ROADS AND INNOVATION

Ajay Agrawal, Alberto Galasso, and Alexander Oettl*

Abstract—We exploit historical data on planned highways, railroads, and exploration routes as sources of exogenous variation in order to estimate the effect of interstate highways on regional innovation: a 10% increase in a region's stock of highways causes a 1.7% increase in regional patenting over a five-year period. In terms of the mechanism, we report evidence that roads facilitate local knowledge flows, increasing the likelihood that innovators access knowledge inputs from local but more distant neighbors. Thus, transportation infrastructure may spur regional growth above and beyond the more commonly discussed agglomeration economies predicated on an inflow of new workers.

I. Introduction

THE literature linking transportation infrastructure to growth focuses on agglomeration economies as the mechanism. We report evidence that highlights a different mechanism. In addition to facilitating the flow of human capital into cities (agglomeration), transportation infrastructure, such as interstate highways, lowers the cost of knowledge flows within regions between local innovators. This finding is important because it sheds light on the black box of knowledge spillovers that lie at the heart of innovation and growth in the macroeconomic literature (Romer, 1986; Aghion and Howitt, 1992; Acemoglu & Akcigit, 2012).

One of the most important features of knowledge spillovers is that they are localized. Starting with the seminal work of Jaffe, Trajtenberg, and Henderson (1993), a number of studies have documented that spillovers are constrained by geography. This may explain some of the variation in productivity across regions. Moretti (2011) documents that after adjusting for skill composition, average wages in the highest- and lowest-paying U.S. metropolitan areas differ by approximately a factor of 3. Such dispersion is also evident when one compares innovation outcomes across regions (Agrawal et al., 2014; Carlino & Kerr, 2014). Silicon Valley and Boston are popular examples of outlier regions, significantly more productive than others in terms of innovation. Despite the prominence of such regional disparities, very little is known about the features of the economic and physical environment or the policy tools that affect knowledge flows and trigger economic growth through innovation.

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 have an impact on regional productivity through their effect on employment, private investment, and the returns to schooling. By increasing the circulation of people in a region, they are also likely to facilitate knowledge diffusion and spillovers. We illustrate this channel with an example in figure 1. The five largest knowledge-flow corridors in Boston (as measured by within-MSA citations from 1988 patents) largely coincide with the topology of the city's highway network, suggesting that roads may affect knowledge flow patterns. The effect of roads on knowledge creation and diffusion is the 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 (U.S. Bureau of Transportation Statistics, 2010), and about 20% of the income of the median U.S. household is devoted to road transportation (Duranton & Turner, 2012). Despite the magnitude of these investments, 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 into policies aimed at enhancing the flow of knowledge within cities. Furthermore, our findings offer insights 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 roads on innovation is the simultaneous determination of transportation infrastructure and regional technological development. For example, regional economic growth may boost local innovation and at the same time may induce local governments to invest in infrastructure. To address this problem, we follow a growing literature in urban economics that focuses on the U.S. interstate highways system and exploits instrumental variables for transportation infrastructure (Baum-Snow, 2007; Duranton & Turner, 2012; Duranton, Morrow, & Turner, 2014). While interstate highways affect local circulation of people and goods, they are predominantly built with nonlocal goals. Compared to other local transport infrastructure, such as local roads and subways, their construction is less likely correlated with regional economic shocks that confound empirical estimation. Building on Duranton and Turner (2012), we consider three instruments for the presence of highways. 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

Received for publication January 5, 2015. Revision accepted for publication February 5, 2016. Editor: Philippe Aghion.

*Agrawal: University of Toronto and NBER; Galasso: University of Toronto and NBER; Oettl: Georgia Institute of Technology.

We thank Nathaniel Baum-Snow, Mercedes Delgado, Chris Forman, Avi Goldfarb, Stephan Heblich, Mitchell Hoffman, Bill Kerr, Francesco Lissoni, Maureen McKelvey, Will Strange, and Matthew Turner, as well as seminar participants at UC Berkeley, Bordeaux, Bristol, Catholic University of Chile, DRUID, Georgia Tech, MIT, University of Toronto, and the NYU Economics of Strategy workshop for helpful comments. We gratefully acknowledge funding from the Centre for Innovation and Entrepreneurship at the University of Toronto, the Social Sciences and Humanities Research Council of Canada, and the Kauffman Junior Faculty Fellowship. All errors and omissions are our own.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00619.

The Review of Economics and Statistics, July 2017, 99(3): 417–434

doi:10.1162/REST_a_00619

Figure 1.—Main Knowledge-Flow Corridors of Boston

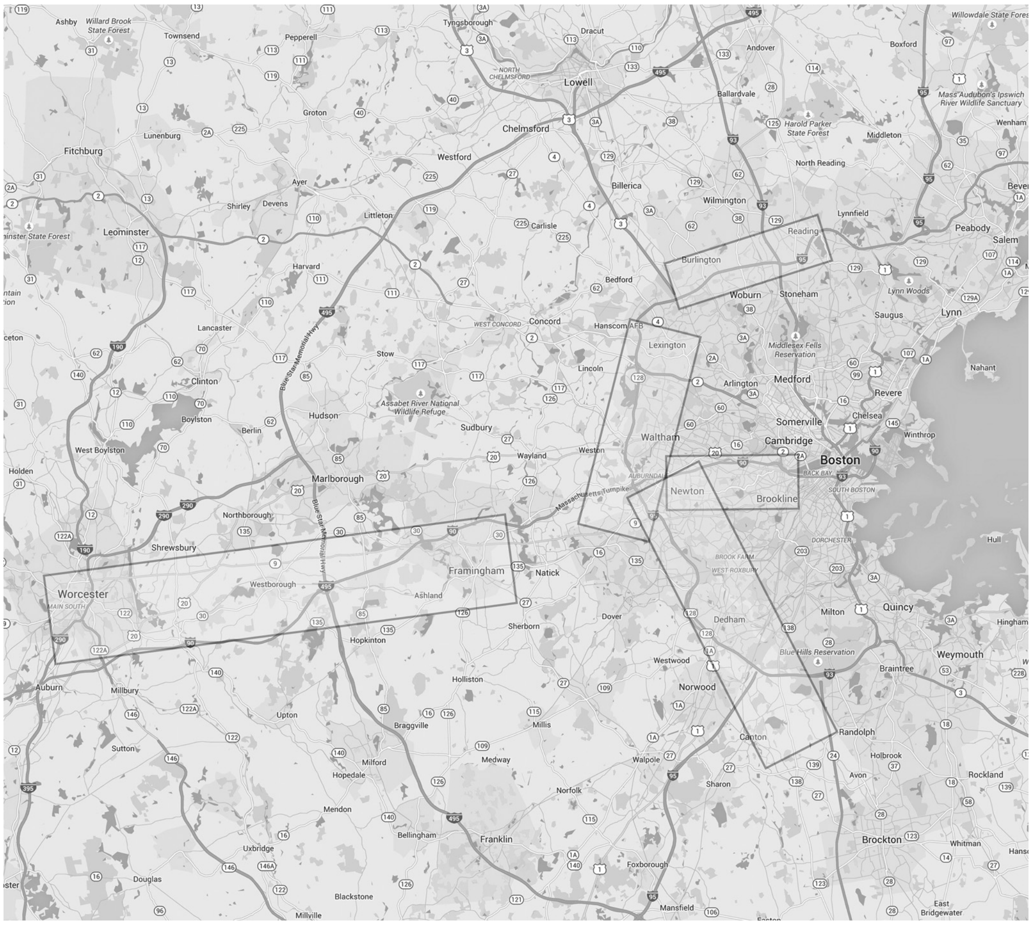

This figure graphically represents five of the largest within-region knowledge corridors in Boston in 1988. We identify these corridors by collecting all patents issued in 1988 where the first inventor is in the Boston MSA and finding all citations that were made to patents where the first inventor was also in the Boston MSA (there are 178 cities/towns within the Boston MSA). We did not examine within-city citations (e.g., Lexington-Lexington). To identify the largest corridors, we aggregate citations to the city-city dyad level (e.g., Worcester-Framingham).

network at the end of the nineteenth century. The third is based on maps of routes of major exploration expeditions of the United States from 1528 to 1850.

We find that a 10% increase in interstate highways leads to a roughly 1.7% increase in innovation as measured by 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 metropolitan statistical area (MSA) or MSA technology class as the units of analysis and that estimates are robust to the inclusion of a large set of control variables that explain persistent productivity differences across regions. We do not find evidence of negative externalities of the stock of highways in one region on innovation in neighboring regions. Moreover, the impact of transport infrastructure does not appear to decline with the diffusion of information and communication technologies.

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 affect innovation: local within-region knowledge flows. 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 & 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 nonmover 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 the provision of highways, consistent with the idea that larger firms are more likely to build on knowledge produced within their own boundaries and thus rely less on that produced by their neighbors.

We conclude 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 nonagglomeration 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 20% of the impact of roads on labor productivity may be due to nonagglomeration channels.

II. Related Literature

Our paper is connected to the literature on the determinants of regional innovation and 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 the 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, Kapur, and McHale (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 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.[1]

The emerging urban economics literature studies the impact of investments in transportation infrastructure on the evolution of metropolitan areas (Redding & Turner, 2015). 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. Chandra and Thompson (2000) study the impact of highways on nonmetropolitan counties and regions. They show that highways have a differential impact across industries and affect the spatial allocation of economic activity. 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.

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 twenty-year period. Duranton, Morrow, and Turner (2014) 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, Ghani, Goswami, and Kerr (2016) show that the Indian highway network had a strong impact on the growth of manufacturing activity.

III. Data

We follow Agrawal et al. (2014) in constructing our sample and begin with the set of 268 metropolitan statistical

[1] In terms of government policies, Marx, Strumkey, and Fleming (2009) show that regional noncompete 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, Schankerman, and Serrano (2013) show that state-level taxes have a strong impact on knowledge diffusion through the decision to trade patent rights.

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, Jaffe, and Trajtenberg (2001).

We obtain information on U.S. patenting activity and on the affiliation and location of patenting inventors in a region from the U.S. 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 with greater impact on social welfare and economic growth are more likely to be patented (Pakes & Griliches, 1980). 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 data set to match cities and townships to counties and, ultimately, MSAs. If a patent has at least one inventor from a particular MSA, we increment the counter for that MSA by 1. Thus, a patent with three inventors located in three different MSAs increments the patent counter for each of those MSAs by 1. However, if all three inventors are located in the same MSA, then the counter for that MSA is incremented by only 1.[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**$_{m,c,t}$: The forward citation weighted sum of distinct patents with primary technology classification c and application year t where at least one inventor is located in MSA m
>
> Patent citations identify prior knowledge on which a patent builds, and prior literature (starting with Pakes & Griliches, 1980) has often employed the number of 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**$_{m,c,t}$: The number of distinct patents with primary technology classification c and application year t where at least one inventor is located in MSA m

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

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, extensively described in Duranton and Turner (2012). All of 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 sociodemographic variables from the 1980 census for each MSA: the share of poor population, the share of college graduates, the share of population employed in manufacturing, and mean income. 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 U.S. Geological Survey and calculate distances between cities using the great circle method as in Singh and Marx (2013).

IV. Empirical Framework

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

$$\log Y_{m,c,1988} = \alpha + \beta \log Highway_{m,1983} + \gamma \log Y_{m,c,1983} + \theta X_{m,c} + \epsilon_{m,c}, \tag{1}$$

where $Y_{m,c,1983}$ is the innovation level in 1983 and $X_{m,c}$ is a vector of additional controls.[3]

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

[3] Lagged dependent variable models are common in the innovation literature because knowledge is a cumulative process and it is natural to consider current knowledge as an input for future knowledge (Aghion & Howitt, 1992). Our results are robust to dropping the lagged dependent variable from the model.

$0 < \gamma < 1$, then patenting in period $t+5$ will be equal to $K_{t+5} = BR_t^{\beta}K_t^{\gamma}$, where $\beta = a(1-\gamma)$ and $B = A^{1-\gamma}$. Taking logs, 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 equation (1) can be rewritten as

$$\begin{aligned}\log Y_{m,c,1988} &- \log Y_{m,c,1983}\\ &= \alpha + \beta \log Highway_{m,1983} + (\gamma - 1)\log Y_{m,c,1983}\\ &\quad + \theta X_m + \epsilon_{m,c}.\end{aligned}$$

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 to 1988.[4]

The main empirical challenge in estimating equation (1) is the possible correlation between unobservables, $\epsilon_{m,c}$, and the level of highways in a region. For example, the local government may react to an economic downturn by building more roads, thus generating 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 in section IVA.

A. *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; Duranton & Turner, 2012; Duranton et al., 2014), this empirical approach is novel in the entrepreneurship and innovation literature. The only exception we are aware of is Glaeser, Kerr, and Kerr (2012), who exploit historical mines as an instrument for entrepreneurship.

To be a valid instrument, a 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).

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. Duranton 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 were 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" (U.S. Federal Works Agency, Public Roads Administration, 1947 press release, cited in Michaels, 2008). Historical evidence discussed in Duranton et al. (2014) confirms that the 1947 highway plan was drawn to this mandate. Moreover, Duranton 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 Duranton 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 fewer. 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 tracks 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 leveling and grading a roadbed.

The U.S. rail network was developed in the middle of an industrial revolution and immediately after the Civil War. At that time, the U.S. 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

[4] Model (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 pretrends in knowledge creation between cities with different levels of highways in 1983. Nonetheless, instrumenting $Highway_{j1983}$ allows us to remove the bias generated by noncommon 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.

technology shocks in the 1980s. As discussed in Duranton and Turner (2012) and Duranton et al. (2014), railroads were developed mainly to transport grain, livestock, and lumber, and it is unlikely that such a flow of agricultural commodities was correlated to innovation activity in the 1980s. Moreover, railroads were typically constructed by private companies that expected to make profits in the short and medium terms.

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* (U.S. Geological Survey, 1970) reporting routes of major expeditions of exploration that occurred during the time period 1528 to 1850. From each map, they count 1 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' choices 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 Duranton et al. (2014) reported, the motivations for these expeditions were very different: searching for gold, establishing fur trading territories, finding emigration routes to Oregon, or expanding the U.S. territory toward 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).

B. *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, New York). Then, we turn our attention to cross-region and technology variation and use MSA class as our unit of analysis (e.g., Rochester, New York—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.—Summary Statistics

Unit of Analysis	Variables	Mean	SD
MSA	*MSA Weighted Patents*$_{1988}$	4,438.53	12,774.52
(N = 220)	*MSA Patents*$_{1988}$	228.69	609.99
	MSA Weighted Patents$_{1983}$	2,660.96	8,154.20
	MSA Patents$_{1983}$	165.00	482.13
	MSA Inventors$_{1983}$	390.05	1,175.62
	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	2,990.63	4,277.54
MSA Class	*MSA Class Weighted Patents*$_{1988}$	993.04	2,386.91
(N = 814)	*MSA Class Patents*$_{1988}$	50.93	116.41
	MSA Class Weighted Patents$_{1983}$	587.63	1,479.83
	MSA Class Patents$_{1983}$	36.44	92.85
	MSA Class Inventors$_{1983}$	105.42	281.59

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, an average annual growth rate of 6.7% per year. In terms of citation-weighted patents, the average MSA in our sample generates 2,661 cites in 1983 and 4,438.5 cites 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.

V. 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

TABLE 2.—ROADS ARE ASSOCIATED WITH MORE CITATIONS AND PATENTS: OLS REGRESSIONS

	(1)	(2)	(3)	(4)
Unit of Analysis	MSA		MSA Class	
Dependent Variable	$\log Cites_{m,1988}$	$\log Patents_{m,1988}$	$\log Cites_{m,c,1988}$	$\log Patents_{m,c,1988}$
$\log Highway_{m,1983}$	0.130**	0.097**	0.250***	0.149***
	(0.064)	(0.038)	(0.082)	(0.043)
$\log Cites_{m,1983}$	0.571***			
	(0.114)			
$\log Patents_{m,1983}$		0.762***		
		(0.073)		
$\log Cites_{m,c,1983}$			0.323***	
			(0.052)	
$\log Patents_{m,c,1983}$				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

All specifications are estimated by ordinary least squares. $\log Cites_{m,t}$ refers to the citation-weighted count of patents applied for (and subsequently granted) in period t in MSA m. $\log Patents_{m,t}$ refers to the count of patents applied for (and subsequently granted) in period t in MSA m. $\log Cites_{m}, c, t$ refers to the citation-weighted count of patents applied for (and subsequently granted) in period t in class c of MSA m. $\log Patents_{m}, c, t$ refers to the count of patents applied for (and subsequently granted) in period t in class c of MSA m. $\log Highway_{m,1983}$ 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 1973, 1978, and 1983. For MSA class regressions, Inventor controls include the log of the inventors in the MSA class in 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. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

the count of citation-weighted patents in 1983; the number of inventors in the MSA in 1983, 1978, and 1973; and a number of MSA geography variables (the share of land that overlays an aquifer, elevation range, an index of terrain ruggedness, and heating and cooling degree days). The specification in column 2 is similar to the one in column 1, but we measure innovation with unweighted 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 since 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, we expect unobservable factors to be correlated with both the levels of interstate highway and innovation in a region for a number of reasons. To address this endogeneity, we turn to an instrumental variable estimation.

We examine the correlation between the historical variables and the stock of interstate highway in 1983, the key empirical variation exploited in our first-stage regressions. We report these correlations in table B.1 in the online appendix. 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 railroad kilometers, and the index of exploration routes between 1528 and 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 highway 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.[5]

Table 3 presents the IV regressions. In column 1, we estimate the causal impact of the 1983 level of interstate highways 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 unweighted patent counts as measures of innovation. Across all specifications, the first-stage F-statistics pass the weak instrument test, and the overidentification 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

[5] Specifically, we find that the instruments are statistically significant at the 1% level in split-sample regressions across population quartiles, census division dummies, mean income quartiles, and share of employment in manufacturing quartiles.

TABLE 3.—ROADS CAUSE AN INCREASE IN CITATIONS AND PATENTS: IV REGRESSIONS

	(1)	(2)	(3)	(4)
Unit of Analysis	MSA		MSA Class	
Dependent Variable	$\log Cites_{m,1988}$	$\log Patents_{m,1988}$	$\log Cites_{m,c,1988}$	$\log Patents_{m,c,1988}$
$\log Highway_{m,1983}$	0.244**	0.170**	0.347***	0.239***
	(0.106)	(0.080)	(0.105)	(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

All specifications are estimated by two-stage least squares and control for the lagged dependent variables evaluated in 1983. $\log Cites_{m,t}$ refers to the citation-weighted count of patents applied for (and subsequently granted) in period t in MSA m. $\log Patents_{m,t}$ refers to the count of patents applied for (and subsequently granted) in period t in MSA m. $\log Cites_{m}, c, t$ refers to the citation-weighted count of patents applied for (and subsequently granted) in period t in class c of MSA m. $\log Patents_{m}, c, t$ refers to the count of patents applied for (and subsequently granted) in period t in class c of MSA m. $\log Highway_{m,1983}$ refers to the 1983 level of interstate highway kilometers in MSA m. The endogenous variable $\log Highway_{m,1983}$ is instrumented with $\log HighwayPlan_{m,1947}$, $\log Rail_{m,1898}$, and $\log ExplorationIndex_{m}$. For MSA regressions, Inventor controls include the log of the number of inventors in the MSA in 1973, 1978, and 1983. For MSA class regressions, Inventor controls include the log of the number of inventors in the MSA class in 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 elevation, index of MSA terrain ruggedness, and the 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. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

generates a downward bias. This downward bias is in line with other studies that investigate the impact of infrastructure on the 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, Van Reenen, & Zingales, 2013, and Bloom, Schankerman, & Van Reenen, 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, with about 2,000 km of highways in 1983), one with a medium-size network (Seattle, with about 500 km), and one with a small network (Madison, Wisconsin, with roughly 100 km). We consider an increase in the highway network of 100 km and 250 km in each of these metropolitan areas. We report the number of extra patents predicted by our IV estimates in each of these scenarios (appendix table B.2). 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 such 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 overinterpreted.

We perform a variety of tests to confirm the robustness of our main finding. In particular, we show that our baseline estimates are similar in models without lagged dependent variable, with state-fixed effects, and with state-class fixed effects. We also show that results are similar if we remove the five largest patenting MSAs from the sample. (Details for these regressions and additional robustness checks are provided in online appendix tables B.3 and B.4).

Following previous literature, our analysis focuses on the effect of interstate highways on the growth of innovation activity in the period 1983 to 1988, which precedes the large-scale diffusion of the Internet and other information and communication technologies (ICT). In principle, access to ICT may amplify or reduce the effect of roads depending on whether face-to-face interactions and ICT are complements or substitutes in knowledge production.

In the online appendix, we explore this issue with two distinct approaches. First, we contrast the magnitude of the effect across different time periods: 1983–1988, 1988–1993, 1993–1998, and 1998–2003. While the magnitude of the effect of transport infrastructures declines over time, the 1983 highway stock appears to have a long-lasting effect on innovation. For each of the estimates, we cannot reject at the 5% level that they are equal to our baseline effect. This evidence supports the idea that the impact of transportation infrastructure did not disappear in more recent time periods because of the diffusion of ICT.

Second, we collect data on the adoption of ICT across the MSAs in our sample. We obtain ICT data from Forman, Goldfarb, and Greenstein (2002), who construct measures of Internet adoption from the Harte Hanks Market Intelligence Survey. In the online appendix, we describe the data, and in appendix table B.5, we present regressions that include these internet measures. Our findings on the positive impact of transport infrastructure on innovation are robust. Coefficients are statistically and quantitatively similar to those in the baseline model.

A. *Displacement Effects*

The analysis has focused on local effects of transport infrastructure and does not consider the possibility that innovation can be displaced from one location to another. In principle, the provision of transport infrastructures can generate a zero-sum game among regions, with no effect on aggregate national innovation.

To explore this possibility, we follow Moretti and Wilson (2014) and extend our baseline model to include spatial lags. Specifically, we construct a new variable,

$$SpatialHighway_{j1983} = \sum_{i \neq j}^{I} w_{ij} \log(Highway_{j1983}), \qquad (2)$$

which is a weighted average of the highway stock in other MSAs. The weights w_{ij} are the elements of a spatial weighting matrix meant to capture the geographical proximity between pairs of MSAs. They satisfy $\sum_{i \neq j}^{I} w_{ij} = 1$ and are constructed using the inverse of the distance between MSAs' population centroids.

In table 4, we add this control to our baseline model instrumenting both the highway stock in an MSA and its spatial lag exploiting the historical variables. More precisely, we construct spatial weighted averages for the 1947 highway plan, railroad routes, and exploration expeditions and use them as instruments for the 1983 highway spatial lag. The coefficients of the direct effect are robust and stable. We find no evidence of statistically significant spatial spillovers across the various specifications. The MSA class level regressions in particular show that the effect is positive and its magnitude very small.[6]

We interpret these findings as suggesting that highways are unlikely to generate a zero-sum game for national innovation. Intuitively, this may arise because ideas are nonrival and lower communication or collaboration costs facilitate enhanced searching and matching among inventors. In other words, to the extent that a distant idea is more likely to be shared with an inventor in the focal region due to the presence of a highway does not diminish that idea's ability to also be used by someone else. This is also consistent with Duranton and Turner (2012), who show that changes in road infrastructure affect city employment growth by affecting driving within the city and not through market access, and with Duranton et al. (2014), who show that a city's road network does not affect the total value of intercity trade.

Nonetheless, we cannot fully account for the possibility that negative displacement effects as well as positive spillovers, arising from the increased physical availability of highways, may be present. Conducting such a precise evaluation of the national impact of local policies is very challenging and requires the estimation of a dynamic general equilibrium model of the U.S. economy during our time period with externalities across regions.[7]

VI. Local Knowledge Flows

We document a positive causal effect of interstate highways on regional innovation in section V. This finding is in line with the previous literature that has uncovered a positive effect of the stock of highways on urban growth (Duranton & 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 modeled as an inflow of new workers (Duranton & 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 (SD 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 5. Each regression controls for the level of patenting 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 socioeconomic and geographic characteristics of the MSA, in column 2 we add to our control variables a set

[6] In unreported results, we modify our spatial lag variable to vary at the MSA-technology-class level as spillovers may depend on the technological specialization of regions. Results are robust and essentially identical to those presented in table 4.

[7] The research of Kline and Moretti (2014) is an advance in this respect, with the structural estimation of the aggregate effects of one regional policy: the Tennessee Valley Authority. Their framework is not adequate for our data, in which each region is affected by a different policy (i.e., each region has a different level of road infrastructure).

TABLE 4.—ROADS CAUSE AN INCREASE IN LOCAL CITATIONS AND PATENTS, NET OF DISPLACEMENT EFFECTS

	(1)	(2)	(3)	(4)
Unit of Analysis	MSA		MSA Class	
Dependent Variable	$\log Cites_{m,1988}$	$\log Patents_{m,1988}$	$\log Cites_{m,c,1988}$	$\log Patents_{m,c,1988}$
$\log Highway_{m,1983}$	0.258**	0.172**	0.340***	0.237***
	(0.102)	(0.076)	(0.104)	(0.047)
$SpatialHighways_{m,1983}$	−0.326	−0.251	0.080	0.020
	(0.449)	(0.281)	(0.603)	(0.331)
Inventor controls	✓	✓	✓	✓
Geography controls	✓	✓	✓	✓
Class fixed effects			✓	✓
Observations	220	220	814	814
F-statistic	17.33	13.12	16.75	16.35
R^2	0.876	0.939	0.714	0.859

All specifications are estimated by two-stage least squares and control for the lagged dependent variables evaluated in 1983. $\log Cites_{m,t}$ refers to the citation-weighted count of patents applied for (and subsequently granted) in period t in MSA m. $\log Patents_{m,t}$ refers to the count of patents applied for (and subsequently granted) in period t in MSA m. $\log Cites_m, c, t$ refers to the citation-weighted count of patents applied for (and subsequently granted) in period t in class c of MSA m. $\log Patents_m, c, t$ refers to the count of patents applied for (and subsequently granted) in period t in class c of MSA m. $\log Highway_{m,1983}$ refers to the 1983 level of interstate highway kilometers in MSA m. $SpatialHighways_{m,1983}$ refers to the spatially weighted Highways measure defined as $SpatialHighway_{m,1983} = \sum_{i \neq m}^{I} w_{im} \log(Highway_{m,1983})$, which is a weighted average of the highway stock in other MSAs. The weights w_{im} are the elements of a spatial weighting matrix meant to capture the geographical proximity between pairs of MSAs. They satisfy $\sum_{i \neq m}^{I} w_{im} = 1$ and are constructed using the inverse of the distance between MSAs' population centroids. The two endogenous variables $\log Highway_{m,1983}$ and $SpatialHighways_{m,1983}$ are instrumented with $\log HighwayPlan_{m,1947}$, $\log Rail_{m,1898}$, and $\log ExplorationIndex_m$ and these three instruments with their own spatial weights for a total of six instruments. For MSA regressions, Inventor controls include the log of the number of inventors in the MSA in 1973, 1978, and 1983. For MSA class regressions, Inventor controls include the log of the number of inventors in the MSA class in 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 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 5.—ROADS INCREASE THE GEOGRAPHIC DISTANCE OF LOCAL KNOWLEDGE INPUTS

	(1)	(2)	(3)	(4)	(5)	(6)
Sample	All Inventors			Nonmoving Inventors		
Dependent Variable	$\log SameMSAdistance_{m,c,1988}$					
$\log Highway_{m,1983}$	0.228***	0.217***	0.242***	0.111**	0.146***	0.117**
	(0.056)	(0.055)	(0.082)	(0.048)	(0.052)	(0.066)
Class fixed effects	✓	✓	✓	✓	✓	✓
Inventors' controls	✓	✓	✓	✓	✓	✓
Geography controls	✓	✓	✓	✓	✓	✓
Census division controls	✓	✓	✓	✓	✓	✓
Socioeconomic 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

The unit of analysis for all specifications is the MSA class. All specifications are estimated by two-stage least squares and control for the lagged dependent variables evaluated in 1983. Columns 1–3 consist of the full sample; columns 4–6 include only patents of inventors who did not move between 1983 and 1988. $\log SameMSAdistance_{m,c,t}$ refers to the mean for all patents in the focal MSA m in class c of 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 m. $\log Highway_{m,1983}$ refers to the 1983 level of interstate highway kilometers in the MSA. The endogenous variable $\log Highway_{m,1983}$ is instrumented with $\log HighwayPlan_{m,1947}$, $\log Rail_{m,1898}$, and $\log ExplorationIndex_m$. Inventor controls include the log of the inventors in the MSA class in 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 elevation, index of MSA terrain ruggedness, and number of MSA heating and cooling degree days. Census division controls include dummy variables for each of the nine census division. Socioeconomic controls (from the 1980 census) include the share of the poor population in the MSA, the share of college graduates, the share of population employed in manufacturing, mean income in the MSA, and a measure of housing segregation computed by Cutler and Glaeser (1997). Extra geography controls include the square of the share of each MSA's land that overlays an aquifer, the square of the MSA's elevation, the square of the MSA's terrain ruggedness index, and the product of the terrain ruggedness index and elevation. Population controls include the log of the MSA's population in 1920, 1930, 1940, and 1950. We add a 1 to all patent, citation, and inventor count variables before taking the log to include observations with values of 0. The number of usable observations is reduced in columns 4–6 as some MSA classes did not include any nonmoving inventors between 1983 and 1988 or inventors who patented in both 1983 and 1988. Robust standard errors clustered at the MSA level are in parentheses. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

of additional geographic variables (in particular, nonlinear 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 also may 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 nonmover MSA inventors. These are inventors active in both 1983 and 1988 who did not change their location over this five-year period. Columns 4 to 6 present results for this sample of nonmover 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 nonmover inventors to cite more distant nonmover local inventors. Overall, the fact that the impact of highways on citations among nonmover 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

TABLE 6.—ROADS INCREASE THE NUMBER OF PATENTS THAT BUILD ON LOCAL KNOWLEDGE INPUTS
Dependent Variable: $\log SameMSApatents_{m,c,1988}$

	(1)	(2)	(3)	(4)
Sample	All Inventors		Nonmoving Inventors	
Subsample	All Assignees	New Assignees	All Assignees	New Assignees
$\log Highway_{1983}$	0.177***	0.146***	0.059***	0.022**
	(0.042)	(0.036)	(0.023)	(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

The unit of analysis for all specifications is the MSA class. All specifications are estimated by two-stage least squares and control for the lagged dependent variables evaluated in 1983. Columns 1 and 2 consist of the full sample, while columns 3 and 4 include only patents of inventors who did not move between 1983 and 1988. The odd-numbered columns include patents assigned to all assignees, while the even-numbered columns include only assignees who are new to (have never been cited by) the citing inventors. $\log SameMSApatents_{m,c,t}$ refers to the number of patents in MSA class m, c that cite at least one patent in the same MSA m. $\log Highway_{m,1983}$ refers to the 1983 level of interstate highway kilometers in the MSA m. The endogenous variable $\log Highway_{m,1983}$ is instrumented with $\log HighwayPlan_{m,1947}$, $\log Rail_{m,1898}$, and $\log ExplorationIndex_m$. Inventor controls include the log of the number of inventors in the MSA class in 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 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. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

infrastructure enables innovators to access more distance local knowledge.

We next investigate the extent to which highways affect the growth of patents that build on local knowledge. To do so, we identify all patents that cite at least one patent in the same MSA class. In table 6, we explore the relationship between road infrastructure and innovation that builds on local knowledge in both the full sample and the sample of nonmover 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 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 on 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 columns 1 and 2 but using our nonmover inventor sample. The estimates are qualitatively similar but smaller in magnitude.

Overall, the results in tables 5 and 6 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 VII, 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 consider only 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.

A. *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 section, following a structure familiar 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 7, we explore the effect of highways on the distance of local citations. As in our previous analysis, we

TABLE 7.—ROADS INCREASE THE GEOGRAPHIC DISTANCE OF LOCAL KNOWLEDGE INPUTS: PATENT-LEVEL ANALYSIS

Dependent Variable	(1) log$Distance_{pq}$	(2) log$Distance_{pq}$	(3) log$Distance_{pq}$	(4) log$Distance_{pq}$ other MSAs
logHighway$_{m,1983}$	0.407***	0.431***	0.085**	−0.010
	(0.061)	(0.061)	(0.041)	(0.031)
Citing patent class fixed effects	✓	✓	✓	✓
Cited patent class fixed effects	✓	✓	✓	✓
Cited patent year fixed effects	✓	✓	✓	✓
Geography controls		✓	✓	✓
Inventor controls			✓	✓
Observations	9,141	9,141	9,141	7,593
R^2	0.128	0.143	0.178	0.154
F-statistic	48.82	55.75	8.35	11.44

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$Distance_{pq}$ 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. log$Highway_{m,1983}$ refers to the 1983 level of interstate highway kilometers in the MSA. The endogenous variable log$Highway_{m,1983}$ is instrumented with log$HighwayPlan_{m,1947}$, log$Rail_{m,1898}$, and log$ExplorationIndex_m$. Inventor controls include the log of the number of inventors in the MSA class in 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 elevation, index of MSA terrain ruggedness, and number of MSA heating and cooling degree days. We add a 1 to all inventor count variables before taking the log to include observations with values of 0. We cluster robust standard errors at the MSA level and present them in parentheses. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

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 0.85%.

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 c_1 to patent q in technology field c_2. For each citation, we identify local citations made by other 1988 patents in field c_1 to patents in technology field c_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 rises 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 noncited local patents of the same cohort. We construct the control group 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

$$CitationDummy_{pqj} = \alpha + \beta_1 \log Highway_{m,1983} + \beta_2 \log Distance_{pq} + \theta X_{pqm} + \epsilon_{pqm}, \quad (3)$$

where *CitationDummy* is an indicator variable that equals 1 if patent p from MSA m cites patent q also located in MSA m. 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.

We report the estimates of these regressions in table 8. 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 0 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

TABLE 8.—ROADS INCREASE THE PROBABILITY OF BUILDING ON LOCAL KNOWLEDGE: PATENT-LEVEL ANALYSIS
Dependent Variable: $\mathbb{1}(\text{Citation}_{pq})$

	(1)	(2)	(3)	(4)
Estimation	OLS		IV	
$\log Highway_{m,1983}$	0.030***	0.006**	0.035***	0.008**
	(0.009)	(0.003)	(0.009)	(0.004)
$\log Distance$	−0.077***	−0.083***	−0.078***	−0.083***
	(0.012)	(0.013)	(0.012)	(0.013)
Citing patent class effects	✓	✓	✓	✓
Cited patent class effects	✓	✓	✓	✓
Cited patent year effects	✓	✓	✓	✓
Inventor controls		✓		✓
Geography controls		✓		✓
Observations	18,282	18,282	18,282	18,282
R^2	0.038	0.041	0.038	0.041
F-statistic			45.45	8.36

The unit of analysis for all specifications is the citing patent p–cited patent q dyad. Columns 1–2 and 3–4 are estimated by ordinary and two-stage least squares, respectively. The sample consists of all citations made to patents in the same MSA by patents applied for (and subsequently granted) in 1988. For each realized citing-cited patent dyad, we also identify a control cited patent that has the same application year and three-digit USPTO technology classification as the realized cited patent but was not cited by the focal citing patent. The dependent variable is a dummy set to 1 if the citing patent–cited patent dyad is a realized citation and 0 if the citing patent–cited patent dyad is a control dyad. By construction, the mean of *Citation* is 0.5. $\log Highway_{1983}$ refers to the 1983 level of interstate highway kilometers in the MSA. The endogenous variable $\log Highway_{m,1983}$ is instrumented with $\log HighwayPlan_{m,1947}$, $\log Rail_{m,1898}$, and $\log ExplorationIndex_m$. $\log Distance$ refers to the distance in kilometers between the first inventors of the citing and cited patents. Inventor controls include the log of the number of inventors in the MSA class in 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 elevation, index of MSA terrain ruggedness, and the number of MSA heating and cooling degree days. We add a 1 to all inventor count variables before taking the log to include observations with values of 0. We cluster robust standard errors at the MSA level and present them in parentheses. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

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 0.08 percentage points, which is 0.16% of the mean citation probability.[8]

We also exploit a different approach that builds on the methodology developed in Duranton and Overman (2005) and Akcigit and Kerr (2010) to examine how the distance between local citations departs from random counterfactuals. More specifically, we undertake Monte Carlo simulations, where we construct a series of random citations between the patents in each MSA class. For each observed local citation in which patent p from MSA m cites patent q (also located in MSA m), we draw two local patents p' and q' with the same application years and technology classes of the citing pair. We include the original citation among the possible pool of local citations and draw with replacement. We measure from this simulated counterfactual the average distance between local citations in the MSA class and repeat this procedure 1,000 times.

Appendix table B.6 uses these simulations to provide additional evidence in support of our findings. First, the table shows that the mean distance between local citations in our data increases with the stock of interstate highway in the MSA. Second, it shows that the observed distance between a patentee and his local citation is on average below the distance between simulated citations. This supports the idea that geography constrains the flow of knowledge and that the observed citation behavior is not randomly determined. The differences, reported in the third column of the table, are statistically significant at the 0.01 level. That column also shows that the average difference between observed and simulated distances decreases with the stock of interstate highways. This finding confirms the idea that roads facilitate local circulation flows. As the stock of transport infrastructure increases, local citation patterns appear less constrained by geography and their behavior becomes more similar to a random diffusion process.[9]

B. Heterogeneous Effects and Interpretation

Overall, the results in tables 5 to 8 provide direct evidence that transportation infrastructure affects the flow of local knowledge and facilitates citations among local inventors. In this section, we present further indirect evidence that local roads increase local knowledge flows by reducing the cost of interaction among innovators.

The reduction in travel time generated by transportation infrastructure is likely to be more valuable for technologies where the frontier shifts quickly. In such cases, direct access to the source of new knowledge is especially beneficial. Therefore, we expect roads to have a greater impact on innovation in fields characterized by fast technological turnover. Caballero and Jaffe (1993), Hall et al. (2001), and Mehta, Rysman, and Simcoe (2010) exploit the citation age-profile of patents to measure the speed of technology obsolescence. These studies show that in all industries, old knowledge eventually is made obsolete by the emergence of newer, superior knowledge, but in two industries, technological turnover is much faster than the rest: the communication and computer industry and the electrical and electronics industry. We label these technology classes as high-velocity technologies and contrast them with the other low-velocity technologies in split-sample regressions presented in columns 1 and 2 of table 9. The estimates show that the impact of highways is concentrated on high-velocity

[8] 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 employs five dummy variables for quintile intervals of distance. The coefficient remains very similar to the one reported in column 4 of table 8. Second, we confirm robustness to introducing socioeconomic 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).

[9] We also use the 1,000 simulated counterfactual citations generated through the Monte Carlo simulation to construct 95% confidence bands for the distance between cited and citing local patent. We can reject random behavior for more than 50% of local citations in MSAs in the first tercile for highway stock. The fraction decreases to 40% in the second tercile and to 22% in the third tercile. This confirms the idea that transport infrastructures render local knowledge flow less constrained by distance.

TABLE 9.—ROADS HAVE THE BIGGEST IMPACT ON CITATIONS WITH HIGH-VELOCITY TECHNOLOGIES AND WHEN STARS LIVE IN THE MSA
Dependent Variable: $\log Cites_{m,c,1988}$

Sample	(1) Low-Velocity Classes	(2) High-Velocity Classes	(3) MSAs without Stars	(4) MSAs with Stars
$\log Highway_{m,1983}$	0.178	0.633**	0.219	0.358**
	(0.161)	(0.263)	(0.274)	(0.154)
Class fixed effects	✓	✓	✓	✓
Inventor controls	✓	✓	✓	✓
Geography controls	✓	✓	✓	✓
Population controls	✓	✓	✓	✓
Observations	572	242	307	507
R^2	0.731	0.696	0.422	0.757
F-statistic	12.13	9.13	10.46	10.73

All specifications are estimated by two-stage least squares and control for the lagged dependent variables evaluated in 1983. $\log Cites_{m,c,t}$ refers to the citation-weighted count of patents applied for (and subsequently granted) in period t for MSA m in class c. Using the Hall et al. (2001) NBER technology categories, we classify the Chemicals, Drugs & Medical, Mechanical, and Other categories as low velocity and the Computer & Communications and Electrical & Electronic categories as high velocity. In addition, we identify all inventors above the 90th percentile in the citation-weighted patenting distribution of the focal technology class in 1983. Columns 3 and 4 include all MSA classes that do not have any stars and those that do, respectively. $\log Highway_{m,1983}$ refers to the 1983 level of interstate highway kilometers in the MSA. The endogenous variable $\log Highway_{m,1983}$ is instrumented with $\log HighwayPlan_{m,1947}$, $\log Rail_{m,1898}$, and $\log ExplorationIndex_m$. For MSA regressions, Inventors controls include log of the inventors in the MSA in 1973, 1978, and 1983. For MSA class regressions, Inventor controls include the log of the number of inventors in the MSA class in 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 elevation, index of MSA terrain ruggedness, and number of MSA heating and cooling degree days. Population controls include the log of the MSA's population in 1920, 1930, 1940, 1950, 1960, and 1970. We add a 1 to all patent, citation, and inventor count variables before taking the log to include observations with values of 0. We cluster robust standard errors at the MSA level and present them in parentheses. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

TABLE 10.—ROADS HAVE THE BIGGEST IMPACT ON CITATIONS IN LOW-DENSITY MSAS AND THE ACTIVITY OF SMALL FIRMS
Dependent Variable: $\log Cites_{m,c,1988}$

Sample	(1) Low-Density MSAs	(2) High-Density MSAs	(3) Small Firms	(4) Large Firms
$\log Highway_{m,1983}$	0.411***	0.189	0.367***	0.158
	(0.126)	(0.133)	(0.106)	(0.102)
Class fixed effects	✓	✓	✓	✓
Inventor controls	✓	✓	✓	✓
Geography controls	✓	✓	✓	✓
Socioeconomic controls	✓	✓	✓	✓
Observations	370	444	814	814
R^2	0.599	0.811	0.704	0.719
F-statistic	38.56	16.78	21.82	21.11

All specifications are estimated by two-stage least squares and control for the lagged dependent variables evaluated in 1983. $\log Cites_{m,c,t}$ refers to the citation-weighted count of patents applied for (and subsequently granted) in period t for each MSA class. Columns 1 and 2 include MSAs that are below and above the mean inventor density $\left(\frac{inventors_{m,1983}}{miles^2_m}\right)$, respectively. Column 3 includes only patents produced by small firms (below the 75th percentile of the firm-size distribution; approximately five or fewer inventors). Column 4 includes patents produced by large firms (above the 97th percentile of the firm-size distribution; approximately 54 or more inventors). Both firm-size constructions follow Agrawal et al. (2014). $\log Highway_{m,1983}$ refers to the 1983 level of interstate highway kilometers in the MSA. The endogenous variable $\log Highway_{m,1983}$ is instrumented with $\log HighwayPlan_{m,1947}$, $\log Rail_{m,1898}$, and $\log ExplorationIndex_m$. For MSA regressions, Inventor controls include the log of the number of inventors in the MSA in 1973, 1978, and 1983. For MSA class regressions, Inventor controls include the log of the number of inventors in the MSA class in 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 elevation, index of MSA terrain ruggedness, and number of MSA heating and cooling degree days. Socioeconomic controls (from the 1980 Census) include the share of the poor population in the MSA, the share of college graduates, the share of population employed in manufacturing, mean income in the MSA, and a measure of housing segregation computed by Cutler and Glaeser (1997). 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. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

technologies. Our findings imply that a 10% increase in 1983 interstate highways has no effect on innovation in low-velocity fields but generates a 6.3% increase in citation-weighted patents in the high-velocity fields of computers and electronics.[10]

A well-known feature of science is that the distribution of output is highly skewed across scientists and inventors in the right tail of the output distribution. *Stars* have disproportionately large knowledge spillover effects (Agrawal et al., 2017). This suggests that highways should have a larger impact on innovation in the presence of star inventors. To explore such heterogeneity, we construct a measure of the number of star inventors in each MSA class. We define a star inventor as an inventor above the 90th percentile in the patenting distribution of the technology class in year 1983. On average, an MSA class has roughly seven star inventors, but about 38% of the MSA classes do not have any star inventors. In columns 3 and 4 of table 9, we show that the impact of highways on innovation differs dramatically depending on the presence of star inventors in the MSA class. The regressions indicate a larger effect of transportation infrastructure on patent productivity for inventors located in regions where at least one star is active in the technology class.

We expect the impact of transportation infrastructure to depend on how densely populated a region is. Controlling for the number of inventors in a region, we expect highways to have a larger impact on innovation in regions where inventors are more spread out ("low density") because in such regions, interaction requires traveling a longer distance. In columns 1 and 2 of table 10, we present split-sample regressions distinguishing between technology classes in high- and low-density MSAs. We classify an MSA as "high density" if its ratio of inventors per square mile is above the sample mean. The estimates show that the impact of highways is concentrated in low-density MSAs, suggesting that transportation infrastructure provides greater benefit to knowledge flows when local interaction among innovators requires substantial traveling.[11]

Finally, the impact of transport infrastructure may also differ across firms of different size. In particular, large firms may be less sensitive to highway provision because they generate more inventions internally and are thus more likely to circulate knowledge within their boundaries and less likely to rely on knowledge flows from their neighbors. To assess such heterogeneity, in columns 3 and 4 of table 10, we distinguish between the patenting activity of large and small labs. We construct lab size following Agrawal et al. (2014), who exploit the distribution of lab sizes in each technology class. In a twenty-year sample, they show that across the

[10] This is consistent with the simple model described in section IV in which high velocity can be interpreted as a low value of the parameter γ. When γ is small, past patenting outcome has a low impact on current patenting levels, and reversion to steady-state knowledge production is fast. The simple model suggests that the impact of highways on innovation will be larger for high-velocity technology fields because $\beta = a(1 - \gamma)$.

[11] From a theoretical perspective, the relationship between density and highways is nonlinear. In a simple model of productive interaction with transportation costs, the marginal impact of extra roads is larger in regions that are more densely populated when the stock of roads is low. Nonetheless, the model predicts a larger impact of roads in low-density regions when the stock of roads is large enough.

various class-years, the median size is about 5 inventors, the 75th percentile is about 9 inventors, and the 97th percentile is roughly 54 inventors. We use this distribution to define large and small labs. A large lab is a lab where the number of inventors is above the 97th percentile in the technology class-year distribution. We define a lab as small if the number of inventors is below the 75th percentile. The regressions show a positive highway effect on small firm innovation and no effect on large lab innovation. This suggests that large firms may have an advantage in accessing local knowledge and that the impact of roads on knowledge flows is largely through between-firm rather than within-firm flows.[12]

Overall, these findings indicate that certain regional characteristics are important determinants of the relationship between transportation infrastructure and innovation. The heterogeneity that we examine suggests that an important mechanism driving innovation and growth is the greater circulation of local knowledge caused by the presence of roads. Of course, other factors can also affect the impact of roads on innovation. One is a growth in employment generated by an influx of new workers. A second is a change in the trade pattern across metropolitan areas. Duranton and Turner (2012) and Duranton et al. (2013) provide evidence in support of these two effects of transportation infrastructure.[13]

Our findings are relevant to current policy debates on the role of transportation infrastructure. Our estimates are consistent with the assertion that highways have a meaningful positive effect on economic growth. However, our results also show that the effect can be very different across regions and that the impact of highway provision crucially depends on characteristics of the local environment such as technology specialization, inventor quality, and the manner in which regional R&D manpower is organized.

VII. Employment Growth versus Innovation

Duranton and Turner (2012) show that highways increase regional employment. Perhaps this channel explains a large fraction of the effect we document in our paper. Our goal in this section is to disentangle the two effects of roads (employment growth, also referred to as agglomeration economies, versus patenting, a proxy for innovation) through a simple calibration of a theoretical model. A calibration exercise is a natural approach in our setting because the impacts of roads on innovation and employment growth are jointly determined and to distinguish the two effects empirically is challenging.

[12] We also run a number of split-sample regressions that explore heterogeneity across MSAs of different size. We find very little difference in the impact of roads across cities of different sizes.

[13] We also explore whether the effect of highways is driven by better matches between inventors and firms through job hopping. We measure within-MSA inventor moves following Agrawal et al. (2014) and find evidence that highways cause an increase in job hopping. We also show that our baseline results are robust to controlling for this measure of job hopping. We leave for future research the study of the impact of transport infrastructure on R&D labor reallocation.

Before developing the theoretical model, we run a number of regressions to provide support for the idea that the effect we estimate cannot be entirely explained by the increase in employment documented in Duranton and Turner (2012). Estimates are presented in appendix table B.7. We show that the effect of highways is only marginally reduced once we control for the growth rate of employment, population, and inventors in the 1980s. These estimates should be interpreted with caution because of the obvious endogeneity of these variables. Nevertheless, they suggest that the effect of interstate highways on innovation is unlikely to be entirely driven by an increase in MSA-level employment.

A. *Calibration of an Urban Economy Model*

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:

$$U(w, M) = \ln w + \ln M. \tag{4}$$

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. We assume capital to be perfectly mobile. If we normalize the price of Y to 1 (sold on global market) and denote with r the (nationwide) cost of capital and with w the wage, the model implies the following inverse labor demand curve,

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

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:

$$\ln A = \zeta\ln P + \sigma\ln L. \tag{6}$$

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

Our reduced-form analysis, together with the findings of Duranton and Turner (2012), indicate that both P and L are affected by the provision of roads R. We incorporate 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 depends only 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 equation (7). In the appendix, we extend the model by incorporating additional inputs in the patent production function, equation (8), as well as by relaxing the inelastic labor supply function, equation (7), by 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 a competitive effect (negative term) and a productivity effect (positive term) 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.[14]

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 with respect to patents, ζ, 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, Hornbeck, & Moretti, 2010), which imply elasticities of -0.13 and 0.45, respectively.

[14] In this simple model, the combination of these three effects is also equal to the impact of roads on welfare if road construction is not costly. This equivalence does not hold in the extensions where more structure is imposed on the migration process.

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 11, with this parameterization 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 11, 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, equation (7), microfounding the workers' migration process. Following Duranton and Turner (2012), we assume a pool of people in the rural area receives utility $\overline{U}$ and that cities draw their new workers from this 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 equation (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 nonresidents 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, equation (8), allowing labor, L, to affect patenting. Specifically, we assume

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

Combining this formula with equations (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 overinterpreted. Our model does not consider a number of additional channels through which highways can affect urban growth, such as their impact on the matching process between firms and workers or on the within-city movement of goods and services. Nonetheless, the estimates show that roads may affect productivity through multiple channels and that nonagglomeration forces may explain an important fraction of the productivity gains generated by transportation infrastructure.

Table 11.—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

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 indicates the percentage of the productivity effect explained by greater innovation (as opposed to labor agglomeration).

VIII. Conclusion

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 local knowledge inputs from neighbors located farther away. 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 policymakers. They suggest that the 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., new venture incubators, science parks, technology clusters) may vary across regions and depend on the availability of transportation infrastructure.

REFERENCES

Acemoglu, Daron, and Ufuk Akcigit, "Intellectual Property Rights Policy, Competition and Innovation," *Journal of the European Economic Association* 10 (2012), 1–42.

Aghion, Philippe, and Peter Howitt, "A Model of Growth through Creative Destruction," *Econometrica* 60 (1992), 323–351.

Aghion, Philippe, John Van Reenen, and Luigi Zingales, "Innovation and Institutional Ownership," *American Economic Review* 103 (2013), 277–304.

Agrawal, Ajay, and Iain Cockburn, "The Anchor Tenant Hypothesis: Exploring the Role of Large, Local, R&D-Intensive Firms in Regional Innovation Systems," *International Journal of Industrial Organization* 21 (2003), 1227–1253.

Agrawal, Ajay, Iain Cockburn, Alberto Galasso, and Alexander Oettl, "Why Are Some Regions More Innovative Than Others? The Role of Small Firms in the Presence of Large Labs," *Journal of Urban Economics* 81 (2014), 149–165.

Agrawal, Ajay, Devesh Kapur, and John McHale, "How Do Spatial and Social Proximity Influence Knowledge Flows? Evidence from Patent Data," *Journal of Urban Economics* 64 (2008), 258–269.

Agrawal, Ajay, K., John McHale, and Alex Oettl, "How Stars Matter: Recruiting and Peer Effects in Evolutionary Biology," *Research Policy* 46 (2017), 853–857.

Akcigit, Ufuk, and William R. Kerr, "Growth through Heterogeneous Innovations," NBER technical report (2010).

Baum-Snow, Nathaniel, "Did Highways Cause Suburbanization?" *Quarterly Journal of Economics* 122 (2007), 775–805.

——— "Urban Transport Expansions, Employment Decentralization, and the Spatial Scope of Agglomeration Economies," Brown University working paper (2013).

Belenzon, Sharon, and Mark Schankerman, "Spreading the Word: Geography, Policy and Knowledge Spillovers," this REVIEW 95 (2013), 884–903.

Bloom, Nicholas, and John Van Reenen, "Patents, Real Options and Firm Performance," *Economic Journal* 112 (2002), C97–C116.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen, "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica* 81 (2013), 1347–1393.

Burchfield, Marcy, Henry G. Overman, Diego Puga, and Matthew A. Turner, "Causes of Sprawl: A Portrait from Space," *Quarterly Journal of Economics* 121 (2006), 587–633.

Caballero, Ricardo J., and Adam B. Jaffe, "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 (1993), 15–74.

Carlino, Gerard, and William Kerr, "Agglomeration and Innovation," NBER working paper 20367 (2014).

Catalini, Christian, "Microgeography and the Direction of Inventive Activity," MIT Sloan working paper (2013).

Chandra, Amitabh, and Eric Thompson, "Does Public Infrastructure Affect Economic activity? Evidence from the Rural Interstate Highway System," *Regional Science and Urban Economics* 30 (2000), 457–490.

Cutler, David M., and Edward L. Glaeser, "Are Ghettos Good or Bad?" *Quarterly Journal of Economics* 112 (1997), 827–872.

Duranton, Gilles, Peter Morrow, and Matthew A. Turner, "Roads and Trade: Evidence from the US," *Review of Economic Studies* 81 (2014), 681–724.

Duranton, Gilles, and Henry G. Overman, "Testing for Localization Using Microgeographic Data," *Review of Economic Studies* 72 (2005), 1077–1106.

Duranton, Gilles, and Matthew A. Turner, "Urban Growth and Transportation," *Review of Economic Studies* 79 (2012), 1407–1440.

Feldman, Maryann P., *The Geography of Innovation* (Dordrecht: Kluwer, 1994).

Feldman, Maryann P., and David B. Audretsch, "Innovation in Cities: Science-Based Diversity, Specialization and Localized Competition," *European Economic Review* 43 (1999), 409–429.

Fernald, John G., "Roads to Prosperity? Assessing the Link between Public Capital and Productivity," *American Economic Review* 89 (1999), 619–638.

Forman, Chris, Avi Goldfarb, and Shane Greenstein, "Digital Dispersion: An Industrial and Geographic Census of Commerical Internet Use," NBER technical report (2002).

Furman, Jeffrey L., Michael E. Porter, and Scott Stern, "The Determinants of National Innovative Capacity," *Research Policy* 31 (2002), 899–933.

Galasso, Alberto, Mark Schankerman, and Carlos Serrano, "Trading and Enforcing Patent Rights," *RAND Journal of Economics* 44 (2013), 275–312.

Ghani, Ejaz, Arti Grover Goswami, and William R. Kerr, "Highway to Success: The Impact of the Golden Quadrilateral Project for the Location and Performance of Indian Manufacturing," *Economic Journal* 126 (2014), 317–357.

Glaeser, Edward Ludwig, Sari Pekkala Kerr, and William Kerr, "Entrepreneurship and Urban Growth: An Empirical Assessment

with Historical Mines," Harvard Business School working paper 13-015 (2012).
Gray, C., *Gray's New Trunk Railway Map of the United States* (Washington, DC: Library of Congress Catalogue, 1989).
Greenstone, Michael, Richard Hornbeck, and Enrico Moretti, "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings," *Journal of Political Economy* 118 (2010), 536–598.
Hall, Bronwyn H., Adam B. Jaffe, and Manuel Trajtenberg, "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools," NBER working paper 8498 (2001).
Henderson, Vernon, "The Urbanization Process and Economic Growth: The So-What Question," *Journal of Economic Growth* 8 (2003), 47–71.
Jacobs, J., *The Economy of Cities* (New York: Random House, 1969).
Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson, "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations," *Quarterly Journal of Economics* 108 (1993), 577–598.
Kerr, William, and Scott Kominers, "Agglomerative Forces and Cluster Shapes," this REVIEW 97 (2015), 877–899.
Kline, Patrick, and Enrico Moretti, "Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority," *Quarterly Journal of Economics* 129 (2014), 275–331.
Kortum, Samuel, and Josh Lerner, "Assessing the Contribution of Venture Capital to Innovation," *RAND Journal of Economics* 31 (2000), 674–692.
Marx, Matt, Deborah Strumsky, and Lee Fleming, "Mobility, Skills, and the Michigan Non-Compete Experiment," *Management Science* 55 (2009), 875–889.
Mehta, Aditi, Marc Rysman, and Tim Simcoe, "Identifying the Age Profile of Patent Citations: New Estimates of Knowledge Diffusion," *Journal of Applied Econometrics* 25 (2010), 1179–1204.
Michaels, G., "The Effect of Trade on the Demand for Skill: Evidence from the Interstate Highway System," this REVIEW 90 (2008), 683–701.
Moretti, Enrico, "Local Labor Markets" (pp. 1237–1313), in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, vol. 4 (Amsterdam: Elsevier, 2011).
Moretti, Enrico, and Daniel J. Wilson, "State Incentives for Innovation, Star Scientists and Jobs: Evidence from Biotech," *Journal of Urban Economics* 79 (2014), 20–38.
Pakes, Ariel, and Zvi Griliches, "Patents and R&D at the Firm Level: A First Report," *Economics Letters* 5 (1980), 377–381.
Redding, Steve, and Matthew Turner, "Transportation Costs and the Spatial Organization of Economic Activity," in Gilles Duranton, J. Vernon Henderson, and William C. Strange, eds., *Handbook of Regional and Urban Economics*, vol. 5 (Dordrecht: North-Holland, 2015).
Romer, Paul M., "Increasing Returns and Long-Run Growth," *Journal of Political Economy* 94 (1986), 1002–1037.
Rosenthal, Stuart S., and William C. Strange, "Evidence on the Nature and Sources of Agglomeration Economies" (pp. 2119–2171), in J. Vernon Henderson and Jacques-François Thisse, eds., *Handbook of Regional and Urban Economics*, vol. 4 (Amsterdam: Elsevier, 2004).
Singh, Jasjit, and Matt Marx, "Geographic Constraints on Knowledge Spillovers: Political Borders vs. Spatial Proximity," *Management Science* 59 (2013), 2056–2078.
U.S. Bureau of Transportation Statistics, "Transportation Statistics Annual Report" (2010).
U.S. Geological Survey, *The National Atlas of the United States of America* (Washington, DC: U.S. Government Printing Office, 1970).