



Paper to be presented at
DRUID15, Rome, June 15-17, 2015
(Coorganized with LUISS)

Roads and Innovation

Ajay Agrawal
University of Toronto
Rotman
ajay.agrawal@rotman.utoronto.ca

Alberto Galasso
University of Toronto
Rotman
alberto.galasso@rotman.utoronto.ca

Alexander Oettl
Georgia Institute of Technology
Scheller College of Business
oettl@gatech.edu

Abstract

We estimate the effect of within-MSA interstate highways on regional innovation by exploiting historical data on planned portions of the highway system, railroads, and exploration routes as sources of exogenous variation. We report evidence that a 10% increase in a region's stock of highways causes a 1.7% increase in regional patenting over a five-year period. We also show that roads facilitate local knowledge flows, allowing innovators to access more distant local knowledge inputs. This suggests that transportation infrastructure may spur regional growth above and beyond the more commonly discussed agglomeration economies predicated on the inflow of new workers.

Roads and Innovation

December 21, 2014

Abstract

We estimate the effect of within-MSA interstate highways on regional innovation by exploiting historical data on planned portions of the highway system, railroads, and exploration routes as sources of exogenous variation. We report evidence that a 10% increase in a region's stock of highways causes a 1.7% increase in regional patenting over a five-year period. We also show that roads facilitate local knowledge flows, allowing innovators to access more distant local knowledge inputs. This suggests that transportation infrastructure may spur regional growth above and beyond the more commonly discussed agglomeration economies predicated on the inflow of new workers.

JEL Classifications: O33, O47, L91

Keywords: innovation, transportation, highways, regional growth

1 Introduction

A striking feature of economic geography is the large variation in productivity across regions. Moretti (2011) documents that after adjusting for skill composition, average wages in the highest- and lowest-paying US metropolitan areas differ by approximately a factor of three. Such dispersion is also evident when one compares innovation outcomes across regions (Agrawal et al., 2014; Carlino and Kerr, 2014). Silicon Valley and Boston are popular examples of outlier regions, significantly more productive than others in terms of innovation.

Such regional disparities have led to a variety of policies focused on enhancing local economic activity. One of the main policies that local governments implement to spur regional economic growth is the provision of infrastructure that reduces local transportation costs. Transportation infrastructure, such as roads, may impact regional productivity through its effect on employment, private investment, and the returns to schooling. They are also likely to facilitate knowledge diffusion and spillovers, which the modern macroeconomic literature identifies as a key source of economic growth (Romer, 1986). The effect of roads on knowledge creation and diffusion is the main focus of this project.

Transportation infrastructure represents a large portion of the U.S. economy. The estimated value of the U.S. road capital stock is roughly \$5 trillion (US Bureau of Transportation Statistics, 2010), and about 20% of the income of the median U.S. household is devoted to road transportation (Duranton and Turner, 2012). Despite the magnitude of these investments along with mounting empirical evidence that knowledge flows are geographically localized, as reported in the seminal work of Jaffe et al. (1993) and a number of subsequent studies, the impact of transportation infrastructure on knowledge creation and diffusion has been overlooked by the innovation literature. We aim to address this here. Our research provides insight for policies aimed at enhancing the flow of knowledge within cities. Furthermore, our findings offer insight for managers who make location and technology strategy decisions because regional knowledge flows are key determinants of firm survival and competitive advantage.

A major identification challenge in estimating the effect of highways on innovation is the simultaneous determination of transportation infrastructure and regional technological development. For example, economic growth in a city may drive both growth in local innovation as well as increased investments in infrastructure in that region. Thus, we must be cautious in interpreting the causal effect of roads on innovation. To solve

this problem, we follow a growing literature in urban economics that exploits instrumental variables for transportation infrastructure (Baum-Snow (2007); Michaels (2008); Duranton and Turner (2012); Duranton et al. (2013)). Building on Duranton and Turner (2012), we consider three instruments for the presence of road infrastructure. The first is based on the 1947 plan of the U.S. interstate highway system. The second is derived from a map of the U.S. railroad network at the end of the 19th century. The third is based on maps of routes of major exploration expeditions of the U.S. from 1518 to 1850.

We find that a 10% increase in interstate highways leads to a roughly 1.7% increase in innovation measured by the patenting activity in the region. This is a large effect, comparable to more than a 3% increase in regional corporate R&D investments. We show that the results are similar using MSA or MSA-technology class as units of analysis and that the estimates are robust to the inclusion of a large set of control variables that can explain persistent productivity differences across regions.

In principle, there are a number of mechanisms through which transportation infrastructure may affect the creation and diffusion of knowledge. We highlight one particular channel through which roads have an impact on innovation: the facilitation of local *within-region* knowledge flow. Specifically, we show that in regions where the stock of transportation infrastructure is larger, innovators build on local knowledge that is geographically more distant. Research in urban economics has emphasized that transportation infrastructure generates regional growth through agglomeration economies, typically modeled as an inflow of new workers (Duranton and Turner, 2012). An important feature of the channel that we highlight is that it does not require an influx of new innovators. Our findings are robust to focusing on a sample of *non-mover* inventors, whose locations do not change during the period of our study. This reinforces our view that transportation infrastructure facilitates the circulation of local knowledge even in the absence of an inflow of new labor, the mechanism typically linked to agglomeration

forces.

Our analysis documents a greater propensity to build on more distant local knowledge in regressions with regions as the unit of analysis (the standard approach in the urban economics literature) as well as in patent level regressions (the standard approach in the innovation literature). At the disaggregated patent level, we show that, conditioning on the distance between two inventors located in the same region, the probability of a citation between two of their patents increases with the stock of highways in the region.

We also provide additional indirect evidence supporting the idea that roads increase local knowledge flows. First, we show that roads have a greater impact on innovation in fields where the technology frontier shifts more quickly such that rapid access to new knowledge is more valuable. Second, we show that the effect of transportation infrastructure is larger in regions characterized by the presence of *star* inventors who generate more significant spillovers. Third, we show that highways have a larger impact on innovation in regions characterized by low density where inventors are likely to be more spread out. Finally, we show that large firms' innovation is less sensitive to highway provision, consistent with the idea that larger firms are more likely to build upon knowledge produced within their own boundaries and thus rely less on that produced by their neighbors.

We conclude the paper with an illustrative quantitative estimation. We develop a simple structural model in which transportation infrastructure affects productivity through two distinct channels. The first is an agglomeration force: roads increase the local supply of labor, which increases labor productivity. The second is a non-agglomeration knowledge channel: roads allow greater patenting because they facilitate knowledge flows even when the supply of labor is fixed. Calibration of the model suggests that about 30% of the impact of roads on labor productivity may be due to

non-agglomeration channels.

We organize the paper as follows. We summarize the related literature in Section 2 and describe our data in Section 3. We introduce the empirical framework in Section 4. We present the results showing the impact of roads on the growth of regional innovation in Section 5 and the results showing the impact of roads on the flow of local knowledge in Section 6. In Section 7, we present a structural model, which decomposes the impact of roads on labor productivity. Finally, we provide brief concluding remarks to close the paper. We include extensions of the theoretical model as well as additional empirical results in the Appendices.

2 Related Literature

Our paper is connected to two literatures: 1) the determinants of regional innovation, and 2) the impact of transportation infrastructure on regional growth.

Building on seminal work by Jaffe et al. (1993) and Feldman (1994), the regional innovation literature identifies a number of factors that may increase innovation in a geographic area by affecting the localization of knowledge spillovers. For example, Feldman and Audretsch (1999) show that diversity of economic activities in a region better promotes innovation. This provides support to Jacobs (1969), who argues that the exchange of complementary knowledge across industries is central to the creation of new economic knowledge and thus growth. Agrawal et al. (2008) provide evidence that social/ethnic proximity substitutes for geographic proximity in terms of its influence on regional knowledge flow patterns, suggesting that the dispersion of socially proximate individuals maximizes regional innovation. Kerr and Kominers (2015) study how the shape of spatial clusters of firms depends on agglomerative forces and on interaction costs. Catalini (2013) provides evidence that microgeographic forces also affect idea

recombination and the direction of inventive activity.

In terms of the effect of industrial organization on regional innovation, Agrawal and Cockburn (2003) report evidence in support of the anchor tenant hypothesis that large, local, R&D-intensive firms have a positive impact on regional innovation. Agrawal et al. (2014) extend this work, showing that local innovation is affected by the organization of R&D manpower in the region and in particular that innovation output is higher in regions that include not only large R&D-intensive firms but also small ones that thicken the market for ancillary services, thereby lowering the cost of spin-outs. In terms of government policies, Marx et al. (2009) show that regional non-compete regulations affect inventor mobility and knowledge spillovers. Belenzon and Schankerman (2013) show that local policies can promote commercial development and diffusion of university innovations. Galasso et al. (2013) show that state-level taxes strongly impact knowledge diffusion through the decision to trade patent rights.

The emerging urban economics literature studies the impact of investments in transportation infrastructure on the evolution of metropolitan areas. Fernald (1999) is the first paper that tries to identify the causal impact of infrastructure on regional productivity. Focusing on the differential impact of highways on productivity growth in industries that have different levels of vehicle intensity, he shows that industries with a lot of vehicles benefited disproportionately from road-building. He interprets this finding as suggestive of the positive impact of changes in road stock to regional productivity. Baum-Snow (2007) exploits the planned proportion of the interstate highway system as a source of exogenous variation to estimate the impact of transportation infrastructure on suburbanization. He finds that one new highway passing through a central city reduces its population by about 18%. Baum-Snow (2013), exploiting the same instrument, shows that the construction of highways causes a large and significant job displacement in city centers but has only a minor impacts on jobs in the suburbs. Michaels (2008) studies

the impact of interstate highways on rural counties. He finds that highways generate an increase in trade-related activities, such as trucking and retail sales.

Duranton and Turner (2012) exploit interstate highway system plans, railroads, and exploration maps as instruments to study the impact of highways on regional growth. They find that a 10% increase in highway stock in a city causes about a 1.5% increase in employment over a 20-year period. Duranton et al. (2013) study the impact of interstate highways on the level and composition of trade for U.S. cities. They find that highways have no effect on the total value of exports and that cities with more highways specialize in sectors producing heavy goods. Finally, Donaldson (2014) exploits archival data from colonial India to estimate the effect of India's railroad network. He finds that railroads cause a decrease in trade costs and interregional price gaps and increase interregional and international trade. He also shows that railroads are associated with an increase in real income levels.

3 Data

We follow Agrawal et al. (2014) in constructing our sample and begin with the set of 268 Metropolitan Statistical Areas (MSAs) defined in 1993 by the U.S. Office of Management and Budget and the set of six one-digit technology classes described in Hall et al. (2001).¹

We obtain information on U.S. patenting activity and on the affiliation and location of patenting inventors in a region from the United States Patent and Trademark Office (USPTO) data. While these data are complete and detailed, two key qualifications should be kept in mind. First, not all inventions are patented. Although this presents a significant limitation to these data, the innovation literature has shown that technologies

¹This classification scheme includes: chemicals, computers and communications, drugs and medical, electrical and electronic, mechanical, and others.

with greater impact on social welfare and economic growth are more likely to be patented (Pakes and Griliches (1980); Griliches (1990)). Second, the coding of inventor location, affiliation, and identity is likely to generate random measurement error in our constructs.

As in Agrawal et al. (2014), we use inventor address information to assign a patent to an MSA, exploiting the U.S. National Geospatial-Intelligence Agency dataset to match cities and townships to counties and ultimately MSAs. If a patent has at least one inventor from a particular MSA, then we increment the counter for that MSA by one. Thus, a patent with three inventors located in three different MSAs increments the patent counter for each of those MSAs by one. However, if all three inventors are located in the same MSA, then the counter for that MSA is only incremented by one.²

We construct our measures using all patents with at least one inventor with a U.S. address. We exclude patents that cannot be attributed to an MSA (due to incomplete address information or a location outside an MSA) and patents assigned to universities and hospitals. While the USPTO is the original source of our patent data, we complement these data with classification data from the NBER (technology classification, assignee name).

We measure innovative activity, our main dependent variable, using a citation-weighted count of U.S. patents:

Weighted Patents_{jkt}: the forward citation weighted sum of distinct patents with primary technology classification k and application year t where at least one inventor is located in MSA j .

Patent citations identify prior knowledge upon which a patent builds, and prior literature (starting with Pakes and Griliches (1980)) has often employed the number of

²Agrawal et al. (2014) show that differences in the variables are negligible if they are constructed using only data from the first inventor.

forward-citations received by a patent as an indirect measure of patent value. We also consider an unweighted patent count as an additional innovation metric:

Patents_{jkt}: the number of distinct patents with primary technology classification k and application year t where at least one inventor is located in MSA j .

Our main explanatory variable is the total number of kilometers of interstate highway in the region in 1983 constructed from the Highway Performance and Monitoring System data that are extensively described in Duranton and Turner (2012). All our results are robust to using an alternative lane-weighted measure of highway stock.

Following Agrawal et al. (2014), we construct variables for the number of inventors in 1983, 1978, and 1973 at the MSA and MSA-class levels. As additional control variables, following Duranton and Turner (2012), we use the logarithms of MSA historical population levels. We also exploit a number of variables describing the physical geography of MSAs as controls. Burchfield et al. (2006) show that the spatial structure of a region is strongly shaped by the availability of groundwater, so we exploit the share of each MSA's land that overlays an aquifer. Following Duranton and Turner (2012), we also use controls for MSA elevation, ruggedness of MSA terrain, and MSA climate (heating and cooling degree days). We exploit a variety of socio-demographic variables from the 1980 census: the share of poor population in the MSA, the share of college graduates, the share of population employed in manufacturing, and mean income in the MSA. We also employ a measure of housing segregation computed by Cutler and Glaeser (1997). In some regressions, we use indicator variables for each of the nine census divisions.

Finally, our analysis of local knowledge flows calculates the distances between cities/towns within an MSA. For each city we identify its centroid geographic coordinates from the US Geological Survey and calculate distances between cities using the great circle method as in Singh and Marx (2013) and Li et al. (2014).

4 Empirical Framework

Our main econometric model focuses on the relationship between measures of innovative activity Y_{jk1988} in MSA-class jk in 1988 and the level of interstate highway in MSA j in 1983, $Highway_{j1983}$. Our main specification takes the following form:

$$\log Y_{jk1988} = \alpha + \beta \log Highway_{j1983} + \gamma \log Y_{jk1983} + \theta X_{jk} + \epsilon_{jk} \quad (1)$$

where Y_{jk1983} is the innovation level in 1983 and X_{jk} is a vector of additional controls.

This empirical specification is consistent with a simple model in which the deterministic innovation level in an MSA, K_t^* , is related to the level of highways, R_t , by the following relationship $K_t^* = AR_t^a$. The rate of innovation adjustment depends on how far out of steady state a region is. If we define the adjustment as $K_{t+5} = K_t^{*1-\gamma} K_t^\gamma$ with $0 < \gamma < 1$, then patenting in period $t + 5$ will be equal to:

$$K_{t+5} = BR_t^\beta K_t^\gamma, \quad (2)$$

where $\beta = a(1 - \gamma)$ and $B = A^{1-\gamma}$. Taking logs of (2), we obtain the estimated regression (1). The parameter of interest is β , which in this simple model describes the rate at which knowledge creation responds to highway provision. More specifically, an unbiased estimate of β answers the following question: *Does the level of MSA highway stock in 1983 affect innovation growth during the period 1983-1988?*

Notice that (1) can be re-written as:

$$\log Y_{jk1988} - \log Y_{jk1983} = \alpha + \beta \log Highway_{j1983} + (\gamma - 1) \log Y_{jk1983} + \theta X_j + \epsilon_{jk},$$

Therefore, there is no loss of generality in interpreting β as a coefficient linking the 1983

Highway level with innovation growth for the period 1983-1988.³

The main empirical challenge in estimating equation (1) is the possible correlation between unobservables, ϵ_{jk} , and the level of highways in a region. For example, the local government may react to an economic downturn by building more roads, and this would generate a negative correlation between roads and innovation. In this case, OLS estimates would underestimate the causal impact of highways on innovation. To address such a concern, we exploit three instrumental variables that we discuss in detail the following sub-section.

4.1 Instrumental Variables

We exploit three instrumental variables (IVs) constructed using archival data on historical transportation infrastructure. While a number of studies in the urban economics literature use historical data as a source of exogenous variation (Baum-Snow (2007); Michaels (2008); Duranton and Turner (2012); Duranton et al. (2013); Donaldson (2014); Duranton et al. (2013)), this empirical approach is novel in the entrepreneurship and innovation literature. The only exception we are aware of is Glaeser et al. (2012), who exploit historical mines as an instrument for entrepreneurship.

To be a valid instrument, an historical variable must not only be a good predictor of the level of interstate highways in 1983 but also be orthogonal to the structural equation error term. We now describe the historical data and discuss their validity as instrumental variables. All three instruments were constructed by Duranton and Turner (2012).

³Model (1) differs from difference-in-differences estimators typically used in the innovation literature. First, the treatment variable, R , is a continuous variable and not a dummy. Second, because our sample covers only two periods we cannot test the assumption of common pre-trends in knowledge creation between cities with different levels of highways in 1983. Nonetheless, instrumenting $Highway_{j1983}$ allows us to remove the bias generated by non-common trends and identify the causal effect of highways on innovation. Therefore the interpretation of our estimates is not substantially different from the typical interpretation in a difference-in-difference model.

The 1947 Plan of the Interstate Highway System

Our first instrumental variable is a measure of the total number of kilometers of highway planned at the national level in 1947. Durant and Turner (2012) construct this variable from a digital image of the 1947 highway plan for which they calculate kilometers of interstate highway in each MSA. Many of the highways planned in 1947 were ultimately built, and the correlation between log 1983 interstate highway kilometers and log 1947 planned highway kilometers is 0.62.

The orthogonality of this instrument relies on the fact that the 1947 proposal was a myopic plan, based on the defense needs and economic conditions of the mid 1940s that are likely to be uncorrelated with innovation activity in the 1980s. Specifically, the goal of the 1947 plan was to “connect by routes as direct as practicable the principal metropolitan areas, cities and industrial centers, to serve the national defense and to connect suitable border points with routes of continental importance in the Dominion of Canada and the Republic of Mexico” (United States Federal Works Agency, Public Roads Administration, 1947). Historical evidence discussed in Durant et al. (2013) confirms that the 1947 highway plan was drawn to this mandate. Moreover, Durant and Turner (2012) show that 1947 planned highways are uncorrelated with population growth in the 1940s and 1950s, confirming that planners in 1947 tried to connect population centers, not anticipate future growth.

The instrumental variable estimation of equation (1) requires orthogonality of the dependent variable and the instruments conditional on control variables, not unconditional orthogonality. As Durant and Turner (2012) point out, this is an important distinction. For example, MSAs with more roads in the 1947 plan may be larger and thus have more inventors than MSAs receiving less. If innovation growth depends on the number of inventors in the MSA and there is persistence in the R&D labor force, then

the 1947 planned highway system may predict innovation growth directly through its ability to predict the R&D labor force in 1983. To address this concern and reduce the threat to the validity of the instrument, we follow the urban economics literature and include in the estimation a large set of appropriate controls (in particular the historical number of inventors and population levels).

Railroad Routes in 1898

The second instrument is based on the map of major railroad lines from about 1898 (Gray, 1989). Duranton and Turner (2012) calculate the kilometers of 1898 railroad track contained in each MSA by converting this map into a digital image. The correlation between log 1983 interstate highways kilometers and log 1898 railroad kilometers is equal to 0.53. Such high correlation is driven by the fact that old railroads are a natural location for modern roads because they do not require levelling and grading a roadbed.

The US rail network was developed in the middle of an industrial revolution and immediately after the Civil War. At that time, the US economy was smaller and more agricultural than the one of the 1980s, and this substantially reduces the concern of correlation between railroads in 1898 and technology shocks in the 1980s. As discussed in Duranton and Turner (2012) and Duranton et al. (2013), railroads were developed mainly to transport grain, livestock, and lumber, and it is unlikely that such a flow of agricultural commodities is correlated to innovation activity in the 1980s. Moreover, railroads were typically constructed by private companies expecting to make profits in the short and medium term.

The validity of the instrument again hinges on its orthogonality conditional on the control variables. A possible concern is that cities with more kilometers of railroad track in 1898 were more productive, and persistent productivity differences may be correlated with greater innovation in the 1980s. To address this concern, we will show that our

results are robust to including direct measures of productivity (e.g., historical growth in the number of inventors, income per capita, the share of adult population with a college degree).

Routes of major exploration expeditions, 1528-1850

The final instrument is an index that measures the number of routes of major exploration expeditions that crossed each MSA. Duranton and Turner (2012) digitize a number of maps from the National Atlas of the United States of America (1970) reporting routes of major expeditions of exploration that occurred during the time period 1528-1850. From each map, they count one kilometer for each pixel crossed by an exploration route in each MSA and then construct their measure by summing those counts across all maps.

The correlation between the exploration route index and 1983 kilometers of interstate highway is equal to 0.43. Such correlation is driven by the fact that good routes for explorers moving on foot, horseback, or wagons are likely to be good routes for cars.

Exogeneity of this variable rests on the assumption that explorers' choice of routes are not related to anything that affects the innovation activity of regions a few centuries in the future, save the suitability of a place for roads. As reported in Duranton et al. (2013), the motivations for these expeditions were very different: searching for gold, the establishment of fur trading territories, finding emigration routes to Oregon, or expanding the U.S. territory towards the Pacific Ocean.

There is a concern that exploration routes may be more prominent in the presence of rivers or lakes that in turn may generate persistent differences in regional productivity. To address this issue, we include in our regressions a number of direct controls for the geography of the region (e.g., the share of MSA land that overlays an aquifer, MSA elevation range, an index of terrain ruggedness, heating and cooling degree days).

4.2 Summary Statistics

We focus on two units of analysis. First, we study cross-region variation and use MSAs as our unit of analysis (e.g., Rochester, NY). Then, we turn our attention to cross-region and technology variation and use MSA-class as our unit of analysis (e.g., Rochester, NY - electronics). Following Duranton and Turner (2012), we drop MSAs with no interstate highways in 1983. We also drop MSA-classes with no inventors in 1983. This leaves us with 220 observations in the MSA sample and 814 observations in the MSA-class sample.

Table 1 reports summary statistics for the sample employed in the MSA-level analysis. The average MSA in our sample generates 165 patents in 1983 and 229 patents in 1988. This represents an average annual growth rate of 6.7% per year. In terms of citation-weighted patents, the average MSA in our sample generates 2661 in 1983 and 4438.5 in 1988, reflecting an average annual growth rate of 7.8% per year. The average MSA has roughly 247 kilometers of interstate highway and approximately 390 inventors in 1983. We similarly report key descriptive statistics for the MSA-Class unit of analysis.

5 Regional Innovation Growth

We start our analysis by documenting the strong positive impact of regional highway stock on regional innovation. Our first set of results confirms the positive effect of roads on economic growth unveiled in Duranton and Turner (2012). The key difference with their analysis is that we look at economic growth through the lens of innovation outcomes whereas Duranton and Turner (2012) exploit employment data.

Columns (1) and (2) in Table 2 contain our first set of results, which show a robust positive association between highways and innovation in MSA-level regressions. We estimate these models using OLS with robust standard errors. In Column (1), the dependent variable is the logarithm of the citation-weighted patent count or, equivalently,

the logarithm of total forward citation count for issued patents applied for by all inventors in the MSA in the year 1988. Column (1) shows a positive correlation between interstate highway kilometers in 1983 and the level of innovation in 1988, controlling for the count of citation-weighted patents in 1983, the number of inventors in the MSA in 1983, 1978, and 1973, and a number of geography variables (the share of MSA land that overlays an aquifer, MSA elevation range, an index of terrain ruggedness, heating and cooling degree days). The specification in Column (2) is similar to the one in Column (1), but innovation is measured with un-weighted patent counts. Overall, these regressions show a strong positive correlation between transportation infrastructure and regional innovation. The magnitude of the coefficient in Column (1) shows that a 10% increase in interstate highway stock is associated with a 1.3% increase in innovative output. In Columns (3) and (4), to account for across-MSA technological heterogeneity, we move to a more disaggregated level and study the association between interstate highways and innovation at the MSA-class level. We cluster standard errors at the MSA level in these regressions as our main independent variable varies at the MSA level. Overall, the regressions in Columns (3) and (4) confirm at a more disaggregated level the main finding of the regressions at the MSA level: transportation infrastructure is positively associated with regional innovation.

The results in Table 2 are to be interpreted as correlations between road infrastructure and innovation, not causal impacts. As we argue above, there are a number of reasons why we expect unobservable factors to be correlated with both the levels of interstate highway and innovation in a region. This intuition is confirmed by a Hausman test that provides strong evidence against the exogeneity of transportation infrastructure. To address this endogeneity, we now turn to an instrumental variable estimation.

We examine the correlation between the historical variables and the stock of interstate highway in 1983, which is the key empirical variation exploited in our first stage

regressions. The table confirms the results in Baum-Snow (2007) and Duranton and Turner (2012), showing a large positive correlation between the stock of interstate highway in 1983 and the three instruments: 1947 planned interstate highway kilometers, 1898 kilometers of railroad, and the index of exploration routes between 1528-1850. The regressions show how each of these variables is strongly correlated with the endogenous variable and confirm the correlation when we include all three instruments. In unreported regressions we find that historical infrastructures are a strong predictor of modern-day highways stocks for multiple subsets of the sample. This suggests that the treatment effects we estimate below represent averages across a broad set of MSAs and thus can be interpreted as average treatment effects.⁴

Table 3 presents the IV regressions. In Column (1), we estimate the causal impact of the 1983 level of interstate highway in the MSA on MSA citations in 1988. The coefficient of 0.244 implies that 10% more interstate highways in 1983 leads to 2.44% more citation-weighted patents after five years. The regression controls for historical inventor levels and geographic variables. Column (2) confirms the results with un-weighted patent counts as measures of innovation. Across all specifications, the first stage F-statistics pass the weak instrument test, and the over-identification test (Hansen’s J statistic) gives a p-value of roughly 0.20, which supports the exogeneity of the instruments. Columns (3) and (4) confirm the positive impact of roads on regional innovation at the more disaggregated MSA-technology class level.

Across the specifications, IV estimates are larger than the corresponding OLS coefficients, indicating that endogeneity generates a downward bias. This downward bias is in line with other studies that investigate the impact of infrastructure on the

⁴Specifically, we find that the instruments are statistically significant at the 1 percent level in split-sample regressions across: (i) population quartiles; (ii) census division dummies; (iii) mean income quartiles and; (iv) share of employment in manufacturing quartiles.

economic growth of a region. To explain this difference between OLS and IV, Duranton and Turner (2012) show that MSAs that experienced negative population shocks tend to have larger road-building sectors. This suggests that the bias is driven by governments reacting to low employment with road building plans.

The estimated effect is large. Column (2) indicates that a 10% increase in interstate highways in 1983 leads to 1.7% more patents after five years. Estimates from the economics of innovation literature suggest an elasticity of corporate patenting to R&D expenditure close to 0.5 (see Aghion et al. (2013) and Bloom et al. (2013) for recent estimates). Therefore, a 1.7% increase in patenting is roughly equivalent to a 3.4% increase in regional corporate R&D investment.

We exploit these results to perform a few illustrative policy simulations that estimate the impact of enlarging the highway system in three representative metropolitan areas. We focus on one MSA with a large highway network (Los Angeles, CA with about 2000 km of highways in 1983), one with a medium-size network (Seattle, WA with about 500 km) and one with a small network (Madison, WI with roughly 100 km). We consider an increase in the highway network of 100 km and 250 km in each of these metropolitan areas. Moreover, by exploiting the figures reported in Kortum and Lerner (2000) on the R&D expenditure per patent in 1988, we transform each effect into an “equivalent R&D subsidy,” (i.e., the extra R&D investment required to generate an equivalent increase in patenting). These calculations suggest that the effects of transport infrastructure on innovation are not trivial. For example, a 100 km increase for the Los Angeles highway system appears roughly equivalent to a \$44 million R&D subsidy. Even in a small metropolitan area as Madison, a 100 km increase in highways is roughly equivalent to a \$17 million R&D subsidy. These estimates are only illustrative and should not be over-interpreted.

6 Local Knowledge Flows

We document a positive causal effect of interstate highways on regional innovation in Section 5. This finding is in line with the previous literature that has uncovered a positive effect of the stock of highways on urban growth (Duranton and Turner, 2012). In principle, there are many mechanisms through which transportation infrastructure affects the creation and diffusion of knowledge. An important economic channel emphasized in previous research is that transportation infrastructure generates regional growth through agglomeration economies, typically modelled as an inflow of new workers (Duranton and Turner, 2012). In this section, we provide evidence of a different mechanism through which roads may affect innovation and growth: their ability to ease the flow of local knowledge, which may serve as an important input to local innovation. A key feature of this channel is that it does not require an inflow of new innovators, and therefore it is conceptually different from traditional agglomeration forces.

We look at the impact of highways on knowledge flows within an MSA. More specifically, we study whether an increase in the stock of highways affects the way in which local innovators rely on each other’s knowledge to spur innovation. To this end, for each patent in an MSA class, we compute the distance between the location of the inventor and the location of the inventors of patents cited by the focal patent and located in the same MSA. For each MSA class, we then compute the average distance between the innovators and their within-MSA cited technologies. The average distance between a patent and its within-MSA citations in 1988 is 35.6 kilometers (std. dev. 21.7). For regions above the median level of highways in 1983, the distance is 46.0 kilometers compared to 23.5 kilometers for regions below the median.

We report regression results illustrating the impact of interstate highways on within-MSA citations distance in Table 4. Each regression controls for the level of patent-

ing in 1983, the average within-MSA citation distance in 1983, historical inventor levels, geographic variables, and technology field effects. Column (1) shows a strong positive effect of highways on citation distance. The estimate indicates that a 10% increase in 1983 highways causes a 2.3% increase in the average distance between innovators and the local inputs cited in their patents. To take into account that distance may depend on socio-economic and geographic characteristics of the MSA, in Column (2) we add to our control variables a set of additional geographic variables (in particular non-linear effects of the basic geographic measures and interactions). Despite the very large number of covariates in this specification, the results are robust. In Column (3), we show that results are similar if we add controls for historical population levels.

The regressions in Columns (1) to (3) show that innovators increase the distance traveled for local inputs in the presence of greater highway stock. This effect may arise mechanically if highway provision increases the dispersion of innovators. But it may also indicate easier access to more distant local knowledge, which generates greater diffusion of local knowledge. To better assess the impact of highways on local knowledge diffusion, in each MSA-class, we identify a set of *non-mover* MSA inventors. These are inventors who were active both in 1983 and in 1988 and who did not change their location over this five-year period. Columns (4) to (6) present results for this sample of *non-mover* inventors. The estimates are qualitatively and quantitatively similar to those we report in Columns (1) to (3). Specifically, these findings show that highway provision induces *non-mover* inventors to cite more distant *non-mover* local inventors. Overall, the fact that the impact of highways on citations among non-mover inventors is similar to the impact for the overall sample indicates that the highway effect is not mechanically driven by increasing dispersion of innovators but rather suggests that transportation infrastructure enables innovators to access more distance local knowledge.

We next investigate the extent to which highways impact the growth of patents

that build upon local knowledge. To do so, we identify all patents that cite at least one patent in the same MSA-class. In Table 5, we explore the relationship between road infrastructure and innovation that builds on local knowledge both in the full sample and in the sample of *non-mover* inventors. In addition, we contrast the propensity to build on local knowledge with the propensity to build on new sources of local knowledge (i.e., to cite a firm that was not cited by previous patents of the inventor). In Column (1), we show that a 10% increase in 1983 highways causes a 1.77% increase in patents that draw upon local knowledge. In Column (2), we show that a 10% increase in 1983 highways causes a 1.46% increase in patents that cite patents in the same MSA by firms new to the inventors. In Columns (3) and (4), we replicate the results in (1) and (2) but using our *non-mover* inventor sample. The estimates are qualitatively similar but smaller in magnitude.

Overall, the results in Tables 4 and 5 provide direct evidence that highways shape the propensity of innovators to rely on local knowledge. Local innovators appear more likely to rely on new and more distant local knowledge in the presence of greater transportation infrastructure. This suggests that an easier flow of local knowledge may be a significant mechanism through which road infrastructure affects local growth. Building on this insight, in Section 7 we present an illustrative estimation of a structural model that aims to quantify the relative importance of highways in terms of traditional agglomeration forces versus facilitating knowledge flows in generating productivity gains.

Our analysis has focused on the impact of interstate highways on local knowledge flows, measured by citations among inventors located in the same MSA. It is natural to expect interstate highways to also affect knowledge exchange between inventors across MSAs. In our regressions, we only consider local knowledge flows because the analysis of citation patterns within an MSA requires milder assumptions on the exogeneity of the historical instrumental variables. To analyze knowledge flows between two MSAs, we

need to assume that our IVs are not correlated with future patent citations between two regions. Railroads, expedition routes, and interstate highways were built and planned to connect principal metropolitan areas. In this respect, unobserved heterogeneity affecting the historical flow of people, goods, and knowledge between two MSAs may be associated with the presence of railroads, routes, or planned highways connecting them and may have a long-lasting impact correlated with future knowledge flow between the two MSAs. By focusing on local (i.e., within MSA) knowledge flows, our empirical analysis rests on the weaker assumption that the historical instrumental variables are not correlated with future citation patterns among inventors located in the same MSA.

6.1 Patent Level Analysis

The regressions presented in the previous sections rely on data aggregation at the MSA or MSA-class level, a standard approach in the urban economics literature. In this subsection, following a familiar structure to the economics of innovation literature, we move to patent-level regressions in order to further study how the provision of transportation infrastructure affects local knowledge spillovers. This finer unit of analysis allows us to introduce a large number of additional explanatory variables measuring characteristics of the patent, its inventors, and its citations. These additional controls reduce the level of unobserved heterogeneity and limit the likelihood of violation of the exclusion restriction.

To perform the patent level analysis, we identify all the local citations made by granted patents with application year 1988 (excluding self-citations). This leads to a sample of 10,776 citations from 1988 patents to other patents located in the same MSA. Exploiting these data, we follow two approaches to study local knowledge flows at the patent level. In Table 6, we explore the effect of highways on the distance of local citations. As in our previous analysis, we instrument Highways with the historical measures

and cluster standard errors at the MSA level. All regressions control for two-digit technology effects of the citing patent, the cited patent, and grant year effects of the cited patent. Because of the introduction of these dummy variables, the sample drops to 9,141 observations since some patents cannot be mapped to two-digit NBER classifications. We include controls for geographic characteristics of the MSA in Column (2). In Column (3), we control for the number of MSA inventors. Across all specifications, we find a strong positive effect of transportation infrastructure on the distance between a patent and its local citations. The marginal effect in Column (3) indicates that a 10% increase in highway stock increases the distance of local citations by 1.26%.

A possible interpretation is that the presence of a large highway stock disproportionately attracts inventors from technology fields that benefit less from close locations. To address this concern, we exploit our data to run a placebo test. Consider a local citation in our sample from patent p belonging to technology field k_1 to patent q in technology field k_2 . For each citation, we identify local citations made by other 1988 patents in field k_1 to patents in technology field k_2 , which are located in other MSAs and compute the average citation distance. In Column (4), we estimate the effect of highways on the distance of citations by patents in the same technology field but located in different MSAs. The coefficient is small and statistically insignificant. This exercise shows that the (instrumented) highway stock variable is uncorrelated with distance at the technology class level and confirms the exogeneity of the historical transportation infrastructure.

As a second approach, rather than estimating the effect of roads on the distance between inventors and their local inputs we hold constant the distance and examine the extent to which the probability of a citation between inventors increases with increasing transportation infrastructure. To do this we use the empirical methodology developed by Jaffe et al. (1993), which has become a classic approach in the economics of innovation

literature. The idea is to compare the characteristics of local patents cited by 1988 patents and a control group of non-cited local patents of the same cohort. The control group is constructed as follows. For each local citation, we randomly select another local patent that is not cited by the focal patent but has the same application year and three-digit patent classification. Following Jaffe et al. (1993) and Belenzon and Schankerman (2013), we run a series of linear probability models that relate a dummy variable for whether a patent is cited to a set of control variables. The specification is:

$$CitationDummy_{pqj} = \alpha + \beta_1 \log Highway_{j1983} + \beta_2 \log Distance_{pq} + \theta X_{pqj} + \epsilon_{pqj}, \quad (3)$$

where *CitationDummy* is an indicator variable that equals one if patent p from MSA j cites patent q also located in MSA j . The additional controls, X , include dummies for the technology field of the focal patent, dummies for the tech field, and grant year of the cited/control patents.

The estimates of these regressions are reported in Table 7. Across all specifications, we find a strong significant negative association between distance and the citation dummy. This result confirms the findings in Jaffe et al. (1993) and Belenzon and Schankerman (2013) and is typically interpreted as evidence that knowledge spillovers are geographically bounded. In Columns (1) and (2), we also document a positive association between citation and the local stock of highways. Notice that, by construction, the rough correlation between the MSA Highway stock and the citation dummy is zero in our sample. The positive coefficient on MSA Highway in Column (1) indicates that, conditioning on the distance between two inventors located in the same MSA, a citation is more likely when the stock of highways in the region is greater. In Column (2), we show that the correlation between highways and citations decreases in magnitude when

we control for the size of innovative activity in the region (i.e., the number of inventors in the MSA in 1983-1978-1973) and MSA geography characteristics. Nonetheless, the coefficient remains statistically significant at the 0.05 level. In Columns (3) and (4), we exploit our historical instruments to address the potential endogeneity of the highway stock. Results are robust, confirming the positive effect of transportation infrastructure on local knowledge flows. The estimates in Column (4) imply that a 10% increase in highway stock increases the citation probability by 0.12 percentage points, which is 0.24% of the mean citation probability.

We confirm the robustness of these findings in a variety of unreported regressions. First, following Belenzon and Schankerman (2013), we replace the distance measure with a flexible specification that allows for non-linear effects of distance. Specifically, we employ five dummy variables for quintile intervals of distance. The coefficient remains very similar to the one reported in Column (4) of Table 7. Second, we confirm robustness to introducing socio-economic controls and census division dummies. Third, we find that the effect is larger (almost double) if we drop patents from MSAs without highway stock. Fourth, we show that results are qualitatively and quantitatively similar when we replace the highway measure with the alternative lane-weighted measure developed by Duranton and Turner (2012).

7 Roads and Labor Productivity: Distinguishing the Role of Knowledge from other Agglomeration Economies

We follow Kline and Moretti (2014) and model an MSA as a small, open economy where firms take as given the prices of capital K , labor L , and output Y . The utility of workers

is a function of wages w and amenities M that takes the following specification:

$$U(w, M) = \ln w + \ln M. \quad (4)$$

Output is produced with a Cobb-Douglas technology:

$$Y = AK^\alpha F^\beta L^{1-\alpha-\beta},$$

where F is a fixed factor and A is total factor productivity. Capital is assumed to be perfectly mobile. If we normalize the price of Y to 1 (sold on global market), denote with r the (nationwide) cost of capital and with w the wage, the model implies the following inverse labor demand curve:

$$\ln w = \Theta - \frac{\beta}{1-\alpha} \ln L + \frac{1}{1-\alpha} \ln A, \quad (5)$$

where $\Theta = \ln(1 - \beta - \alpha) + (\alpha \ln \alpha + \beta \ln F - \alpha \ln r)/(1 - \alpha)$. We assume that the total factor productivity depends on two variables: patents, P , and labor, L , according to the following specification:

$$\ln A = \zeta \ln P + \sigma \ln L. \quad (6)$$

The parameter σ captures the strength of agglomeration economies, a concept studied and documented in the urban economics literature (Rosenthal and Strange (2004); Duranton and Turner (2012); Kline and Moretti (2014)). The parameter ζ describes the impact of patenting on productivity, a concept studied and documented by innovation economists (Bloom and Van Reenen (2002); Furman et al. (2002)).

Our reduced form analysis, together with the findings by Duranton and Turner (2012), indicate that both P and L are affected by the provision of roads R . We incor-

porate this effect by assuming a simple (reduced-form) specification for patenting and labor supply with roads as the only input:

$$L = bR^\mu \tag{7}$$

$$P = aR^\theta. \tag{8}$$

These simple functional forms are sufficient to highlight the differential impact of roads on productivity through agglomeration and innovation channels. Moreover, by assuming that labor supply only depends on roads, the impact of roads on wages can be interpreted as a welfare effect without having to specify a worker migration model. Duranton and Turner (2012) explain that typically migration models with sticky labor adjustments lead to reduced form equations similar to (7). In the Appendix, we extend the model incorporating additional inputs in the patent production function (8) as well as relaxing the inelastic labor supply function (7) introducing more structure in the workers' migration process.

Combining the above formulas we obtain:

$$\frac{d \ln w}{d \ln R} = -\frac{\beta\mu}{1-\alpha} + \frac{\zeta\theta}{1-\alpha} + \frac{\mu\sigma}{1-\alpha},$$

which decomposes the impact of roads on wages into: (i) a competitive effect (negative term) and (ii) productivity effect (positive terms) that arises from the impact of roads on patenting and an additional agglomeration effect. Specifically, roads have three distinct effects on wages. First, they attract labor, which reduces wages. Second, they facilitate knowledge flows, which lead to greater patenting and higher productivity that increases wages. Third, they increase productivity through agglomeration, which increases wages.⁵

⁵In this simple model, the combination of these three effects is also equal to the impact of roads on

Our regression estimates, together with those of Duranton and Turner (2012), provide natural structural estimations for $\mu = 0.15$ and $\theta = 0.24$. Following Kline and Moretti (2014), we set $\beta = 0.47$ and $\alpha = 0.68$. We obtain the elasticity of TFP respect to patents, ς , from Furman et al. (2002) that estimate $\zeta = 0.11$. Exploiting these parameters, we can rewrite:

$$\frac{d \ln w}{d \ln R} = -0.22 + 0.08 + 0.47\sigma.$$

This shows that the impact of roads on wages crucially depends on the strength of agglomeration forces. For example, in the absence of agglomeration economies ($\sigma = 0$), the model would predict a road elasticity of wages equal to -0.14. Estimates in the literature range from $\sigma = 0.03$ (Henderson (2003)) to $\sigma = 1.25$ (Greenstone et al. (2010)), which imply elasticities of -0.13 and 0.45, respectively.

Duranton and Turner (2012) estimate a labor elasticity of wages equal to 0.03 that combined with their estimate of $\mu = 0.15$ implies a road elasticity of wages equal to 0.2 and a corresponding $\sigma = 0.72$ that is roughly in the middle of the estimates in the literature. As illustrated in Table 10, with this parametrization for σ , we find that the total productivity effect is 0.42 and that the patenting channel accounts for 19% of this effect.

In Table 10, we also illustrate the decomposition in three extensions of the baseline model. We discuss the details of each model in the Appendix. In the first extension, we relax the inelastic labor supply function (7) microfounding the workers' migration process. Following Duranton and Turner (2012), we assume that there is a pool of people in the rural area that receives utility \bar{U} and that cities draw their new workers from this

welfare if road construction is not costly. This equivalence does not hold in the extensions where more structure is imposed on the migration process.

rural pool. Duranton and Turner (2012) explain how this assumption is consistent with U.S. data showing that most immigration to cities is drawn from rural areas and from abroad. We also extend (4) to allow roads to increase the attractiveness of a city by reducing travel costs:

$$U(w, M, R) = \ln w + \ln M + t \ln R.$$

In this model, migration occurs until utility between residents and non-residents is equalized. The positive impact that roads have on the utility of residents is compensated by a reduction in wages triggered by migration. Calibration of this model leads to a lower estimate for the productivity effect with patenting explaining 50% of it.

The second extension considers an alternative patent production function (8), allowing labor, L , to affect patenting. Specifically, we assume:

$$\ln P = \theta \ln R + \lambda \ln L.$$

Combining this formula with (7), (6) and (5), we obtain a slightly stronger patenting effect that now explains 21% of the productivity effect.

The final extension combines the migration process with the alternative patent production function. The role of patenting is even more pronounced in this setting because of the lower impact of agglomeration economies in the migration model. The estimates imply that 56% of the productivity effect is due to patenting.

These calculations are only illustrative and should not be over-interpreted. Nonetheless, they show how roads may affect productivity through multiple channels and that non-agglomeration forces may explain an important fraction of the productivity gains generated by transportation infrastructure.

8 Concluding Remarks

We estimate the causal effect of within-MSA interstate highways on regional innovation. The identification strategy exploits variation in historical data on planned portions of the interstate highway system, railroads, and exploration routes. There are two key findings. First, in terms of the magnitude of the main effect, a 10% increase in a region's stock of highways causes a 1.7% increase in regional innovation growth over a five-year period. Second, in terms of the mechanism, transportation infrastructure facilitates the flow of local knowledge by lowering the cost and thus increasing the returns to accessing more distant local knowledge inputs. This finding suggests that roads may spur regional growth even in the absence of agglomeration economies that arise from the inflow of new workers, the mechanism typically considered in the literature.

Our findings have implications for policy makers. They suggest that the set of tools available to spur regional innovation are much broader than targeted R&D subsidies and tax credits and may include the provision of infrastructure that facilitates the flow of knowledge. Our analysis also suggests that the returns to particular regional innovation policies (e.g., cultivating entrepreneurial ventures or organizing business innovation summits) may vary across regions and depend on the availability of transportation infrastructure.

References

- Aghion, P., J. Van Reenen, and L. Zingales (2013). Innovation and institutional ownership. *American Economic Review* 103(1), 277–304.
- Agrawal, A. and I. Cockburn (2003, November). The anchor tenant hypothesis: exploring the role of large, local, r&d-intensive firms in regional innovation systems. *International Journal of Industrial Organization* 21(9), 1227–1253.
- Agrawal, A., I. Cockburn, A. Galasso, and A. Oettl (2014). Why are some regions more innovative than others? The role of small firms in the presence of large labs. *Journal of Urban Economics* 81(0), 149 – 165.
- Agrawal, A., D. Kapur, and J. McHale (2008, September). How do spatial and social proximity influence knowledge flows? Evidence from patent data. *Journal of Urban Economics* 64(2), 258–269.
- Agrawal, A. K., J. McHale, and A. Oettl (2014, March). Why stars matter. Working Paper 20012, National Bureau of Economic Research.
- Baum-Snow, N. (2007, 05). Did Highways Cause Suburbanization? *The Quarterly Journal of Economics* 122(2), 775–805.
- Baum-Snow, N. (2013). Urban Transport Expansions, Employment Decentralization, and the Spatial Scope of Agglomeration Economies. Working Paper.
- Belenzon, S. and M. Schankerman (2013). Spreading the word: Geography, policy and knowledge spillovers. *Review of Economics and Statistics* 95(3), 884–903.
- Bloom, N., M. Schankerman, and J. Van Reenen (2013). Identifying technology spillovers and product market rivalry. *Econometrica* 81(4), 1347–1393.
- Bloom, N. and J. Van Reenen (2002). Patents, real options and firm performance. *The Economic Journal* 112(478), C97–C116.
- Burchfield, M., H. G. Overman, D. Puga, and M. A. Turner (2006). Causes of sprawl: A portrait from space. *The Quarterly Journal of Economics* 121(2), 587–633.
- Caballero, R. J. and A. B. Jaffe (1993). How high are the giants’ shoulders: An empirical assessment of knowledge spillovers and creative destruction in a model of economic growth. *NBER Macroeconomics Annual* 8, pp. 15–74.
- Carlino, G. and W. Kerr (2014). Agglomeration and innovation. NBER working paper 20367.
- Catalini, C. (2013). Microgeography and the direction of inventive activity. MIT Sloan working paper.

- Cutler, D. M. and E. L. Glaeser (1997). Are ghettos good or bad? *The Quarterly Journal of Economics* 112(3), 827–872.
- Donaldson, D. (2014). Railroads of the raj: Estimating the impact of transportation infrastructure. *American Economic Review*, Forthcoming.
- Duranton, G., P. Morrow, and M. A. Turner (2013). Roads and trade: Evidence from the us. *The Review of Economic Studies*, Forthcoming.
- Duranton, G. and M. A. Turner (2012). Urban Growth and Transportation. *Review of Economic Studies* 79(4), 1407–1440.
- Feldman, M. P. (1994). *The Geography of Innovation*. Dordrecht, Kluwer.
- Feldman, M. P. and D. B. Audretsch (1999). Innovation in cities: Science-based diversity, specialization and localized competition. *European Economic Review* 43(2), 409 – 429.
- Fernald, J. G. (1999, June). Roads to Prosperity? Assessing the Link between Public Capital and Productivity. *American Economic Review* 89(3), 619–638.
- Furman, J. L., M. E. Porter, and S. Stern (2002). The determinants of national innovative capacity. *Research policy* 31(6), 899–933.
- Galasso, A., M. Schankerman, and C. Serrano (2013). Trading and enforcing patent rights. *RAND Journal of Economics* 44, 275–312.
- Glaeser, E. L., S. Pekkala Kerr, and W. Kerr (2012). Entrepreneurship and urban growth: An empirical assessment with historical mines. Harvard Business School Working Papers 13-015, Harvard Business School.
- Gray, C. (1989). *Gray's New Trunk Railway Map of the United States*. New York, Library of Congress Catalogue.
- Greenstone, M., R. Hornbeck, and E. Moretti (2010). Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. *Journal of Political Economy* 118(3), 536–598.
- Griliches, Z. (1990). Patent statistics as economic indicators: A survey. *Journal of Economic Literature* 28(4), 1661–1707.
- Hall, B. H., A. B. Jaffe, and M. Trajtenberg (2001, October). The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools. Working Paper 8498, National Bureau of Economic Research.
- Henderson, V. (2003). The urbanization process and economic growth: The so-what question. *Journal of Economic Growth* 8(1), 47–71.

- Jacobs, J. (1969). *The Economy of Cities*. Random House: New York.
- Jaffe, A. B., M. Trajtenberg, and R. Henderson (1993, August). Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations. *The Quarterly Journal of Economics* 108(3), 577–98.
- Kerr, W. and S. Kominers (2015). Agglomerative forces and cluster shapes. *Review of Economics and Statistics*. Forthcoming.
- Kline, P. and E. Moretti (2014). Local economic development, agglomeration economies, and the big push: 100 years of evidence from the tennessee valley authority. *The Quarterly Journal of Economics* 129(1), 275–331.
- Kortum, S. and J. Lerner (2000). Assessing the contribution of venture capital to innovation. *RAND Journal of Economics* 31(4), 674–692.
- Li, G.-C., R. Lai, A. D’Amour, D. M. Doolin, Y. Sun, V. I. Torvik, A. Z. Yu, and L. Fleming (2014). Disambiguation and co-authorship networks of the u.s. patent inventor database (1975–2010). *Research Policy* 43(6), 941 – 955.
- Marx, M., D. Strumsky, and L. Fleming (2009, June). Mobility, skills, and the michigan non-compete experiment. *Manage. Sci.* 55(6), 875–889.
- Mehta, A., M. Rysman, and T. Simcoe (2010). Identifying the age profile of patent citations: new estimates of knowledge diffusion. *Journal of Applied Econometrics* 25(7), 1179–1204.
- Michaels, G. (2008, November). 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. (2011). Local Labor Markets. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Volume 4, Chapter 14, pp. 1237–1313. Elsevier.
- Pakes, A. and Z. Griliches (1980). Patents and R&D at the firm level: A first report. *Economics Letters* 5(4), 377–381.
- Romer, P. M. (1986, October). Increasing Returns and Long-run Growth. *Journal of Political Economy* 94(5), 1002–37.
- Rosenthal, S. S. and W. C. Strange (2004). Evidence on the nature and sources of agglomeration economies. In J. V. Henderson and J.-F. Thisse (Eds.), *Handbook of Regional and Urban Economics*, Volume 4, pp. 2119–2171. Elsevier.
- Singh, J. and M. Marx (2013). Geographic constraints on knowledge spillovers: Political borders vs. spatial proximity. *Management Science* 59(9), 2056–2078.
- US Bureau of Transportation Statistics (2010). Transportation statistics annual report.

Table 1: Summary Statistics

Unit of Analysis	Variables	mean	Std. Dev.
MSA N = 220	MSA Weighted Patents ₁₉₈₈	4438.53	12774.52
	MSA Patents ₁₉₈₈	228.69	609.99
	MSA Weighted Patents ₁₉₈₃	2660.96	8154.20
	MSA Patents ₁₉₈₃	165.00	482.13
	MSA Inventors ₁₉₈₃	390.05	1175.62
	MSA Highway ₁₉₈₃ (km)	247.30	300.36
	1947 Planned Highways (km)	118.46	129.47
	1898 Railroads (km)	290.19	301.66
	1528-1850 Exploration route index	2990.63	4277.54
MSA-Class N = 814	MSA-Class Weighted Patents ₁₉₈₈	993.04	2386.91
	MSA-Class Patents ₁₉₈₈	50.93	116.41
	MSA-Class Weighted Patents ₁₉₈₃	587.63	1479.83
	MSA-Class Patents ₁₉₈₃	36.44	92.85
	MSA-Class Inventors ₁₉₈₃	105.42	281.59

Notes: It may be that the means of the MSA-level variables are greater than six times the means of the MSA-class-level variables as some patents cannot be disaggregated to technology classifications.

Table 2: Roads are associated with more citations and patents - OLS Regressions

	(1)	(2)	(3)	(4)
Unit of Analysis	MSA		MSA-Class	
Dependent Variable	$\log\text{Cites}_{m,1988}$	$\log\text{Patents}_{m,1988}$	$\log\text{Cites}_{m,c,1988}$	$\log\text{Patents}_{m,c,1988}$
$\log\text{Highway}_{m,1983}$	0.130** (0.064)	0.097** (0.038)	0.250*** (0.082)	0.149*** (0.043)
$\log\text{Cites}_{m,1983}$	0.571*** (0.114)			
$\log\text{Patents}_{m,1983}$		0.762*** (0.073)		
$\log\text{Cites}_{m,c,1983}$			0.323*** (0.052)	
$\log\text{Patents}_{m,c,1983}$				0.517*** (0.048)
Inventor Controls	✓	✓	✓	✓
Geography Controls	✓	✓	✓	✓
Class Fixed Effects			✓	✓
Observations	220	220	814	814
R^2	0.878	0.940	0.715	0.861

Notes: All specifications are estimated by ordinary least squares. $\log\text{Cites}_{m,t}$ refers to the citation-weighted count of patents applied for (and subsequently granted) in period t in MSA m . $\log\text{Patents}_{m,t}$ refers to the count of patents applied for (and subsequently granted) in period t in MSA m . $\log\text{Cites}_{m,c,t}$ refers to the citation-weighted count of patents applied for (and subsequently granted) in period t in class c of MSA m . $\log\text{Patents}_{m,c,t}$ refers to the count of patents applied for (and subsequently granted) in period t in class c of MSA m . $\log\text{Highway}_{m,1983}$ refers to the 1983 level of interstate highway kilometers in MSA m . For MSA-level regressions, Inventor controls include the log of the inventors in the MSA in years 1973, 1978, and 1983. For MSA-class regressions, Inventor controls include the log of the inventors in the MSA-class in years 1973, 1978, and 1983 and the total number of inventors in the MSA in 1983. Geography controls include the share of each MSA's land that overlays an aquifer, MSA's elevation, index of MSA terrain ruggedness, and number of MSA heating and cooling degree days. We add a 1 to all patent, citation, and inventor count variables before taking the log to include observations with values of 0. Robust standard errors clustered at the MSA-level are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Roads cause an increase in citations and patents - IV Regressions

	(1)	(2)	(3)	(4)
Unit of Analysis	MSA		MSA-Class	
Dependent Variable	$\log\text{Cites}_{m,1988}$	$\log\text{Patents}_{m,1988}$	$\log\text{Cites}_{m,c,1988}$	$\log\text{Patents}_{m,c,1988}$
$\log\text{Highway}_{m,1983}$	0.244** (0.106)	0.170** (0.080)	0.347*** (0.105)	0.239*** (0.046)
Inventor Controls	✓	✓	✓	✓
Geography Controls	✓	✓	✓	✓
Class Fixed Effects			✓	✓
Observations	220	220	814	814
F-statistic	23.22	19.20	22.92	21.79
R^2	0.876	0.939	0.714	0.859

Notes: Robust standard errors clustered at the MSA-level are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Roads increase the geographic distance of local knowledge inputs

Sample	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable	All Inventors			Non-Moving Inventors		
	logSameMSAdistance _{m,c,1988}					
logHighway _{m,1983}	0.228*** (0.056)	0.217*** (0.055)	0.242*** (0.082)	0.111** (0.048)	0.146*** (0.052)	0.117* (0.066)
Class Fixed Effects	✓	✓	✓	✓	✓	✓
Inventors Controls	✓	✓	✓	✓	✓	✓
Geography Controls	✓	✓	✓	✓	✓	✓
Census Division Controls	✓	✓	✓	✓	✓	✓
Socio-Economic Controls		✓	✓		✓	✓
Extra Geography Controls		✓	✓		✓	✓
Population Controls			✓			✓
Observations	814	814	814	495	495	495
R^2	0.558	0.564	0.560	0.295	0.312	0.294
F-statistic	21.19	21.57	12.07	23.83	22.84	8.793

Notes: Robust standard errors clustered at the MSA-level are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Roads increase the number of patents that build upon local knowledge inputs

	(1)	(2)	(3)	(4)
Sample	All Inventors		Non-Moving Inventors	
Subsample	All Assignees	New Assignees	All Assignees	New Assignees
Dependent Variable	logSameMSApatents _{m,c,1988}			
logHighway ₁₉₈₃	0.177*** (0.042)	0.146*** (0.036)	0.059*** (0.023)	0.022** (0.011)
Inventor Controls	✓	✓	✓	✓
Geography Controls	✓	✓	✓	✓
Class Fixed Effects			✓	✓
Observations	814	814	814	814
F-statistic	22.81	22.30	22.84	22.69
R ²	0.826	0.821	0.693	0.437

Notes: Robust standard errors clustered at the MSA-level are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Roads increase the geographic distance of local knowledge inputs – Patent-level analysis

Dependent Variable	(1) logDistance _{pq}	(2) logDistance _{pq}	(3) logDistance _{pq}	(4) logDistance _{pq} other MSAs
logHighway _{m,1983}	0.407*** (0.061)	0.447*** (0.060)	0.126*** (0.049)	-0.006 (0.030)
Citing Patent Class Fixed Effects	✓	✓	✓	✓
Cited Patent Class Fixed Effects	✓	✓	✓	✓
Cited Patent Year Fixed Effects	✓	✓	✓	✓
Geography Controls		✓	✓	✓
Inventor Controls			✓	✓
Observations	9141	9141	9141	7593
R ²	0.129	0.142	0.176	0.155
F-stat	48.99	65.98	10.55	14.01

Notes: The unit of analysis for all specifications is the citing patent p - cited patent q dyad. All specifications are estimated by two-stage least squares. The sample consists of all citations made to patents in the same MSA by patents applied for (and subsequently granted) in 1988. $\log\text{Distance}_{pq}$ refers to the distance in kilometers between the first inventors of the citing patent p and cited patent q . The dependent variable in Column 4 consists of the mean distance of within-MSA citations for all patents applied for in 1988 (and subsequently granted) that are in the same technology class but in different MSAs from the focal MSA. We cluster robust standard errors at the MSA-level and present them in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Roads increase the probability of building upon local knowledge – Patent-level analysis

	(1)	(2)	(3)	(4)
Estimation	OLS		IV	
Dependent Variable	$\mathbb{1}(\text{Citation}_{pq})$			
logHighway _{<i>m</i>,1983}	0.030*** (0.009)	0.007** (0.004)	0.035*** (0.009)	0.012*** (0.004)
logDistance	-0.077*** (0.012)	-0.083*** (0.013)	-0.078*** (0.012)	-0.083*** (0.013)
Citing Patent Class Effects	✓	✓	✓	✓
Cited Patent Class Effects	✓	✓	✓	✓
Cited Patent Year Effects	✓	✓	✓	✓
Inventor Controls		✓		✓
Geography Controls		✓		✓
Observations	18282	18282	18282	18282
R^2	0.0381	0.0412	0.0381	0.0412
F-stat			45.63	10.53

Notes: We cluster robust standard errors at the MSA-level and present them in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Structural decomposition of the impact of roads on labor productivity

Model	Road Elasticity of Wages	Competitive Effect	Productivity Effect	Productivity Effect Explained by Innovation (percent)
Baseline	0.20	-0.22	0.42	19
Migration	-0.06	-0.22	0.16	50
Inventors	0.20	-0.22	0.42	21
Migration + Inventors	-0.06	-0.22	0.16	56

Notes: The table presents the estimates from three alternative structural models decomposing the impact of roads on wages through their interaction with labor competition and labor productivity. The last column of the table indicates the percentage of the productivity effect explained by greater innovation (as opposed to labor agglomeration).