should arguably be matched, as for example, Atlanta ATL airport.

¹⁰In Appendix C we include a table with the 168 MSAs, those connected at least once and those connected in the four years.

¹¹For example, two routes always connected non-stop as Chicago to Los Angeles (2,800 km apart) and New York to Boston (300 km apart) got respectively 47% and 23% reduction in travel time.

¹²In 1951 the regulation set a 8-hour maximum flight time allowed for a crew in a 24-hour period. In 1956, the new regulation up to 10 hour flights for transcontinental flights which made it possible to connect New York City and Los Angeles with a non stop flight in 8 hours 30 minutes.

The change in travel time in non-stop flights is also reflected in the travel time for connecting flights. Figure 1 shows, relative to 1951, the average change in travel time for all MSA-pairs, including non-stop and connecting flights. The continuous line shows the observed change and the dashed line shows the change in the instrumental travel time explained below. Between 1951 and 1956, there is an average reduction in observed travel time of 9.2% which is roughly constant for all distances over 500km. Between 1951 and 1961, there is a reduction in travel time that is increasing with distance. The average decrease in travel time is of 16.8%, while the reduction is of 29.4% for a distance of more than 2,000km and 39.2% for a distance of 4,250-4,500km. Between 1951 and 1966, there is an even stronger decrease in travel time at all distances. The average reduction in travel time is 28.7% across all distances, 40.8% for a distance of more than 2,000km and 48.4% for a distance of 4,250-4,500km. The increased adoption of jets for short distance flights implied that both non-stop flights at short distance and connecting flights at farther distance had a decrease in travel time.

Figure 12 in Appendix C shows that the change in travel time is accompanied by a reduction of the amount of legs needed to connect two MSAs at every distance. However, between 1956 and 1961 we do not observe a big reduction in the amount of legs, implying that the decrease in travel time observed in Figure 1 between 1956 and 1961 comes from a source other than the amount of legs. In Appendix Figure 13 we open up the change in travel time by the way an MSA pair was connected in 1951 and 1966: either directly (non-stop flight) or indirectly (connecting flight).¹³ Importantly, MSA-pairs more than 2,000km apart that were connected indirectly in both periods got an average reduction in travel time of 42%. This fact shows that in case we are concerned about endogeneity of jet adoption in non-stop flights, there would still be a large decrease in travel time when focusing on routes that require connecting flights.

It could be the case that a reduction in the amount of legs or an increase in frequency

¹³Much of the increase in travel time for MSA pairs less than 250km apart comes from routes that in 1951 were operated non-stop while in 1966 were operated with connecting flights. Appendix Figure 14 repeats the exercise discarding layover time in all time periods. By comparing Figure 13 and Figure 14 we can disentangle the effect of layover time and the change in in-flight time. For MSA pairs less than 250km that changed from direct to indirect connection, 80% of the increase in travel time is due to the increase in layover time (which was previously zero as it was a non-stop flight), and 20% is due to the increase of in-flight time.



Figure 1: Observed and Instrumental Travel Time between US MSAs. Change in travel time within MSA-pairs relative to 1951, including non-stop and indirectly connected MSA-pairs. Each dot represents the average change for MSA-pairs within a 250km bin (e.g. 2,000km -2,250km). The solid line shows the observed change in travel time. The dashed line shows the change in the counterfactual travel time used as an instrument. Counterfactual travel time is calculated by fixing the MSA-pair flight path to the 1951 path and applying the average speed and takeoff/landing time of each year (see Section 3.2 for details).

of flights reduces layover time, which then translates into a reduction of travel time. In Appendix Figure 15 we compare the change in travel time from 1951 to 1966 with a counterfactual change in travel time in which we eliminate layover time in both time periods. We observe that the average change in travel time is stronger at every distance in the counterfactual scenario without layover time. This implies that, if any, layover time attenuated the reduction in travel time.

3.2 Constructing an instrument

One endogeneity concern in the estimation stage would be that the change in travel time is not as good as random (i.e. driven by a technological difference between propeller and and jet airplanes) but rather by endogenous decisions of agents (e.g. airlines or the regulator) that may be correlated with our outcomes of interest (i.e. knowledge flows and creation of new knowledge). In this section we construct an instrumental travel time that is based on the preexisting flight routes and the time-varying nationwide roll out of jets. In the instrumental travel time we fix the flight path to 1951 and for each year we simulate the travel time that would have happened if all routes were operated with the *average airplane* of that year, dropping layover travel time.¹⁴ In this way, the instrument abstracts from the endogenous decisions of two agents: First, regulator's decision on the opening/closure of routes. Second, airlines' decision about to which routes allocate jet vs propeller airplanes and scheduling (frequency of flights and layover time). The key source of variation in the instrument is that as the speed of airplanes increases, the importance of the number of stops relative to the distance flown changes, and this provides a decrease in travel time that is larger for MSA-pairs located farther apart and/or with fewer stops. The two identifying assumptions for the instrument to be valid are that 1951 routes do not yet incorporate changes in expectation of the arrival of jets, and that the nationwide roll out of jets is not driven by any single route.

In Borenstein and Rose (2014) it is argued that, due to strict regulation, it was difficult for airlines to adapt their flight network when technology to fly changed. However, we may be concerned that the decision of the regulator to grant new routes could be targeted to specific pairs or correlated with unobservable variables that also affect the creation and diffusion of knowledge.¹⁵ Hence, as the first step in the construction of our instrument, we *fix paths* to the fastest path in 1951.^{16,17}

Airlines could decide on two factors that affect travel time: the type of airplane (jet vs. propeller) operated in each route and scheduling, which consists on the frequency

¹⁴Fixing the flight path between two (indirectly connected) airports to the fastest path in 1951 means that if in 1951 to go from New York City to San Francisco the fastest path included a stop in Chicago, then the instrumental travel time follows the same path New York-Chicago-San Francisco in all years.

¹⁵For example, the regulator could have targeted the opening of new routes between places in order to boost their economic activity.

¹⁶By fixing the path we are also not allowing for the opening/closure of new routes. For example, in the instrument there are no non-stop transcontinental routes.

¹⁷In our estimations we exploit time variation for identification. Hence, if pre-existing routes affect the levels at the origin-destination level, this does not drive our identification. However, we may be concerned that pre-existing routes could affect future growth and not only levels. To address this concern, in robustness analysis Appendix Table 17 we estimate the elasticity of citations to travel time using only MSA-pairs that are always indirectly connected. Results are consistent with baseline estimation.

of flights and layover time in case of connecting flights.¹⁸ We may be concerned that, as with the regulator, airlines' decisions could be correlated with unobservables that also affect the creation and diffusion of knowledge.¹⁹ The second step in the construction of our instrument is to discard layover time (hence discarding all scheduling decisions) in all time periods, and assume that in each year all routes are operated with a *fictitious average airplane* of the year. Hence, the change in instrumental travel time in a route is independent of the type of airplane used in the route and it only depends on the nationwide roll out of jets.

To construct the instrumental travel time, we first estimate a linear regression of travel time on flight distance separately for each year, using only the fastest non-stop flights for each origin-destination airport pair.²⁰ These yearly regressions provide the parameters for the fictitious average airplane of each year: the intercept represents the takeoff and landing time, while the slope indicates the (inverse) cruising speed. Second, we use these regressions to predict travel times for all non-stop flights in each year. Third, we then compute the fastest paths for all airport pairs in 1951 (including indirect connections) using the Dijkstra algorithm based on the 1951-predicted travel times. Fourth, for each year, we calculate the predicted travel time along the fastest 1951 path for all airport pairs.²¹ Therefore, the instrumental travel time reflects across-time variation within the fixed 1951 path. For instance, if the Dijkstra algorithm in 1951 determined that the optimal path from A to B was via C, the instrumental travel times for A-B in subsequent years is calculated as the sum of the fictitious travel times for A-C and C-B obtained in the second step. Layover time is set to zero for all years.

The across-time source of variation in the instrument is the time varying importance

¹⁸In 1961, all non-stop flights of more than 3,000km had at least one jet operating within them, while in 1966 it was the case in all non-stop flights of more than 2,000km. Therefore the endogeneity of jet adoption is a smaller concern for long distance flights.

¹⁹For example, airlines may have decided to prioritize the allocation of jets to routes which had a higher share of business travel, which may be correlated with the diffusion of knowledge.

²⁰These regressions use all routes available in each year. Results, presented in Appendix Table 7, show that the implied average flight speed increased from 412 km/h in 1951 to 453 km/h in 1956, 758 km/h in 1961, and 876 km/h in 1966. Meanwhile, the intercept fluctuates from 25.3 minutes in 1951 to 29.9 minutes in 1966.

²¹Some MSAs have multiple airports. Within each MSA-pair, we use the fastest airport-pair path. The instrumental travel time holds the 1951 airport-pair fixed for each MSA-pair.

of in-flight travel time relative to the number of stops required to go from one MSA to another. In shorter flights, the amount of stops (intercept) have a larger share of the total instrumental travel time and changes in flight speed (slope) have less of an influence. In long distance flights the flight speed becomes more relevant. As estimated flight speed more than doubles over the time period, longer flights have a larger reduction in travel time in the instrument. However, differences in the amount of stops required also leads to variation in changes of travel time conditional on the route distance. Long distance routes and routes with less amount of stops have a larger reduction of travel time in the instrument.²² Appendix Figure 13 shows that this is also a pattern observed in the data.

Figure 1 shows that the instrumental travel time follows pretty closely the observed change in travel time in each year. Especially, it replicates the pattern of a stronger decrease in travel time for MSAs located farther apart. This finding shows that most of the change in travel time that we observe is due to the change in speed of airplanes, and that the endogeneity concern is limited for MSAs located far away from each other.

In Appendix C we present other two counterfactual travel times: one in which we fix airplanes to be the average airplane of 1951 and allow routes to evolve, and another in which both the average airplane and routes are varying. These two counterfactuals together with the one presented in this section allow us to decompose the change in travel time by the change in routes and the change in speed of airplanes. We obtain that around 90% of the change in travel time is due to the change in speed of airplanes, while around 10% of the change is due to the change in the flight routes. Appendix Figure 17 shows that the share is roughly constant for all distances. This finding confirms that most of the observed changes in travel time are due to improvements in flight technology.

²²For example, using the coefficients in Appendix Table 7, the instrumental travel time for a pair of airports located 300 km apart connected non-stop in 1951 and 1966 would be 69.1 minutes and 50.3 minutes respectively, indicating a 27.2% reduction in travel time. For a pair of airports located 2,000 km apart connected non-stop, the instrumental travel time would be 317.3 minutes and 165.9 minutes, indicating a 47.7% reduction in travel time. Assuming both pairs of airports had two intermediary stops in a straight line between the origin and destination in both years (such that the origin-destination distance and travel distance are the same), the reduction in travel time would be 8% for the 300 km apart airports and 38.6% for the 2,000 km apart airports.

4 Patent data

We use patent data as our source of innovation information.²³ We construct a dataset of all corporate patents granted by the United States Patent and Trademark Office (USPTO) with filing year between 1949 and 1968, which includes for each patent: filing year, technology classification, location of the inventors when they applied for the patent, owner of the patent and citations to other patents also granted in the United States.²⁴ This dataset provides the distribution of patents and citations over the geographic space. Details on the construction of the data are provided in Appendix D.

We select and aggregate our patent data sample as follows. We drop patents that are owned by universities or government organizations. To count patents that are classified into multiple technology categories, we do a fractional count by assigning proportionally a part of the patent to each category. Citations are counted as the multiplication of the technology weight of the citing and cited patents. We drop patents (and their citations) that have inventors in multiple MSAs and citations in which the citing owner is the same as the cited owner.²⁵

We aggregate the patent data to 4 time periods of 5 years each, with the center of each period being the year of travel time data collected. The periods are: 1951 (which contains the years 1949-1953), 1956 (1954-1958), 1961 (1959-1963) and 1966 (1964-1968). We consider only patents in MSAs that are matched to an airport in the four periods. The final dataset contains 108 MSAs with patents and travel time.²⁶

²³For a discussion on using patents as a measure of innovation see Fagerberg et al. (2005) Chapter 6 and Carlino and Kerr (2015) Chapter 6.2.3. Feldman (1994) finds that the correlation between the location where new products are introduced to the market and patents is 0.8.

²⁴Filing year, also called application year, is the closest date to the date of invention that is present in the data and it represents the date of the first administrative event in order to obtain a patent. Front page patent citations were made compulsory in 1947 Gross (2019) and there were very few citations prior to that date, limiting the analysis going backwards. At the same time, Appendix Figure 19 shows that the location non-matching rate of patents (likely due to registration changes in the USPTO) increases up to 20% at the beginning of 1970s limiting the analysis going forward.

²⁵3% of patents have inventors in more than one MSA. Working with multi-MSA patents requires an assumption on how to compute distance and travel time between the citing and cited patents, as they do not have a single origin-destination location pair. We hence prefer to abstract from multi-MSA patents.

²⁶Appendix Table 8 shows amount of patents, citations and quartiles of citation distance in each of the steps of the sample selection.

4.1 Descriptive statistics: Patents

This section presents three facts about US patents over our sample period: First, initially less innovative locations had a higher patenting growth rate implying a convergence rate between MSAs in the lowest and highest quartiles of 5.3% per year. Second, the South and the West had a yearly growth rate around 2 percentage points higher than the Northeast and the Midwest, leading to a geographic shift of innovation. Third, the mass of citations shifted towards longer distances, with the third quartile of citation distance increasing by 39%.

We compute descriptive statistics by technology category. In here we present averages across technologies. More information on the presented descriptive statistics and a disaggregation by technology are included in Appendix D.

Fact 1: Initially less innovative locations had a higher patenting growth rate. Figure 2 shows the geographic distribution of patenting in 1951. Darker colors refer to a higher level of *initial innovativeness*, which is defined as the amount of patents filed by inventors in the MSA in 1951.²⁷ We observe that MSAs in the top quartile of patenting are concentrated in the Northeast (which includes New York) and the Midwest (which includes Chicago), with few additional MSAs in the West.²⁸

Figure 3 shows the geographic distribution of patenting growth in 1951-1966.²⁹ We observe a striking pattern relative to Figure 2: high growth MSAs were those that were initially less innovative. High growth happens in initially less innovative locations in the South and the West but also in the Northeast. The average yearly growth rate of MSA-technologies in the lowest quartile of initial innovativeness is 7.2% while it is 1.9%

²⁷To compute the level of initial innovativeness we only use patents filed in 1951 (years 1949-1953). We aggregate patents to the MSA-technology level and then compute the quantile-position of each MSA in the technology. Lower values of quantile-position refers to lower amount of patents in the technology (relative to other MSAs). Each MSA has a different value of quantile-position in each of the 6 technology categories. To obtain the MSA level quantile we take the average quantile across technologies within the MSA. Finally we classify MSAs into quartiles depending on whether the average quantile is higher or lower than the thresholds 0.25, 0.50, 0.75.

²⁸The top 5 patenting MSAs in 1951 were: New York City (25% of all patents), Chicago (11%), Los Angeles (8%), Philadelphia (6%) and Boston (4%).

²⁹We compute the growth rate of patenting in each technology within a MSA and then take the average across technologies within the MSA.

in the highest quartile.³⁰ The percentage point difference between the two growth rates implies that locations in the lowest quartile converged towards locations in the highest quartile at a speed of 5.3% per year.³¹ The convergence in patenting across MSAs is consistent with *The Postwar Decline in Concentration, 1945-1990* described in Andrews and Whalley (2021).

Appendix Figure 24 confirms this pattern presenting the MSA's ranking of innovativeness in 1951 and its subsequent patenting growth rate in 1951-1966. MSAs that were initially more innovative are those that saw lower values of subsequent patenting growth. At the mean, lowering initial innovativeness by 10 positions in the ranking was associated with a subsequent 0.42 percentage points higher yearly growth rate of patenting.



Figure 2: Geography of Patenting 1951 In Figure 2 each MSA is colored according to the amount of patents applied by inventors residing in that MSA in the period 1949-1953, while in Figure 3 according to the growth rate in patents applied 1949-1953 to 1964-1968. 4th Quartile is respectively the MSAs with the largest amount of patents and MSAs with the highest growth rate. Patenting and quartiles are computed only for MSAs that are used in the regression analysis. MSAs for which we do not observe an airline operating in all time periods are marked as Missing.

Fact 2: The South and the West of the US had a higher patenting growth rate.

Figure 3 shows that MSAs located in the South and the West of the US had a higher

³⁰We first compute the 1951-1966 growth rate (19-year growth rate) for each MSA-technology. We then take averages across MSAs within a quartile-technology, and after take averages across technologies within a quartile. Finally, we convert the 19-year growth rate into an average yearly growth rate.

³¹We note that the aggregate growth of patents is much smaller than the across MSAs unweighted average, and this is exactly because initially less innovative MSAs grew faster. If we compute the growth rate in nationwide amount of patents in each of the technologies and then average across technologies we obtain a yearly growth rate of 1.5%.

patenting growth rate in 1951-1966. We classify MSAs using Census Regions of the US (Midwest, Northeast, South and West) and aggregate patents within each region-technology-year.³² The share of patents filed by inventors located in the Midwest and the Northeast decreased from 75% in 1951 to 68% in 1966. In the period 1951-1966 the South and the West increased their amount of patenting by 80%, while the Midwest and the Northeast had a 22% growth.³³ Translated into yearly growth rates, the South and the West grew three times as fast as the Midwest and the Northeast (3.13% vs. 1.08% per year).³⁴

Fact 3: Distance of citations increased. In our analysis we use citations as a proxy for knowledge diffusion. According to Jaffe et al. (1993) *"a citation of Patent X by Patent Y means that X represents a piece of previously existing knowledge upon which Y builds."* (page 580).³⁵ We compute the distance between the citing inventor and the cited inventor. Figure 4 shows the evolution over time of the first, second and third quartile of citation distance.³⁶ We observe that 25% of citations happened between inventors located less than 300km apart throughout our sample period. For the middle 50% of citations we observe that over time inventors cited other inventors located farther away. The third quartile of citation distance increased from 1,590km in 1951 to 2,212km in 1961, a 39% increase in the distance.³⁷ In other words, the mass of citations shifted towards longer distances.

In Figure 5 we present the share of citations by distance range between the citing and cited inventors.³⁸ The distance cutoffs where chosen in order to have a balanced

³²In Appendix E we present a map with the four Census Regions. Some MSAs belong to multiple Census Regions. In here we present descriptive statistics duplicating such MSAs (assigning the MSA to both Census Regions). Statistics dropping such MSAs are quantitatively similar.

³³Growth rates are computed by region-technology and then averaged across technologies within region. ³⁴3.13% $\approx 1.80^{(1/19)} \times 100, 1.08\% \approx 1.22^{(1/19)} \times 100$

³⁵Jaffe et al. (1993) discusses the reasons why to cite and why not to cite. Using a survey of inventors, Jaffe et al. (2000b) find that there is communication among inventors and citations are a "noisy signal of the presence of spillovers."

³⁶We compute distance between MSA centroids.

³⁷As a reference, the distance from New York City NY to other places is: Boston MA 300km, Chicago IL 1,140km, Dallas TX 2,200km, San Francisco CA 4,130km. The quantile 0.10 of was at 0km in every period, implying that 10% of citations took place within MSA. The quantile 0.90 was between 3,538km and 3,716km over the sample period.

³⁸While Figure 4 shows how the distance of each quantile changes over time, Figure 5 shows the mass of citations (and hence the quantile to which belongs) in a certain distance cutoff. For example, in 1951

share of citations in the initial time period, and considering the changes in travel time presented in Section 3.1. The share of citations that happen between inventors located more than 2,000km apart grew from 20.9% in 1951 to 27.1% in 1966. The 6.2 percentage points increase represents an increase of 30% of the share of citations at more than 2,000km.



Figure 4: Quantiles of citation distance

Figure 5: Share of citations by distance

5 Diffusion of knowledge

In this section we show that the reduction in travel time led to an increase in knowledge diffusion, especially over long distances. In doing so we estimate the parameter β highlighted in equation (2): the elasticity of knowledge diffusion to travel time.

To perform the analysis we merge the Air Travel and Patent datasets to obtain a final dataset that contains for each patent owner-location, the amount of patents filed in a certain 5-year period and technology class, the amount of citations to other patents with their respective owner identifier, location and technology class, and the travel time to every location. We label a patent owner a *firm* and call *research establishment* a firm-MSA pair for MSAs in which it has inventors applying for patents. We aggregate

the share of citations in the 0-300km range was 31.6%, which is equal to saying that the quantile 0.316 in 1951 was 300km.

citations to the citing-cited establishment-technology within each period. We assume that passengers take a return flight, hence we make travel times symmetric.³⁹

5.1 Diffusion of knowledge: Baseline estimation

We estimate a gravity equation which relates citations between two establishmentstechnologies with their pairwise travel time.⁴⁰ We estimate the following regression:

$$citations_{FiGjhkt} = \exp \left[\beta \log(\text{travel time}_{ijt}) + FE_{FiGjhk} + FE_{Fiht} + FE_{Gjkt}\right] \times \varepsilon_{FiGjhkt}$$
 (1)

where *citations*_{*FiGjhkt*} is the amount of citations from patents filed by the establishment of firm *F* in location *i*, technology *h* and time period *t*, to patents filed by establishment of firm *G* in location *j* and technology *k*. We call *Fi* the research establishment of firm *F* in location *i*. travel time_{*ijt*} is the air travel time (in minutes) between location *i* and *j* at time period *t*. The parameter of interest in the regression is β , which represents the elasticity of citations to travel time.⁴¹ If citations are affected negatively by travel time we would expect a negative value of β .

Given the panel structure of our data, we can include the fixed effect FE_{FiGjhk} that absorbs any time invariant citation behavior within the *citing establishment-technology and cited establishment-technology*. This fixed effect flexibly controls for persistent relationships within an establishment pair that would lead to relatively more (or less) citations. That includes characteristics like physical distance, but also pre-existing commercial relationships between establishments. The fixed effects FE_{Fiht} and FE_{Gjkt} control for the time changing general level of citations specific to each establishment and technology. For example FE_{Fiht} controls for the fact that if *Fih* files more patents

³⁹travel time_{ijt} = $\frac{travel time_{ijt}^{original} + travel time_{jit}^{original}}{2}$ where travel time_{ijt} stands for the travel time between MSA *i* and *j* at time period *t*.

⁴⁰For explanation and micro foundations of the gravity equation see Head and Mayer (2014) and references thereof. While variation in travel time is at the MSA-pair level, we estimate the regression at the more granular level of the establishment-pair as this allows to control for establishment-level shocks.

⁴¹A 1 percent increase in travel time has an effect of β percent increase (or decrease in the case of a negative β) in citations.

in a given period, it would mechanically make more citations to every establishment. On the other hand, FE_{Gjkt} controls for Gjk filing more patents or higher quality patents that would receive more citations from every establishment.⁴²

The inclusion of FE_{FiGjhk} implies that only variation across time within an establishmentpair is used for identification. By additionally including the fixed effect FE_{Fiht} , the across-time variation is compared only between citing-cited establishment-technology pairs FiGjhk within a citing establishment-technology Fih in period t. As we also include FE_{Gjkt} , the comparison is done while controlling for the size of the cited establishment-technology Gjk in period t. Put differently and simplifying slightly, the identification of β relies on changes in citations and travel time within an establishmentpair, relative to another establishment-pair with the same citing establishment, conditional on the two cited establishments' sizes.

Following Silva and Tenreyro (2006), we estimate the gravity equation by Poisson Pseudo Maximum Likelihood (PPML).⁴³ This estimation methodology has two advantages over a multiplicative model that is then log-linearized to obtain a log-log specification. First, it only requires the conditional mean of the dependent variable to be correctly specified, while the OLS estimation of the log-linearized model would lead to biased estimates in the presence of heteroskedasticity. Second, it allows to include zeros in the dependent variable, which is especially relevant when using disaggregated data. One downside of estimating PPML with the fixed effects that we include is that both coefficients and standard errors have to be corrected due to the incidental parameter problem (Weidner and Zylkin (2021)). We follow Weidner and Zylkin (2021) to use split-panel jackknife bias-correction on the coefficients and Dhaene and Jochmans (2015) to bootstrap standard errors which we also bias-correct with split-panel jackknife.⁴⁴

Whenever *FiGjhk* has positive citations in at least one period and missing value in another, we impute zero citations in the missing period.⁴⁵ Travel time is set to one

⁴²In the International Trade literature, the parallel of the fixed effects (simplified for exposition) would be: FE_{ij} country-pair fixed effect, FE_{jt} origin-time fixed effect and FE_{it} destination-time fixed effect.

⁴³We use the package *fixest* (Bergé (2018)) in R to estimate high dimensional fixed effects generalized linear models *feglm* with Poisson link function.

⁴⁴Details on the bias correction and bootstrap procedures are provided in Appendix F.

⁴⁵We do not impute zeros in *FiGjhk* that are always zero, as those observations would be dropped due

	PP:	ML	IV PPML			
Dep. variable:	citations					
	(1)	(2)	(3)	(4)		
log(travel time)	-0.084^{***}		$-0.161^{***}_{(0.031)}$			
log(travel time) \times 0-300km		-0.015		-0.185 $_{(0.153)}$		
log(travel time) \times 300-1,000km		-0.085^{***}		-0.155^{***}		
log(travel time) \times 1,000-2,000km		-0.096^{***}		-0.132^{**}		
log(travel time) \times +2,000km		$-0.166^{***}_{\scriptscriptstyle{(0.035)}}$		$-0.206^{***}_{(0.042)}$		
Control residuals 1st stage	-	-	Yes	Yes		
N obs. effective	5,147,161	5,147,161	5,147,161	5,147,161		
R2	0.88	0.88	0.88	0.88		

***p < 0.01; **p < 0.05; *p < 0.10

Table 1: Elasticity of citations to travel time

Column (1) shows the result of Poisson Pseudo Maximum Likelihood (PPML) estimation of *citations*_{*FiGjhkt*} = exp [βlog (travel time_{*ijt*}) + *FE*_{*FiGjhk*} + *FE*_{*Fiht*} + *FE*_{*Gjkt*}] × $\varepsilon_{FiGjhkt}$, for citations of patents filed by establishment of firm *F* in location *i*, technology *h* and time period *t*, to patents filed by establishment of firm *G* in location *j* and technology *k*. travel time_{*ijt*} is the travel time in minutes between location *i* and *j* at time period *t*, and it is set to 1 when *i* = *j*. When *FiGjhk* has positive citations in at least one period and no citations in another, we attribute zero citations in the missing period. Column (2) includes the interaction of travel time_{*ijt*} with a dummy for distance bin between the citing establishment *Fi* and the cited establishment *Gj*. Column (3) and (4) show the result of two step instrumental variables estimation, where *log*(travel time_{*ijt*}) is instrumented with *log*(travel time_{*ijt*}), the ficitious travel time that would have taken place fixing the flight path of 1951 and operating in each year with the *average airplane* of the year. Columns (3) and (4) include as controls residuals of first stage. Bootstrap standard errors are presented in parentheses. The coefficients and standard errors in columns (1) and (2) are jackknife bias-corrected. R2 is computed as the squared correlation between observed and fitted values.

minute whenever $i = j.^{46}$

Column (1) in Table 1 presents the results of estimating equation (1). The value of the elasticity of citations to travel time is estimated to be -0.084, statistically significant at the 1% level. Given the average reduction in travel time of 31.4% in the full estimating sample, the elasticity implies that citations increased on average 2.6% as consequence of the reduction in travel time. If we consider the average decrease in travel time across all MSAs in the baseline travel time data, the implied increase is 2.4%.⁴⁷

The importance of air transport relative to other means of transport potentially depends on the distance to travel. Also, we observed in Section 3.1 that the improvements in air travel time depended on the distance to travel, with a difference in jet adoption for travel distances under and over 2,000km. Taking these two characteristics into account, we estimate a variation of equation (1) in which we allow the elasticity of citations to travel time to vary by distance interval between the locations of citing and cited establishments.⁴⁸ Column (2) in Table 1 shows the result of this estimation.⁴⁹ The estimated value of the elasticity in absolute terms increases with distance, reaching -0.166 for distances of more than 2,000km. Between 1951 and 1966 the average change in travel time in the full estimating sample is 47.8% for a distance of more than 2,000km. The estimated elasticity implies that citations between establishments at more than 2,000km apart increased by 7.9% due to the decrease in travel time. In total citations at more than 2,000km increased by 22.9%, implying that the change in travel time can account for 34.6% of the observed increase. If instead we consider the 40.8% average

to not being able to identify FE_{FiGihk} .

⁴⁶We measure air travel time in minutes. In our sample 14.9% of citations happen within the same MSA. The inclusion of those citations in the estimation increases the amount of observations available to identify of FE_{Fiht} and FE_{Gjkt} , and hence keeping them increases the amount of FiGjhkt that remain in the effective sample to identify β . In order to include them we then need to impute a within-location travel time. We assume that within-location (air) travel time is not changing across time periods. Nonetheless, the identification of β is not affected by the value chosen for the within-location (time invariant) travel time, as β is identified by across time variation. In the appendix we show results using other values of (time invariant) within MSA travel time and the coefficients remain equal.

⁴⁷These values come from the multiplication of the elasticity of citations to travel time -0.084 and the average change in travel time between 1951 and 1966: -31.4% in the full estimating sample and -28.7% in the raw data of travel time across MSAs.

⁴⁸We compute distance between the geographical center of each MSA.

 ⁴⁹The share of observations (citations) in each distance interval in the effective sample is: 0-300km 25.4% (27.5%), 300-1,000km 32.5% (30.1%), 1,000-2,000km 19.0% (19.1%), +2,000km 23.1% (23.4%).

reduction in travel time across MSAs in the baseline data, the elasticity implies an increase in citations of 6.8%, accounting for 29.7% of the total citation increase.

In Appendix D we investigate different heterogeneous effects. We estimate an heterogeneous elasticity depending on the level of spatial concentration of the citing technology and the cited technology, we do not find a statistical difference. We also look at whether it is older patents or younger patents that get diffused, finding some slight evidence that it is technologies that take longer time to diffuse that increase more their diffusion with the reduction in travel time. We study citations to and from government patents, and self citations, on the whole we do not find a different pattern from the baseline. We also do not find a particular pattern of the elasticity depending on the citing *firm's size* as measured by the amount of patents filed in 1949-1953. Finally, we estimate the elasticity by citing and cited technology and most of the effect seems to come when the citing and cited technologies are the same.

There are two types of threats to identification in estimating equation (1): (i) the potentially targeted changes in travel time, which could be due to the opening of new routes, the allocation of jets across routes, or changes in scheduling, and (ii) time changes in other variables at the MSA-pair level which also drive the diffusion of knowledge and are correlated with the changes in travel time. In the next sub-section we address the first type of threat by estimating the model by instrumental variables. In the following subsection we address the second type of threat by adding multiple controls. In both cases we show that results do not qualitatively change.

5.2 Diffusion of knowledge: Instrumental variables estimation

As mentioned in Section 3.2, we may be concerned that the timing and allocation of jets to routes and that the opening/closure of routes were not random. In case there is an omitted variable that drives both the change in travel time at the MSA-pair level and the change in citations across establishments within the same MSA-pair, we would estimate biased coefficients. In order to tackle the endogeneity concern due to omitted variable we do an instrumental variables estimation using the instrument proposed in

Section 3.2. To implement the instrumental variables estimation we follow a control function approach described in Wooldridge (2014). We proceed in two steps estimating the following two equations:

$$log(travel time)_{FiGjhkt} = \lambda_2 log(instrumental travel time_{FiGjhkt}) + FE_{FiGjhk} + FE_{Fiht} + FE_{Gjkt} + u_{FiGjhkt}$$
(2)

$$citations_{FiGjhkt} = \exp \left[\beta \log(\text{travel time}_{ijt}) + \lambda \,\hat{u}_{FiGjhkt} + FE_{FiGjhk} + FE_{Fiht} + FE_{Gjkt}\right] \times v_{FiGjhkt}$$
(3)

In a first step we estimate equation (2) and obtain estimated residuals $\hat{u}_{FiGjhkt}$. In a second step we use the estimated residuals as a regressor in equation (3) which *controls* for the endogenous component of travel time. To perform inference we bootstrap standard errors.⁵⁰

Columns (3) and (4) of Table 1 show the results of the instrumental variables estimation. If airlines were allocating jet airplanes to routes that would have witnessed a higher degree of exchange of knowledge even in the absence of jets, then we would expect the instrumental variables estimate to be smaller in absolute terms relative to the baseline coefficient. On the other hand, if the regulator targeted the opening of new routes between places that were in a lower trend of exchange of knowledge, we would expect the instrumented coefficient to be larger in absolute terms. Column (3) estimates the elasticity to be -0.161, bigger in absolute value compared to the non-instrumented estimate. The instrumental variables corrects for a downward bias in absolute terms, which represents evidence in favor of the regulator targeting the opening of new routes between places that had a lower degree of exchange of knowledge.^{51,52}

⁵⁰Appendix F includes details on the bootstrap procedure.

⁵¹The incidental parameter problem is potentially present also in the instrumental variables estimation (IV PPML). However, we are not aware of any bias-correction procedure for IV-PPML. Hence, columns (3) and (4) in Table 1 are not bias-corrected. In column (2) of Table 2 we present the PPML estimation not bias-corrected.

⁵²The literature on weak instruments for non-linear instrumental variables is scarce. The rule of thumb of Staiger and Stock (1997) based on the F statistic is constructed using the bias that a *weak instrument* generates in a linear second stage (see Staiger and Stock (1997), Stock and Yogo (2005) and Sanderson and Windmeijer (2016) for testing for weak instruments in linear IV regression). For informative purposes, in the first stage of the model estimated in column (3) in Table 1 we obtain $\hat{\lambda}_2 = 0.91$ with a

Column (4) of Table 1 presents the coefficients of the instrumental variable estimation by distance between the citing and cited establishments. We observe the presence of a bias in the same direction as in column (3), however the magnitude of the bias is smaller except for the distance bin 0-300km, which is not precisely estimated. In particular, at more than 2,000km, the coefficient is relatively similar to the baseline estimation. Appendix G Tables 14 and 15 present the regression results including coefficients on the residual *controls*. According to Wooldridge (2014), there would be evidence of endogeneity if the parameter λ on controls in equation (3) is estimated to be statistically different from zero. While the control is statistically significant when using only one coefficient for all distances, none of them is statistically significant when opening the coefficient by distance range. Especially, we do not find evidence of endogeneity at +2,000km.

The instrument used in the instrumental variables estimation is constructed using the 1951 flight network. We may be concerned that the 1951 flight network is correlated with future changes of citations.⁵³ In order to address this concern in Appendix G we estimate equation (1) by restricting the sample to establishments in MSA-pairs that are always indirectly connected. Results go in the same direction.

5.3 Diffusion of knowledge: Robustness

We may be concerned that there are other variables that could drive the diffusion of knowledge and at the same time be correlated with the change in travel time. In order to bias the coefficients, such omitted variables should be time-changing at the origin-destination MSA pair and be systematically correlated with the change in MSA-pair air travel time.⁵⁴ We consider three potential variables that could bias our estimates: improvements in highways, improvements in telephone communication and changes

standard error 0.042 (clustered at the non-directional MSA-pair level, *ij* is the same location pair as *ji*), and a within R2 of 0.34 (the share of residual variation explained by the instrument, after projecting out fixed effects). See Appendix Tables 14 and 15 for results of first stage estimations.

⁵³We include a establishment pair fixed effect in the regressions, so a potential correlation between the 1951 flight network and the level of citations between research establishments does not affect our estimation.

⁵⁴Variables that are not time changing or that are time changing at the MSA or establishment level do not represent a threat to identification, as they are flexibly controlled for with the fixed effects.

in flight ticket prices. In Table 2 we show the results controlling for this variables separately, while in Appendix G we include them simultaneously. Estimates are robust to including these controls.

Columns (1) and (2) in Table 2 present the elasticity of citations to travel time by distance bin. In column (1) the elasticity is bias-corrected while in column (2) it is not. We observe that not doing the bias correction does not qualitatively affect the results. Columns (3) to (6) include the additional controls and should be compared to column (2).⁵⁵

First, in 1947 the Congress published the official plan for the Interstate Highway System, a nationwide infrastructure plan to improve existing highways and build new ones (see Baum-Snow (2007), Michaels (2008), Jaworski and Kitchens (2019) and Herzog (2021)). In case the change in travel time by air is correlated with the change in travel time by highway, we would have an omitted variable bias if we include only one of them in the estimation. Taylor Jaworski and Carl Kitchens have graciously shared with us data on county-to-county highway travel time and travel costs for 1950, 1960 and 1970, which we converted to MSA-to-MSA and linearly interpolated to convert to the same years of our air travel data. Hence we have a MSA-to-MSA time-varying measure of highway travel time which we include as control.⁵⁶

Second, other means of communication like telephone lines may have expanded or changed their price during the period of analysis. Haines et al. (2010) contains information on the share of households within each city with telephone lines in 1960. We aggregate the variable to the MSA level. For each MSA-pair, we take the log of the mean share of households with telephone lines.⁵⁷ To include the variable as control we interact it with a time dummy to make the measure time variant. The assumption

⁵⁵The jackknife bias-correction due to the incidental parameter problem is computationally intensive. Due to the computational burden and given that the bias correction does not substantially change the results in the baseline analysis, we have not bias-corrected estimations of robustness analysis in columns (2) to (6) of Table 2.

⁵⁶In Appendix G we show the correlation of MSA-to-MSA change in air travel time and highway travel time.

⁵⁷Data from the 1962 City Data Book which comes from the US Bureau of the Census. $log(mean \ telephone \ share_{ij} = log((telephone \ share_i + telephone \ share_j)/2)$. Using as control the *multiplied telephone share* = $telephone \ share_i \times telephone \ share_j$ gives similar results.

behind the interaction is that, if telephone lines expanded or changed their price over the time period, this time-change specific to each year was proportional to the 1960 log mean share of the MSA-pair.

Third, during the period of analysis ticket prices were set by the Civil Aeronautics Board, so airlines could not set prices of their own tickets. Some airlines included a sample of prices in the last page of their booklet of flight schedules, which we digitized. In Appendix G we document multiple facts about prices. The relevant fact for this section is that during 1962-1963 we observe a drop in prices of around 20% for routes of more than 1,000km distance. We may be concerned that the change in flow of knowledge is actually consequence of the change in prices, which happens to be correlated with the change in travel time. Given that we do not have ticket prices for each route and year, we use an estimated route price which is time varying. We obtain estimated prices by using the sample of prices that we digitized and fitting, for each year, price on a third degree polynomial of distance between origin and destination. We use log of estimated prices as control.⁵⁸

Column (3) to (5) of Table 2 include the described controls. All of the coefficients are in the ball park of the baseline coefficients in both columns (1) and (2).⁵⁹

Fourth, we control for a time varying effect of distance on citations. We may believe that other variables may have an effect on the diffusion of knowledge, and those variables are related to the distance between the citing and cited establishments. In column (6) we include as control log(distance) interacted with a time dummy. We observe that the coefficients reduce in magnitude, potentially due to the fact that the change in travel time is also correlated with distance, hence controlling for a time-varying effect of distance absorbs part of the effect. In spite of that, the coefficient for distance of more than 2,000km remains statistically significant at the 5% level. This result highlights the importance of the origin-destination time varying travel time data when studying the impact of face to face interactions, pointing that travel time and

⁵⁸Standard errors presented are not adjusted by the fact that the regression includes a predicted regressor as control variable.

⁵⁹Assuming the covariance across coefficients of different regressions is zero, none of the coefficients is statistically different from the baseline coefficients either in column (1) or (2).

distance are not equivalent measures. This result differentiates our analysis from the one of Feyrer (2019) who uses two types of time-invariant distance (sea distance and geographical distance) interacted with time dummies to study changes in international trade.

In Appendix G we present additional robustness analysis. We may be concerned that the change in diffusion of knowledge is only consequence of the change in the geographic location of innovation. Hence, we re-estimate equation (1) with different samples: first, using only citing establishments that were present in 1949-1953, and second using only citing and cited establishments that were present in 1949-1953. We find that across sub-samples the coefficient at more than 2,000km remains stable across samples and statistically significant at the 1% level. Next, we estimate a variation of equation (1) in the form of log-log and obtain results that are in the ballpark of the baseline estimation.⁶⁰

6 Creation of knowledge

In this section we interpret the results on increased diffusion of knowledge through the lens of a model of knowledge spillovers. We show that the reduction in travel time to innovative locations led to an increase in knowledge creation. The effect on the creation of knowledge was stronger in initially less innovative locations, leading to convergence across locations in terms of innovation. Additionally, the reduction in travel time contributed to a change in the geographic distribution of knowledge creation, increasing the relative importance of locations in the South and the West of the United States.

In Appendix section A, we present a conceptual framework that guides our empirical specifications. This framework introduces a production function of knowledge with external returns in the form of knowledge spillovers. It shows how a measure of potential knowledge spillovers can be constructed using the elasticity of knowledge diffusion to a barrier to knowledge diffusion, in our case, travel time. Based on this

⁶⁰See Appendix G Table 18. Elasticity at +2.000km is estimated to be -0.161 by OLS.

	PPML		PPML			
	bias-corrected		no			
Dep. variable: <i>citations</i>	<i>citations</i> _{FiGihkt}					
	(1)	(2)	(3)	(4)	(5)	(6)
log(travel time) \times 0-300km	-0.015	$\underset{(0.041)}{-0.014}$	-0.013	$-0.016 \atop (0.041)$	$\underset{(0.041)}{-0.011}$	-0.003
log(travel time) \times 300-1,000km	$-0.085^{***}_{\scriptscriptstyle{(0.024)}}$	-0.095^{***}	-0.091^{***}	-0.091^{***}	-0.096^{***}	-0.066^{**}
log(travel time) \times 1,000-2,000km	-0.096^{***}	-0.092^{**}	-0.086^{**}	$-0.075^{st}_{ m (0.041)}$	-0.100^{**}	-0.037
log(travel time) \times +2,000km	$-0.166^{***}_{\scriptscriptstyle{(0.035)}}$	$-0.177^{***}_{(0.048)}$	$-0.170^{***}_{(0.050)}$	$-0.167^{***}_{(0.049)}$	$-0.185^{***}_{(0.048)}$	$-0.112^{**}_{(0.056)}$
N obs. effective	5,147,161	5,147,161	5,147,161	5,147,161	5,147,161	5,147,161
R2	0.88	0.88	0.88	0.88	0.88	0.88
Controls:						
log(highway time)	-	-	Yes	-	-	-
log(telephone share) $ imes$ time	-	-	-	Yes	-	-
log(price)	-	-	-	-	Yes	-
$\log(distance) \times time$	-	-	-	-	-	Yes

***p < 0.01; **p < 0.05; *p < 0.10

Table 2: Robustness: Elasticity of citations to travel time

Column (1) shows the result of Poisson Pseudo Maximum Likelihood (PPML) estimation of *citations*_{*FiGjhkt*} = exp $[\sum_d \beta_d \mathbb{1}\{distance_{ij} \in d\} \log(\text{travel time}_{ijt}) + FE_{FiGjhk} + FE_{Fiht} + FE_{Gjkt}] \times \varepsilon_{FiGjhkt}$, for citations of patents filed by establishment of firm *F* in location *i*, technology *h* and time period *t*, to patents filed by establishment of firm *G* in location *j* and technology *k*. travel time_{ijt} is the travel time in minutes between location *i* and *j* at time period *t*, and it is set to 1 when i = j. *d* are distance intervals: [0 - 300km], (300km - 1000km], (1000km - 2000km], (2000km - max]. Column (1) presents jackknife bias-corrected coefficients and bias-corrected bootstrap standard errors. Column (2) repeats column (1) without bias-correction. Relative to (2), columns (3) through (6) contain additional controls. Column (3) controls for log highway time between *i* and *j* at period *t*. Column (4) controls for the log of the mean share of households with telephone line in 1960 in *ij* pair interacted with a time dummy. Column (5) controls for log flight ticket price between *i* and *j* at period *t*. Column (6) controls for log distance *ij* interacted with a time dummy. When *FiGjhk* has positive citations in at least one period and no citations in another, we attribute zero citations in the missing period. Columns (2) through (6) present standard errors clustered at the non-directional location in parentheses (*ij* is the same non-directional location pair as *ji*). R2 is computed as the squared correlation between observed and fitted values.

framework, we first construct an empirical measure of potential knowledge spillovers, which we label *Knowledge Access*, and then estimate the elasticity of new knowledge to knowledge access.

We construct a measure of *Knowledge Access* (KA_{iht}) that shows how *easy* it is in time period *t* for research establishments in location *i* and technology *h* to access knowledge created in other locations.⁶¹ We compute *Knowledge Access* as follows:

$$KA_{iht} = \sum_{k} \omega_{hk} \sum_{j, j \neq i} \text{Patent stock}_{jk,t=1953} \times \text{travel time}_{ijt}^{\beta}$$
(4)

where, from right to left, travel time^{β}_{*ijt*} is the travel time between locations *i* and *j* at time period *t*, to the power of the elasticity of diffusion of knowledge to travel time. Patent stock_{*jk*,*t*=1953} is the discounted sum of patents produced in location *j* and technology *k* between 1941 and 1953.^{62,63} ω_{hk} is the share of citations of technology *h* that go to technology *k* at the aggregate level in 1949-1953, similar to an input-output weight.⁶⁴ Then, *KA*_{*iht*} is a weighted sum of the patent stock in each other location and technology, where the weights are how easy it is to access that patent stock (travel time^{β}_{*ijt*}) multiplied by how relevant that knowledge is (ω_{hk}).

In order to reduce concerns of potential endogeneity of accessing knowledge and creating knowledge, we exclude the patent stock in the location itself from the sum (we only use $j \neq i$).⁶⁵

⁶¹This measure is an extension of Appendix equation (2) to an include multiple technology categories and time periods.

⁶²Patent stock_{*jk*,*t*=1953} = $\sum_{y \in [1941, 1953]}$ Patents_{*jky*} × (1 – depreciation rate)^{1953-y}. We use a depreciation rate of 5%, which is in the range of average depreciation rates of R&D found by De Rassenfosse and Jaffe (2017).

⁶³Location *j* and technology *k* would be the *source* location and technology, while *i* and *h* would be the *destination* location and technology.

 $^{^{64}\}omega_{hk} = citations_{hk,t=[1949,1953]}/citations_{h,t=[1949,1953]}$ is included to weight each *source* technology category *k* by how important it is for the *destination* technology category *h*.

⁶⁵The theory makes no distinction on whether the knowledge stock is in *i* or *j*, so in principle we would like to include the patent stock of *i* in the knowledge access of *i*. However, this could lead to econometric problems. First, we do not have exogenous variation of travel time within *i*. Second, if knowledge creation in *i* is a persistent process, by including the patent stock of *i* we would introduce a mechanical relationship between knowledge access and knowledge creation. Hence, our baseline measure of knowledge access of *i* does not consider the patent stock of *i*. This is similar to what Donaldson and Hornbeck (2016) in the case of the empirical approximation of their Market Access measure.

The measure of *Knowledge Access* contains across-time variation within a locationtechnology *ih*, and cross-sectional variation across technologies *h* within a location *i*. The across-time variation is only due to the change in travel time between locations, every other component of the measure is fixed to its 1949-1953 level. The cross-sectional variation comes from a distribution of Patent stock_{*jk*,*t*=1953} within source technologies *k* that is not uniform across source locations *j*, and from the input-output weights ω_{hk} . The joint across-time and cross-sectional variation means that if travel time for *ij* reduces, there will be a differential change in *Knowledge Access* across technologies *h* within location *i* which depends on the initial patent stock and input-output weights.

The degree with which changes in travel time are reflected in access to knowledge depend on how *important* travel time is to get knowledge to diffuse, which is the elasticity of knowledge diffusion to travel time that we estimated in Section 5. As the baseline we use $\beta = -0.206$, which is the elasticity of citations to travel time at more than 2,000 km estimated by IV-PPML. In Appendix Table 27 we do robustness analysis with distance-specific β and in Appendix Table 30 we do sensitivity analysis of the results to changing the value of β .

The measure of *Knowledge Access* allows us to translate changes in travel time between pairs of MSAs into a single location-technology specific characteristic, and to represent it on the same scale as patent growth in Figure 3. Figure 6 depicts the time change in log *Knowledge Access* from 1951 to 1966, averaged across technologies within each MSA. Darker colors represent higher growth in *Knowledge Access*. As with patent growth, we observe that MSAs that had the strongest growth are generally located in the South and the West of the United States, far from the knowledge centers of New York and Chicago. The reduction in travel time was larger between locations far apart, implying that locations which happened to be far from knowledge centers increased relatively more their *Knowledge Access*.



Figure 6: Change in log Knowledge Access 1951 - 1966

6.1 Creation of knowledge: Baseline estimation

With the measure of *Knowledge Access* we then adapt Appendix equation (1) to estimate:

$$Patents_{Fiht} = \exp\left[\rho \log(KA_{iht}) + FE_{Fih} + FE_{it} + FE_{ht}\right] \times \xi_{Fiht}$$
(5)

where Patents_{*Fiht*} are patents applied by establishment of firm *F* in location *i* and technology *h* at time period *t*. The measure of knowledge access KA_{iht} is at the *iht* location-technology-time level, meaning that all establishments within an *iht* share the same level of knowledge access. The parameter of interest ρ is the elasticity of (the creation of new) patents to knowledge access. In the presence of knowledge spillovers as presented in the conceptual framework in Section A, we would expect ρ to be positive and statistically significant.

The fixed effect FE_{Fih} absorbs time invariant characteristics at the firm-locationtechnology level, as for example the productivity of the establishment-technology. This fixed effect is more fine grained than just a location-technology, which would absorb the comparative advantage of a location in a certain technology. The fixed effect FE_{it} absorbs characteristics that are time variant at the location level. For example, an inflow of population that shift demand of technology, or changes in R&D subsidies that are location specific and common across all technologies would be absorbed by this fixed effect. Also, better flight connectivity could spur economic activity as shown in Campante and Yanagizawa-Drott (2017), leading to an increase in patenting activity in the location. While this could potentially be a channel through which jet airplanes affected innovation, this channel wouldn't be variation used for identifying the coefficient of interest. If that increase is general across technologies within the location, then FE_{it} would absorb it. Finally, the fixed effect FE_{ht} absorbs characteristics that are time variant at the technology level. If technologies had different time-trends at the national level, then the fixed effect would control for these trends in a flexible way.

The inclusion of FE_{Fih} implies that only across-time variation within an establishmenttechnology is used to identify ρ . The inclusion of FE_{it} implies that only variation acrosstechnologies within a location-time is exploited, so across-time variation is compared across establishments within a location, and not across locations. The inclusion of FE_{ht} implies that the identifying across-time variation is conditional on aggregate trends of the technology. In short, identification of ρ relies on across-time changes in the amount of patents and knowledge access of an establishment, relative to other establishments in the same location, conditional on aggregate technological trends.

Column (1) in Table 3 shows the result of estimating equation (5). The elasticity of patents to knowledge access is estimated to be 9.11, significant at the one percent level. The average change in knowledge access at the location-technology level is 10.0%, implying that on average the change in travel time predicts a 3.5% average yearly growth rate of patents.^{66,67} The observed average yearly growth rate of new patents at the location-technology is 4.5%.⁶⁸ Comparing the predicted and observed growth rates,

⁶⁶Due to entry, we cannot compute the growth rate at the establishment-technology level for 71% of establishment-technology, given that they had 0 patents in the initial time period. In the case of location-technology, 4% did not have patents in the initial period. We the prefer to interpret coefficients using location-technology growth rates, which we compute using the remaining 95% of location-technologies that had positive patents in the initial time period.

⁶⁷The elasticity of 9.11 predicts an increase of 91.3% over the time period of 19 years (9.11 × 0.10 \approx 0.913, without rounding in intermediate steps), which translates into a 3.5% average yearly growth rate ((1+0.913)^{1/19}-1 \approx 0.035).

⁶⁸From the first time period (1949-1953) to the last time period (1964-1968) we observe an average growth

	PP]	ML	IV PPML		IV PPML centered	
Dependent Variable:	Patents					
	(1)	(2)	(3)	(4)	(5)	(6)
log(knowledge access)	9.11*** (3.29)	8.41** (3.31)	$10.20^{*}_{(5.81)}$	9.39 (5.85)	10.22** (5.06)	7.62 (5.12)
log(knowledge access) \times 3rd quartile		$1.86^{***}_{(0.53)}$		$2.10^{***}_{(0.60)}$		3.54*** (1.11)
log(knowledge access) \times 2nd quartile		$3.42^{***}_{(0.81)}$		3.79*** (0.75)		6.89*** (2.09)
$\log(\text{knowledge access}) \times 1$ st quartile		4.50*** (1.17)		5.20*** (1.01)		8.05*** (2.24)
R2	0.85	0.85	0.85	0.85	0.85	0.85
N obs. effective	991,480	991,480	991,480	991,480	991,480	991,480

***p < 0.01; **p < 0.05; *p < 0.10

Table 3: Effect of knowledge access on patents, by MSA innovativeness quartile

Column (1) shows the result of Poisson Pseudo Maximum Likelihood (PPML) estimation of Patents_{*Fiht*} = $\exp \left[\rho \log(KA_{iht}) + FE_{Fih} + FE_{it} + FE_{ht}\right] \times \xi_{Fiht}$, for patents filed by establishment of firm *F* in location *i*, technology *h* and time period *t*. KA_{iht} is knowledge access of establishments in location *i* technology *h* and time period *t*. Column (2) opens the coefficient ρ by the quartile of innovativeness of location *i* within technology *h*, computed using patents filed in 1949-1953. Higher quartile indicates higher initial level of innovativeness. The fourth quartile is used as reference category. Columns (3) and (4) show the result of two step instrumental variables estimation, where KA_{iht} is instrumented with \widetilde{KA}_{iht} , knowledge access computed using the counterfactual travel time that would have taken place if routes were fixed to the ones in 1951 and each year routes were operated at the average airplane of the year. Columns (5) and (6) use a centered version of \widetilde{KA}_{iht} following Borusyak and Hull (2023) by subtracting the expected instrument which is computed using random flight networks. Standard errors are presented in parentheses. Column (1) and (2) present clustered at the location-technology *ih*. Column (3) and (4) present bootstrap standard errors. R2 is computed as the squared correlation between observed and fitted values.

the improvement in air travel time has the power to account for 78% of the observed average yearly patent growth rate.⁶⁹

We aggregate predicted changes in patent growth at the Census Region level. The change in travel time predicts a yearly growth rate 0.75 percentage points higher in the South and the West relative to the Midwest and Northeast. In the data we observe 2.04 percentage points difference in the growth rate, implying that the change in travel time can account for 37% of the observed differential growth rate.⁷⁰

Section 4.1 showed that in the data, initially less innovative MSAs had a larger growth rate of patenting. In column (2) in Table 3 we investigate if the increase in knowledge access had an heterogeneous effect on the amount of new patents created depending on the initial innovativeness of the location *i* in technology *h*. We compute the quartile of innovativeness of location *i* in technology *h* in the time period 1949-1953 and interact it with $\log(KA_{iht})$.⁷¹ We use as reference category the highest quartile of initial innovativeness, hence the coefficient on $\log(KA_{iht})$ without interaction is the elasticity for the highest quartile. Coefficients on other quartiles should be interpreted relative to the highest quartile.

We find that the coefficients on lower quartiles of initial innovativeness are positive and statistically different from the coefficient in the highest quartile. Thus, knowledge access had a greater effect on patenting for establishments that were located in initially less innovative locations.⁷² Given the difference in the coefficients, the increase in knowledge access predicts an average yearly growth of new patents of 4.5% for the ini-

rate of new patents of 129%. We obtain 0.045 $\approx ((1 + 1.290)^{1/19} - 1)^{6978} \approx (2.5 (4.4 \times 100)^{1/19})^{1/19}$

 $^{^{69}}_{70}78 \approx 3.5/4.4 \times 100$

⁷⁰Using the coefficient of column (1) in Table 3, we compute the MSA-technology predicted level of patents for 1966 and aggregate it at the Census region - technology level. Then, we compute yearly growth rates within each region-technology and take averages across technologies. Next, we take the average between S and W, and MW and NE, and finally compute the differential predicted growth. If we use the quartile-specific coefficients of column (2) in Table 3 we obtain a predicted differential growth rate of 0.87 percentage points, which implies that the change in travel time can account for 43% of the observed differential growth rate.

⁷¹We use the quartiles of innovativeness defined in section 4.1, computed using the amount of patents of location *i* in technology *h* filed in the time period 1949-1953. Each location *i* has (potentially) a different value quartile in each technology *h*. The 1st quartile refers to the 25% initially least innovative MSAs in technology *h*.

⁷²A given percentage change in knowledge access led to a stronger increase in patenting in initially less innovative locations.

tially lowest quartile of innovativeness, while it predicts 3.4% for the highest quartile.⁷³ The change in knowledge access predicts differential growth rate of 1.1 percentage points. In the data we observe that the average yearly growth rate of patents in the lowest quartile is 5.2 percentage points higher than in the highest quartile. Comparing the predicted and observed differential growth rates, the improvement in knowledge access as consequence of the reduction in travel time explains 21% of the difference in growth rates of new patents between locations in the lowest and highest quartile of innovativeness.⁷⁴

In Appendix Table 21 we present results estimating equation 5 weighting patents by quality using the breakthroughness level computed by Kelly et al. (2021).⁷⁵ We find a larger coefficient in magnitude, providing evidence that the results are not driven by the granting of lower quality patents. Interestingly, we find similar results as in the baseline if we use a quality-weighted version of knowledge access. This finding connects to Iaria et al. (2018), which documents how World War I and the subsequent boycott of Central scientists disrupted international knowledge flows. Their analysis reveals that reduced access to frontier research from abroad disproportionately hindered scientists who depended on top-quality foreign research.

6.2 Creation of knowledge: Instrumental variables estimation

As in Section 5, we may be concerned that decisions of the regulator or airlines which affect travel time are endogenous to the diffusion of knowledge and consequently to knowledge access. Therefore, we construct an instrument for knowledge access in

⁷³The change in knowledge access for the lowest quartile is on average 10.15%, which multiplied by the coefficient 12.91 (obtained by doing 8.41+4.50=12.91) gives a predicted growth of 131% over 19 years. Translated into average yearly growth it is $4.5\% = [(1 + 1.31)^{(1/19)} - 1] \times 100$. For the highest quartile, knowledge access changed on average 10.56%, which multiplied by the coefficient 8.41 predicts 89% growth rate, which is 3.4% yearly growth rate.

 $^{^{74}21\% \}approx 1.1/5.2 \times 100$

⁷⁵We use the patent's 5-year percentile of breakthroughness after demeaning by year fixed effects computed by Kelly et al. (2021). The measure of breakthroughness is computed by comparing the patent text of the focal patent with previous and future patents in a 5-year window to find whether the patent introduces new concepts that were not common before but became common after, making a *breakthrough*. Using the measure computed with 10-year data gives similar results. Importantly, the computation of the measure does not use citation data.

which instead of using observed travel time, we use the fictitious travel time presented in section 3.2 in which routes are fixed to the ones in 1951 and each route is operated with the average airplane of the year:

$$\widetilde{KA}_{iht} = \sum_{k} \omega_{hk} \sum_{j, j \neq i} \text{Patent stock}_{jk,t=1953} \times (\text{instrumental travel time}_{ijt}) \beta$$
(6)

Recently, Borusyak and Hull (2023) have pointed out that when multiple sources of variation are combined to define treatment according to a known formula, treatment exposure can be non-random. Failing to account for this difference in expected treatment can create omitted variable bias. Our instrument combines cross-sectional variation stemming both from the 1951 flight network and the spatial distribution of the knowledge stock in the early 1950s, with variation across time, stemming from the nationwide rollout of jets. A MSA like San Francisco that is connected to knowledge hubs like Chicago via non-stop, long-distance flights would have a large reduction in travel time as consequence of jets and see a large increase in knowledge access. On the other hand, a MSA like Boston which is nearby a major innovation hub like New York, would see a lower increase in knowledge access. In a more general way, MSAs located far away from every other location, in particular the innovative ones, have more possible long-distance connections and are thus more prone to have larger increases in knowledge access due to faster airplanes, creating an omitted variable bias if not accounted for.

Following Borusyak and Hull (2023), we recenter our instrument by subtracting the expected value of the instrument. To construct the *expected instrument* we draw a set of random counterfactual networks, and under each of those we compute the counterfactual travel time and counterfactual value of instrumental knowledge access. Counterfactual networks contain the underlying observed geography and pre-existing distribution of knowledge, hence locations farther apart from innovation centers see a larger increase in knowledge access even in random networks. We then take the average across counterfactual networks to obtain the expected instrument.⁷⁶ By recentering

⁷⁶Details on the construction of the centered instrument are presented in Appendix Section G.2.2.

the instrument we purge it from the non-randomness that might be introduced by geography. The recentered instrument is:

$$log(\widetilde{KA}_{iht})_{centered} = log(\widetilde{KA}_{iht}) - \mathbb{E}[log(\widetilde{KA}_{iht})]$$
(7)

We implement the instrumental variables estimation by control function as in Section 5. Results are presented in Table 3. Columns (3) and (4) show results using the non-centered instrument while columns (5) and (6) use the centered version. The coefficients do not show an important change and the convergence prediction obtained using non-instrumented PPML remains valid.^{77,78}

Figure 7 shows in the left panel the patent growth observed in the data (it replicates Figure 3), while in the right panel it is the predicted patent growth. We compute the prediction using the observed change in travel time and quartile specific elasticities of column (2) in Table 3. Similarly to what is observed in the data, the change in travel time predicts a larger patenting growth rate in the South and the West relative to the Northeast and Midwest.



Figure 7: Observed vs. predicted patent growth 1951 - 1966

⁷⁷The first stage of the model estimated in column (3) of Table 3 gives a $\hat{\lambda}_2 = 1.01$ with standard error 0.03 (clustered at the location-technology level *ih*), and a within R2 of 0.53.

⁷⁸Using non-centered IV estimates, the predicted yearly patent growth rate in the lowest quartile is 4.9% while it is 3.7% in the highest quartile. The predicted differential growth rate is then 1.2 percentage points, meaning that the change in knowledge access can explain $(1.2/5.3) \times 100 \approx 23\%$ of the observed differential growth rate.

The result in column (2) implies that a given change in *Knowledge Access* had a stronger effect on patenting growth in less innovative locations. In other words, knowledge spillovers as an externality had a more predominant role in the production of knowledge in locations that initially produced relatively fewer patents. Theoretically, this result implies that the parameter ρ in equation (1) varies depending on the level of previous production of knowledge of location *i*. Empirically the implication is that a given increase in knowledge spillovers leads to innovation convergence across locations. As seen in section 4.1, during 1949-1968 we observe innovation-convergence across locations and that is exactly what the estimated coefficients predict following a reduction in travel time.

In order to understand the convergence result and compare it with other findings in the literature it is important to remember that commercial airplanes during 1950s and 1960s were a means of transportation mainly for people. On the other hand, other transportation improvements as those in water transport, railroads or highways also contain another ingredient: they were used to carry goods. Hence, other means of transportation combine an impact on both face to face interactions and trade. The effect on trade may be particularly relevant in the presence of economics of scale in production, which may act as a force for concentration of production. Pascali (2017) finds that the introduction of the steam engine vessels in the second half of the 19th century had an impact on international trade that led to economic divergence between countries. Faber (2014) finds that the expansion of the highway system in China led to a reduction of GDP growth in peripheral counties, with evidence suggesting a trade channel due to reduction in trade costs. In our setup, the introduction of jet airplanes represented a big shock to the mobility of people while not affecting significantly the transport of merchandise. Therefore, studying the introduction of jet airplanes allows us to focus on improved face to face interactions, while the trade channel would be a second order effect.

6.3 Creation of knowledge: Robustness

In this section we show that the effect of *Knowledge Access* on the creation of new patents and the convergence effect remains after including different controls. Table 4 shows the results.

Jaworski and Kitchens (2019) show that improvements in the Interstate Highway System led to local increases in income through an increased market access. In our set up, if the effect of market access affects innovation in the same way across technologies, then it would be absorbed by the MSA-time fixed effect FE_{it} in equation (5). However, if the effect of market access on innovation varies across technologies, then it would be a confounder. To control for this potential confounder, we compute market access by highway and interact it with a technology dummy. We compute market access as:

Market
$$Access_{it} = \sum_{j} Population_{j,t=1950} \times \tau^{\theta}_{ijt}$$
 (8)

where Population_{*j*,*t*=1950} is population in MSA *j* in 1950, τ_{ijt} are the shipping costs provided in the data of Taylor Jaworski and Carl Kitchens computed using each year's highway driving distance, highway travel time, petrol cost and truck driver's wage. θ is the elasticity of trade to trade costs which we set to -8.28, the preferred value of Eaton and Kortum (2002) and in the range of many other estimates in the literature (Head and Mayer (2014), Caliendo and Parro (2015), Donaldson and Hornbeck (2016)). Columns (3) and (4) of Table 4 show the results, we do not observe an important difference with the baseline estimates.

Campante and Yanagizawa-Drott (2017) shows that better connectivity by airplane leads to an increase in economic activity as measured by satellite-measured night light. Söderlund (2023) shows that an increase in business travel in the late 1980s and early 1990s led to an increase in trade between countries. In a similar way to knowledge access, we could think that better connectivity by airplane could have led to an increase in market access due to a reduction in information frictions, with goods being shipped by land. Similarly to highway market access, if the effect of market access by airplane is common to all technology categories then it would be absorbed by the MSA-time fixed effect FE_{it} . In order to account for a technology-specific effect, we construct a measure of airplane market access and interact it with a technology dummy. The measure of airplane market access is similar to equation (8) where τ is the travel time by airplane and θ is set to -1.304, the elasticity of trade to travel time from Söderlund (2023). The results are shown in columns (5) and (6) of Table 4. While the coefficients in all quartiles are reduced, the estimated value of ρ is positive and significant and the result on convergence remains.

Potential contemporaneous improvements in other means of communication, like telephones, could have spurred the creation of new patents. In columns (7) and (8) we include the log of the MSA's share of households with telephones in 1960 and double-interact it with a technology dummy and a time dummy. The results remain invariant with respect to the baseline.

Another potential explanation for the increase of patenting could be that better connectivity decreased technology-specific financial frictions. The potential reduction in financial frictions, rather than a confounder, would be a mechanism through which airplanes increased innovation. However, according to Jayaratne and Strahan (1996) during 1950s and 1960s interstate lending and bank branching were limited. Prior to the 1970s, banks and holdings were restricted in their geographic expansion within and across state borders. Additionally, the Douglas Amendment to the Bank Holding Company Act prevented holding companies from acquiring banks in other states. Therefore, it is unlikely that interstate bank financing would be a driving force. Nonetheless, if other sector-specific modes of financing like venture capital were active, they could be driving the results. In Appendix G we construct multiple measures of access to capital by using market capitalization of patenting firms listed in the stock market. The results present suggestive evidence that access to capital is not driving the results.

Finally, in Appendix G we include additional robustness checks. We compute different versions of *Knowledge Access*: we use distance-specific β from section 5, we consider the patent stock only of locations *j* far from *i*, we do sensitivity analysis using different values of β . Also, we re estimate the effects by quartile of initial innovativeness using patents per capita. Last, we re-do the baseline regression using OLS estimation. Re-

	PPML							
Dependent Variable:	Patents							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log(knowledge access)	9.11***	8.41**	8.33**	7.40**	5.83*	5.45^{*}	9.30***	8.33***
	(3.29)	(3.31)	(3.30)	(3.31)	(3.20)	(3.22)	(3.08)	(3.08)
$\log(\text{knowledge access}) \times 3rd quartile = 0.50$		1.86***		1.95***		1.86***		2.01***
		(0.53)		(0.52)		(0.53)		(0.51)
log(knowledge access) \times 2nd quartile = 0.25		3.42***		3.50***		3.38***		3.55***
		(0.81)		(0.81)		(0.80)		(0.82)
$\log(\text{knowledge access}) \times 1$ st quartile = 0.00		4.50***		4.62***		4.57***		4.67***
		(1.17)		(1.17)		(1.17)		(1.19)
N obs. effective	991,480	991,480	991,480	991,480	991,480	991,480	991,480	991,480
R2	0.85	0.85	0.85	0.85	0.85	0.85	0.85	0.85
Controls:								
log(Highway market access) $ imes$ technology	-	-	Yes	Yes	-	-	-	-
$\log(\text{Airplane market access}) imes ext{technology}$	-	-	-	-	Yes	Yes	-	-
log(Telephone share) \times technology \times time	-	-	-	-	-	-	Yes	Yes

***p < 0.01; **p < 0.05; *p < 0.10

Table 4: Elasticity of new patents to knowledge access, by MSA innovativeness quartile

Column (1) shows the result of Poisson Pseudo Maximum Likelihood (PPML) estimation of Patents_{*Fiht*} = exp [$\rho log(KA_{iht}) + FE_{Fih} + FE_{it} + FE_{ht}$] × ξ_{Fiht} , for patents filed by establishment of firm *F* in location *i*, technology *h* and time period *t*. *KA*_{iht} is knowledge access of establishments in location *i* technology *h* and time period *t*. Column (2) opens the coefficient ρ by the quartile of innovativeness of location *i* within technology *h*, computed using patents in 1949-1953. Higher quartile indicates higher initial level of innovativeness. The fourth quartile is used as reference category. Relative to columns (1) and (2), columns (3) and (4) control for technology specific effect of log(highway market access), columns (5) and (6) control for technology specific effect of log(airplane market access), columns (7) and (8) control for technology and time specific effect of log(telephone share). Standard errors clustered at the location-technology *ih* are presented in parentheses. R2 is computed as the squared correlation between observed and fitted values.

sults go in the same direction: an increase in knowledge access leads to an increase in patenting and the effect is stronger in initially less innovative locations.

7 Conclusion

This paper studies how frictions to the mobility of people affect the geography of innovation in the context of the early Jet Age in the United States. With newly digitized data on airlines' flight schedules, we construct a dataset of the flight network in the United States during the 1950s and 1960s. We document the large reduction in air travel time that jet airplanes brought about: around 11 hours, a 50% reduction, for coast-to-coast return trip. Combined with patent data, we find that the reduction in travel time increased knowledge diffusion, especially between research establishments located far apart. The increase in diffusion of knowledge created long-distance spillovers and led to the production of new knowledge.

Our results point to jet airplanes as an important driver behind major changes in the geography of innovation in the United States post World War II: a catching up of the South and the West with the Northeast and the Midwest, and initially less innovative MSAs reducing the gap with more innovative ones.

The results provide policy-relevant insights regarding the impact of passenger transport infrastructure on the emergence of technology clusters. Large R&D policies, like the recent CHIPS and Science Act, frequently include a place-based component to increase technology capacity in regions that lag behind (Gruber and Johnson (2019), Gross and Sampat (2023)). Our results show that connectivity to existing clusters can lead to an increase of local innovation and act as a convergence force between regions.

References

Acemoglu, D., U. Akcigit, and W. R. Kerr (2016). Innovation network. *Proceedings of the National Academy of Sciences* 113(41), 11483–11488.

- Agrawal, A., A. Galasso, and A. Oettl (2017). Roads and innovation. *The Review of Economics and Statistics* 99(3), 417–434.
- Andrews, M. J. and A. Whalley (2021). 150 years of the geography of innovation. *Regional Science and Urban Economics*, 103627.
- Arzaghi, M. and J. V. Henderson (2008). Networking off madison avenue. *The Review of Economic Studies* 75(4), 1011–1038.
- Atkin, D., M. K. Chen, and A. Popov (2022). The returns to face-to-face interactions: Knowledge spillovers in silicon valley. Technical report, National Bureau of Economic Research.
- Bahar, D., P. Choudhury, D. Y. Kim, and W. W. Koo (2023). Innovation on wings: Nonstop flights and firm innovation in the global context. *Management Science* 69(10), 6202–6223.
- Bai, J., W. Jin, and S. Zhou (2023). Proximity and knowledge spillovers: Evidence from the introduction of new airline routes. *Management Science*.
- Baum-Snow, N. (2007). Did highways cause suburbanization? *The quarterly journal of economics* 122(2), 775–805.
- Bergé, L. (2018). Efficient estimation of maximum likelihood models with multiple fixed-effects: the R package FENmlm. *CREA Discussion Papers* (13).
- Berger, T. and E. Prawitz (2024). Collaboration and connectivity: Historical evidence from patent records. *Journal of Urban Economics* 139, 103629.
- Borenstein, S. and N. L. Rose (2014, June). *How Airline Markets Work... or Do They? Regulatory Reform in the Airline Industry*, pp. 63–135. University of Chicago Press.
- Borusyak, K. and P. Hull (2023). Nonrandom exposure to exogenous shocks. *Econometrica* 91(6), 2155–2185.
- C.A.B. (1951, 1956, 1961, 1966). Air carrier traffic statistics. Civil Aeronautics Board.

- Cai, J., N. Li, and A. M. Santacreu (2022). Knowledge diffusion, trade, and innovation across countries and sectors. *American Economic Journal: Macroeconomics* 14(1), 104– 145.
- Caliendo, L. and F. Parro (2015). Estimates of the trade and welfare effects of nafta. *The Review of Economic Studies 82*(1), 1–44.
- Campante, F. and D. Yanagizawa-Drott (2017, 12). Long-Range Growth: Economic Development in the Global Network of Air Links*. *The Quarterly Journal of Economics* 133(3), 1395–1458.
- Carlino, G. and W. R. Kerr (2015). Agglomeration and innovation. *Handbook of regional and urban economics* 5, 349–404.
- Chatterji, A., E. Glaeser, and W. Kerr (2014). Clusters of entrepreneurship and innovation. *Innovation policy and the economy* 14(1), 129–166.
- Coscia, M., F. Neffke, and R. Hausmann (2020). Knowledge diffusion in the network of international business travel. *Nature Human Behaviour* 4(10).
- De Rassenfosse, G. and A. B. Jaffe (2017). Econometric evidence on the r&d depreciation rate. Technical report, National Bureau of Economic Research.
- Dhaene, G. and K. Jochmans (2015). Split-panel jackknife estimation of fixed-effect models. *The Review of Economic Studies* 82(3), 991–1030.
- Donaldson, D. and R. Hornbeck (2016). Railroads and american economic growth: A "market access" approach. *The Quarterly Journal of Economics* 131(2), 799–858.
- Duranton, G., P. Martin, T. Mayer, and F. Mayneris (2009). The economics of clusters: evidence from france.
- Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica* 70(5), 1741–1779.
- Faber, B. (2014). Trade integration, market size, and industrialization: evidence from china's national trunk highway system. *Review of Economic Studies* 81(3), 1046–1070.

- Fagerberg, J., D. C. Mowery, and R. R. Nelson (2005). *The Oxford handbook of innovation*. Oxford university press.
- Feldman, M. P. (1994). The geography of innovation, Volume 2. Springer Science & Business Media.
- Feyrer, J. (2019). Trade and income—exploiting time series in geography. *American Economic Journal: Applied Economics* 11(4), 1–35.
- Fogel, R. W. (1963). *Railroads and American economic growth: essays in econometric history*.Ph. D. thesis, Johns Hopkins University.
- Furman, J. L. and S. Stern (2011, August). Climbing atop the shoulders of giants: The impact of institutions on cumulative research. *American Economic Review* 101(5), 1933–63.
- Gertner, J. (2013). *The idea factory: Bell Labs and the great age of American innovation*. Penguin.
- Glaeser, E. (2011). Triumph of the City. Macmillan.
- Gross, D. P. (2019). The consequences of invention secrecy: Evidence from the uspto patent secrecy program in world war ii. Technical report, National Bureau of Economic Research.
- Gross, D. P. and B. N. Sampat (2023). America, jump-started: World war ii r&d and the takeoff of the us innovation system. *American Economic Review* 113(12), 3323–3356.
- Gruber, J. and S. Johnson (2019). *Jump-starting America: How breakthrough science can revive economic growth and the American dream*. Hachette UK.
- Haines, M. R. et al. (2010). Historical, demographic, economic, and social data: the united states, 1790–2002. *Ann Arbor, MI: Inter-university Consortium for Political and Social Research*.

- Head, K. and T. Mayer (2014). Chapter 3 gravity equations: Workhorse, toolkit, and cookbook. In G. Gopinath, E. Helpman, and K. Rogoff (Eds.), *Handbook of International Economics*, Volume 4 of *Handbook of International Economics*, pp. 131–195. Elsevier.
- Herzog, I. (2021). National transportation networks, market access, and regional economic growth. *Journal of Urban Economics* 122, 103316.
- Holbrook, D., W. M. Cohen, D. A. Hounshell, and S. Klepper (2000). The nature, sources, and consequences of firm differences in the early history of the semiconductor industry. *Strategic Management Journal* 21(10-11), 1017–1041.
- Hovhannisyan, N. and W. Keller (2015, March). International business travel: an engine of innovation? *Journal of Economic Growth* 20(1), 75–104.
- Iaria, A., C. Schwarz, and F. Waldinger (2018). Frontier knowledge and scientific production: Evidence from the collapse of international science. *The Quarterly Journal* of Economics 133(2), 927–991.
- I.C.C. (1965). 79th Annual Report. Interstate Commerce Commission.

I.C.C. (1967). 81st Annual Report. Interstate Commerce Commission.

- Jaffe, A. B., M. Trajtenberg, and M. S. Fogarty (2000a). Knowledge spillovers and patent citations: Evidence from a survey of inventors. *American Economic Review* 90(2), 215–218.
- Jaffe, A. B., M. Trajtenberg, and M. S. Fogarty (2000b, May). Knowledge spillovers and patent citations: Evidence from a survey of inventors. *American Economic Review* 90(2), 215–218.
- Jaffe, A. B., M. Trajtenberg, and R. Henderson (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *The Quarterly Journal of Economics* 108(3), 577–598.

- Jaworski, T. and C. T. Kitchens (2019). National policy for regional development: Historical evidence from appalachian highways. *Review of Economics and Statistics* 101(5), 777–790.
- Jayaratne, J. and P. E. Strahan (1996). The finance-growth nexus: Evidence from bank branch deregulation. *The Quarterly Journal of Economics* 111(3), 639–670.
- Kantor, S. and A. T. Whalley (2023). Moonshot: Public r&d and growth. Technical report, National Bureau of Economic Research.
- Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy (2021). Measuring technological innovation over the long run. *American Economic Review: Insights* 3(3), 303–320.
- Kerr, W. R. and F. Robert-Nicoud (2020). Tech clusters. *Journal of Economic Perspectives* 34(3), 50–76.
- Manson, S., J. Schroeder, D. Van Riper, T. Kugler, and S. Ruggles (2020). Ipums national historical geographic information system: Version 15.0.
- Michaels, G. (2008). The effect of trade on the demand for skill: Evidence from the interstate highway system. *The Review of Economics and Statistics 90*(4), 683–701.
- Moretti, E. (2010). Local multipliers. American Economic Review 100(2), 373–377.
- Moretti, E. (2021, October). The effect of high-tech clusters on the productivity of top inventors. *American Economic Review* 111(10), 3328–75.
- Moretti, E. and P. Thulin (2013). Local multipliers and human capital in the united states and sweden. *Industrial and Corporate Change* 22(1), 339–362.
- Nagler, M., M. Schnitzer, and M. Watzinger (2022). Fostering the diffusion of general purpose technologies: Evidence from the licensing of the transistor patents. *The Journal of Industrial Economics* 70(4), 838–866.
- Pascali, L. (2017). The wind of change: Maritime technology, trade, and economic development. *American Economic Review* 107(9), 2821–54.

- Perlman, E. R. (2016). Dense enough to be brilliant: patents, urbanization, and transportation in nineteenth century america. *Work. Pap., Boston Univ.*
- Perlow, L. A. (1999). The time famine: Toward a sociology of work time. *Administrative science quarterly* 44(1), 57–81.
- Roach, M. and W. M. Cohen (2013). Lens or prism? patent citations as a measure of knowledge flows from public research. *Management Science* 59(2), 504–525.
- Sanderson, E. and F. Windmeijer (2016). A weak instrument f-test in linear iv models with multiple endogenous variables. *Journal of econometrics* 190(2), 212–221.
- Silva, J. M. C. S. and S. Tenreyro (2006). The log of gravity. *The Review of Economics and Statistics 88*(4), 641–658.
- Söderlund, B. (2023). The importance of business travel for trade: Evidence from the liberalization of the soviet airspace. *Journal of International Economics* 145, 103812.
- Staiger, D. and J. H. Stock (1997). Instrumental variables regression with weak instruments. *Econometrica* 65(3), 557–586.
- Stock, J. H. and M. Yogo (2005). *Testing for Weak Instruments in Linear IV Regression*, pp. 80–108. Cambridge University Press.
- Storper, M. and A. J. Venables (2004, 08). Buzz: face-to-face contact and the urban economy. *Journal of Economic Geography* 4(4), 351–370.
- Tsiachtsiras, G. (2021). Transportation networks and the rise of the knowledge economy in 19th century france.
- U.S. Department of Commerce (1958). *Travel Survey 1957*. Transport Division, Bureau of Census.
- Wardman, M. (2012). Review and meta-analysis of uk time elasticities of travel demand. *Transportation* 39(3), 465–490.

- Weidner, M. and T. Zylkin (2021). Bias and consistency in three-way gravity models. *Journal of International Economics* 132, 103513.
- Wooldridge, J. M. (2014). Quasi-maximum likelihood estimation and testing for nonlinear models with endogenous explanatory variables. *Journal of Econometrics* 182(1), 226–234.