

DISCUSSION PAPER SERIES

DP16794

Cities and Technological Waves

Martí Mestieri, Enrico Berkes and Ruben Gaetani

INTERNATIONAL TRADE AND REGIONAL ECONOMICS

MACROECONOMICS AND GROWTH

CEPR

Cities and Technological Waves

Martí Mestieri, Enrico Berkes and Ruben Gaetani

Discussion Paper DP16794
Published 08 December 2021
Submitted 07 December 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- International Trade and Regional Economics
- Macroeconomics and Growth

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Martí Mestieri, Enrico Berkes and Ruben Gaetani

Cities and Technological Waves

Abstract

We develop a spatial model of economic growth to study the effect of changes in the technological landscape on the spatial distribution of economic activity. We use this framework to study the evolution of the U.S. economic geography over the twentieth century. In the model, innovation via frictional idea diffusion makes cities trajectories sensitive to “technological waves,” defined as long-term shifts in the importance of different knowledge fields. We calibrate the model using a new dataset of historical geolocated patents, and find that cities differential exposure to technological waves explains between 15% and 20% of the variation in local population growth over the twentieth century. Counterfactual experiments suggest large and heterogeneous geographical effects of future technological scenarios.

JEL Classification: R12, O10, O30, O33, O47

Keywords: cities, population growth, technology diffusion, Innovation, patents

Martí Mestieri - marti.mestieri@northwestern.edu
Northwestern University, FRB Chicago and CEPR

Enrico Berkes - berkes.8@osu.edu
Ohio State University

Ruben Gaetani - ruben.gaetani@rotman.utoronto.ca
University of Toronto

Cities and Technological Waves*

Enrico Berkes[†]

Ruben Gaetani[‡]

Martí Mestieri[§]

7th December 2021

Abstract

We develop a spatial model of economic growth to study the effect of changes in the technological landscape on the spatial distribution of economic activity. We use this framework to study the evolution of the U.S. economic geography over the twentieth century. In the model, innovation via frictional idea diffusion makes cities trajectories sensitive to “technological waves,” defined as long-term shifts in the importance of different knowledge fields. We calibrate the model using a new dataset of historical geolocated patents, and find that cities differential exposure to technological waves explains between 15% and 20% of the variation in local population growth over the twentieth century. Counterfactual experiments suggest large and heterogeneous geographical effects of future technological scenarios.

Keywords: Cities, Population Growth, Technology Diffusion, Innovation, Patents.

JEL Classification: R12, O10, O30, O33, O47.

*We thank Mike Andrews, Paco Buera, Ben Jones, Joel Mokyr, Nicola Persico, Frédéric Robert-Nicoud, Bruce Weinberg, our discussants Klaus Desmet, Ed Glaeser, Xian Jiang, and Sara Mitchell, and seminar attendees at University of Colorado Boulder, University of Toronto, Ohio State University, University of Pittsburgh, Conference of Swiss Economists Abroad, NBER Summer Institute (Urban Economics), VMACS Junior Conference, 2021 AEA Meetings, 2021 European Meeting of the Urban Economics Association, 2021 Canadian Summer Conference in Real Estate and Urban Economics, 2021 Barcelona GSE Summer Forum (Firms in the Global Economy), 2021 SED Meeting, and 2021 North American Meeting of the Urban Economics Association for helpful comments and discussions. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System. All errors are our own.

[†]The Ohio State University, berkes.8@osu.edu.

[‡]University of Toronto, ruben.gaetani@utoronto.ca.

[§]Federal Reserve Bank of Chicago and CEPR, mestieri.marti@gmail.com.

1 Introduction

The economic geography of countries is perpetually evolving. In the United States, many cities and regions that have thrived in the past have progressively lost population and influence to other geographical areas. In recent decades, several cities in the Rust Belt, which had experienced extraordinary growth throughout most of the twentieth century, have entered a prolonged phase of decline. At the same time, a handful of urban areas specialized in knowledge-intensive sectors, such as information technology (IT) and pharmaceuticals, have gained prominence, becoming increasingly attractive for workers and firms (Glaeser and Gottlieb, 2009; Moretti, 2012). The determinants of these rich dynamics in city growth are still a matter of debate and remain a central question in urban economics.

In this paper, we propose that frictions to knowledge diffusion across locations and fields of knowledge make the growth trajectory of cities sensitive to “technological waves,” defined as long-term shifts in the importance of fields of knowledge in the innovation landscape. Leveraging a new dataset of geolocated U.S. patents, we document a robust positive reduced-form relationship between a city’s exposure to technological waves and its ability to attract population over the following decades since the early twentieth century. Motivated by this finding, we then develop a quantitative model that formalizes the feedback between technological waves and the dynamics of city growth. The model combines an economic geography setting with a theory of economic growth that emphasizes the role of recombination, imitation, and diffusion of knowledge, as recently developed by Lucas and Moll (2014), Perla and Tonetti (2014), and Buera and Oberfield (2020), among others.

The quantitative results suggest that frictional knowledge diffusion accounts for most of the reduced-form relationship between exposure to technological waves and local population growth that we document. Overall, the differential exposure of cities to technological waves explains between 15% and 20% of the total variation in local population growth during the twentieth century. Barriers to diffusion across fields of knowledge and geographical areas are both important, each explaining roughly half of this relationship. We also use the model to investigate how alternative scenarios of future technological waves might transform the United States’ economic geography in the coming decades. The model predicts substantial differences in the geographical effects of alternative scenarios such as a rise of autonomous vehicles, medical sciences, or sustainable agriculture.

To begin our analysis, we leverage in Section 2 a new comprehensive dataset of geolocated historical patents to provide reduced-form evidence that over the last century, cities’ growth trajectories were systematically affected by changes in the technological landscape. Specifically, using a shift-share measure of exposure to technological waves (Bartik, 1991), we document that cities whose innovation activities were concentrated in expanding fields experienced systematically higher population growth compared with cities whose innovation activities were concentrated in declining fields. This relationship is stable over time, and it is robust to controlling for proxies of the local density of human capital and the local industry composition. Using the same patents data, we also document that the patterns of knowledge diffusion, as measured via patent citations, are persistently localized, both in the geographical and technological spaces. Taken together, these facts suggest that a city’s ability to seize new technological opportunities depends on the local availability of complementary ideas, and can explain why changes in the technological environment lead to the rise of some cities and the decline of others. In the remainder of the paper, we formalize this mechanism by developing a spatial model of endogenous growth with innovation and frictional idea diffusion, and then use it to quantify our proposed mechanism.

In the model, that we introduce in Section 3, newborn agents make migration and occupational decisions after forming expectations on their lifetime productivity in each location and sector. Productivity is determined by a decision to imitate or innovate. Agents can either imitate an idea drawn from the local knowledge distribution, or innovate by improving upon an idea drawn from the distribution of any other location and sector in the economy. The applicability of an idea is affected by frictions reflecting both geographical and technological distance. These frictions imply that knowledge drawn *within* any location-sector can be converted into new inventions more effectively than knowledge drawn from other locations and sectors. For this reason, a city’s stock of knowledge determines both current productivity and future innovation possibilities, making the local growth trajectory sensitive to changes in the technological centrality of different sectors—what we refer to as “technological waves.” To focus on this novel interplay between economic geography and idea diffusion, the model purposefully abstracts from other drivers of city growth such as endogenous residential amenities.

The framework remains tractable for any arbitrary number of locations, sectors, and time periods, and it has a unique equilibrium with an explicit solution. Absent technological wave shocks, the model features a unique balanced growth path (BGP). The distribution of ideas for each location-sector endogenously retains a Fréchet structure, and it implies an intuitive law of motion of

its scale parameter. This law of motion summarizes the evolution of the distribution of knowledge for each location-sector and its dependence on all other location-sectors. The Fréchet structure also allows us to characterize knowledge flows in closed form through a gravity representation that can be estimated using patent citation data.

Despite its relative parsimony, our model generates linkages across geographical areas and fields of knowledge that imply non-trivial population dynamics. Before turning to the quantitative analysis, we study the mechanics of the model by log-linearizing the equilibrium conditions around the BGP. We derive intuitive theoretical predictions about the relationship between technological waves, the evolution of local productivity, and city population growth. First, the growth rate of productivity in each location-sector can be expressed as the sum of aggregate sectoral shocks weighted by the reliance of local innovation on ideas from each perturbed sector. This implies that cities specialized in expanding (declining) sectors will experience higher (lower) local productivity growth. Second, combining these productivity dynamics with individual migration decisions, we show that a measure of local exposure to technological waves relative to the overall economy is a sufficient statistic to predict local population growth. Third, in the particular case of knowledge flows across sectors being of second-order importance relative to flows within sectors, this measure of exposure becomes the standard shift-share variable that we use in our reduced-form analysis—implying that a city grows if and only if the average of sectoral shocks weighted by the incidence of each sector in the city is larger than the corresponding weighted average for the rest of the economy.

In Section 4, we turn to the quantitative assessment of our proposed mechanism in explaining the evolution of the U.S. economic geography over the twentieth century. We show that the model has a recursive structure that allows us to calibrate the parameters and to recover the unobserved disturbances—including the technological wave shocks—by imposing a small set of transparent assumptions. To validate our quantification exercise, we show that the calibrated model is successful in capturing key moments of the data that are not directly targeted, including the relationship between city size and city average income.

Our quantitative results, which we present in Section 5, suggest that technological waves can account for a significant portion of the variation in population growth across cities over the last century. In our baseline counterfactual, which isolates the effect of technological waves on population growth through the endogenous mechanism of knowledge creation and diffusion, increasing the local exposure to technological waves by one standard deviation increases local population

growth by 15.8% of a standard deviation in the first half of the twentieth century and by 20.0% of a standard deviation in the second half. These estimates imply that the endogenous mechanism of knowledge creation and diffusion can account for most of the reduced-form relationship that we document, in which exposure to technological waves explains 22.5% and 20.7% of the variation in population growth in the first and second half of the twentieth century, respectively. A decomposition exercise reveals that frictions to knowledge diffusion across geographical areas and technological fields roughly equally contribute to explaining this effect.

The mechanism of frictional knowledge creation and diffusion that we study implies that the degree of local diversification plays a central role in determining a city’s resilience in the face of technological wave shocks. Simulations of counterfactual paths of sectoral shocks reveal that more diversified cities experience significantly less volatile growth trajectories. There are two factors behind this relationship that reflect the existence of frictions in the knowledge and geographical dimensions. First, frictions to diffusion across fields of knowledge imply that in response to technological wave shocks, productivity growth is higher in some sectors compared with others. As a result, more diversified cities have a lower chance of large swings (either on the positive or negative side) in their productivity growth, since negative shocks to some sectors are likely to be compensated by positive shocks to other sectors. Second, frictions to diffusion across geographical areas imply that more diversified cities have a broader availability of ideas to draw from, so that (positive or negative) shocks to individual sectors have a weaker impact on the evolution of local productivity.

Finally, we use the quantitative model to explore the predicted city dynamics in the coming decades under different scenarios for the evolution of the technological landscape. In particular, we study which cities benefit in terms of population growth—and which do not—compared with the status quo in the following scenarios: (1) a rise in the importance of transportation-related technologies, due to the emergence of new modes of transportation such as autonomous vehicles; (2) an increase in the centrality of pharmaceuticals and biotech in response to new challenges in global health; (3) a comeback of agriculture as a pivotal sector in the innovation landscape as a result of regulatory changes and increasing demand for sustainable farming. We find that cities in the Rust Belt benefit from the first scenario and experience positive population growth, at the expense of cities in the North-East and the Pacific regions. The second scenario penalizes knowledge hubs specialized in IT-related innovation, favoring more diversified areas such as Boston and the cities in California outside the Silicon Valley. The third scenario prompts a reallocation

of economic activity towards the agricultural areas in the Central states.

Related Literature This paper builds on multiple strands of the literature. First, our theory is based on modelling idea flows across location-sectors, with technological and geographical frictions in knowledge diffusion playing a key role in explaining city dynamics. While a rich body of literature has documented the strength and geographical span of localized knowledge spillovers (among others, [Jaffe et al., 1993](#); [Audretsch and Feldman, 1996](#); [Greenstone et al., 2010](#)), there has been no attempt, to the best of our knowledge, to perform a quantitative assessment of the importance of localized knowledge for understanding long-run city dynamics.

One of the main obstacles for providing such an assessment is the complexity of modeling idea diffusion in a spatial setting. In recent years, two flourishing bodies of literature have provided major methodological advances to help address this complexity. First, a number of papers have developed tractable endogenous growth models that emphasize recombination, imitation, and knowledge diffusion as major drivers of aggregate productivity growth (e.g., [Lucas and Moll, 2014](#); [Perla and Tonetti, 2014](#); [Buera and Oberfield, 2020](#); [Huang and Zenou, 2020](#)). Second, a rich body of work on quantitative spatial economics has developed tools for studying the distribution of economic activity in space, both within cities (e.g., [Ahlfeldt et al., 2015](#); [Heblich et al., 2020](#)) and in a system of locations (e.g., [Allen and Arkolakis, 2014](#); [Desmet et al., 2018b](#)).¹ This paper combines insights from these two strands of the literature and develops a dynamic endogenous growth model in a spatial economy that is highly tractable and can be quantitatively disciplined using data on city population and patents over a long time period. While a number of papers have used detailed data on patenting to study innovation and knowledge flows in firm and industry dynamics (e.g., [Kogan et al., 2017](#), [Akcigit and Kerr, 2018](#); [Cai and Li, 2019](#)) or developed static models that emphasize localized knowledge spillovers as the main determinant of the economic geography (e.g., [Davis and Dingel, 2019](#)), this paper is, to the best of our knowledge, the first attempt at quantitatively assessing the importance of frictions to knowledge diffusion for city dynamics.

Our paper also relates to the vast and rich literature that has investigated the forces governing the long-run evolution of the economic geography, specifically in its propensity to display path dependence and occasional reversals of fortune (e.g., [Brezis and Krugman, 1997](#); [Davis and Weinstein, 2002](#); [Bleakley and Lin, 2012](#); [Kline and Moretti, 2014](#)), as well as in its responsiveness

¹[Buera and Lucas \(2018\)](#) provide a comprehensive review of the body of literature on models of endogenous growth with idea flows, and [Redding and Rossi-Hansberg \(2017\)](#) provide a comprehensive review of the body of literature on quantitative spatial equilibrium models.

to aggregate shocks such as rising sea level (e.g., [Desmet et al., 2018a](#)), and regional or sectoral shocks (e.g., [Caliendo et al., 2018](#); [Hornbeck and Moretti, 2018](#); [Adao et al., 2020](#)). The working hypothesis in this paper is that aggregate changes in the technological landscape, combined with frictional knowledge transmission, have a first-order impact on the geographical distribution of economic activity. The framework can account simultaneously for path dependence and reversal of fortune in city dynamics. While the focus on innovation and frictional idea diffusion is new to this literature, there is a rich body of work that has analyzed the historical dynamics of the U.S. geography, both from an empirical perspective (e.g., [Bostic et al., 1997](#); [Simon and Nardinelli, 2002](#); [Michaels et al., 2012](#); [Desmet and Rappaport, 2017](#)) and from a structural and quantitative viewpoint (e.g., [Duranton, 2007](#); [Desmet and Rossi-Hansberg, 2014](#); [Nagy, 2017](#); [Allen and Donaldson, 2018](#); [Eckert and Peters, 2019](#)).

This paper also contributes to the longstanding debate on the returns to local specialization ([Marshall, 1890](#)) versus urban diversity ([Jacobs, 1969](#)), as well as their effects on city growth. Notable contributions in this literature include [Glaeser et al. \(1992\)](#), whose empirical assessment finds evidence supporting Jane Jacob’s view of urban variety as a key driver of local employment growth, and [Duranton and Puga \(2001\)](#), who develop a model in which diversified and specialized cities coexist in equilibrium.² This paper suggests and quantifies a new channel through which urban diversification affects long-run city growth, namely, by shaping the responsiveness of a city to changes in the surrounding technological landscape.³ In this sense, the model provides a new lens for interpreting the effect of local policies directed at increasing local diversification.

The remainder of the paper is organized as follows. Section 2 introduces the data and presents historical trends and the motivational facts on the relationship between city growth and the technological landscape. Section 3 introduces the model and derives the main theoretical predictions. Section 4 describes the model calibration and Section 5 presents the quantitative results. Section 6 discusses avenues for further research and concludes.

²A comprehensive overview of the patterns of specialization across U.S. locations is provided by [Holmes and Stevens \(2004\)](#).

³Consistently with this interpretation, [Balland et al. \(2015\)](#) find that cities with more diverse knowledge bases are less sensitive to technological crises, defined as sustained declines in patenting activity.

2 Data and stylized facts

Technological change is a slow-moving secular process. To study how the rise and fall of technologies determines the growth and decline of cities, we therefore need to consider a time period long enough to capture multiple episodes of widespread technological transformation. In this paper, we exploit a recently assembled dataset of historical geolocated patents from 1836 through 2015. We define cities as 1990 U.S. commuting zones (CZs), and we keep them fixed throughout the analysis.⁴

2.1 Data sources

To measure innovative activities at the city level, we collect patent data from the Comprehensive Universe of U.S. Patents, or CUSP (Berkes, 2018). The CUSP contains information on the near-universe of patents issued by the U.S. Patent and Trademark Office (USPTO) between 1836 and 2015, with an estimated coverage above 90% in each year. From the CUSP, we gather information on the technology classes and location of the first inventor listed on each patent, as well as the filing date. The CUSP assigns patents to the city of the inventors' residence and does not rely on the county reported in the patent's text. This allows us to build geographically consistent measures of innovation at the commuting zone level over the long time span covered by our study. This paper is the first to exploit the entire time series of geolocated patents provided by the CUSP dataset.⁵

We also collect data on population, human capital, and industry composition at the commuting zone level using the corresponding decennial censuses for each decade between 1870 and 2010⁶ from the Integrated Public Use Microdata Series (IPUMS, Ruggles et al., 2021) and the National Historical Geographic Information System (NHGIS, Manson et al., 2021).⁷ We build a consistent measure of the local density of human capital that combines available information on local literacy and education. To make this measure comparable across decades, we rank cities in terms of the

⁴Although commuting flows have changed over time, assuming a stable geography allows us to abstract from annexations and redefinition of town borders that have been pervasive phenomena throughout the nineteenth and twentieth century.

⁵Berkes (2018) provides details about the data collection procedure, as well as summary statistics and stylized facts related to the underlying data. Andrews (2021), in a comparison of historical patents data, describes it as “currently the gold standard both in terms of completeness and scope of the types of patent information it contains.” Some slices of the data have been used in Berkes and Nencka (2020), who study the effect of Carnegie libraries on the local patenting activity; Clemens and Rogers (2020), who study how procurement policies affect the characteristics of medical innovation; and Babina et al. (2020), who study the effect of the Great Depression on innovative activities in the United States.

⁶For the 2010 observation, we use multi-year averages of the American Community Survey (ACS).

⁷Since data from the 1890 decennial census are not available, we construct the 1890 observations by linearly interpolating the observations from the 1880 and 1900 decennial censuses.

relevant measure for each decade and use the resulting ranking for the analysis.⁸ For each decade, we also collect information on local employment by industry. Appendix D provides further details on the construction of the data.

Since our focus is on measuring long-run technological change, our unit of time throughout the empirical and structural analysis corresponds to 20-year periods, spanning from 1870 through 2010. Patent counts by sector are obtained by adding patents filed in the two decades around the focal year (for example, patents in the 1990 observation correspond to the total patent count between 1980 and 1999). We restrict our sample to the subset of commuting zones in the contiguous United States that accounted for at least 0.02% of the total population for each decade since 1890. This delivers a sample of 373 commuting zones, which jointly account for roughly 87% of the U.S. population in 2010.⁹ Sectors are defined as the technological class-groups obtained by grouping 3-digit International Patent Classification (IPC) categories into 11 class-groups, as detailed in Appendix Table A.1. Among these 11 categories, we have agriculture, health and life-saving inventions, transportation, chemistry, and electricity.¹⁰

2.2 Historical trends

The last 150 years have witnessed major shifts in the technological landscape as measured by changes in the patenting shares across different patent classes. These major shifts are already apparent when comparing patenting output across the main broad IPC classes (which are coarser than our baseline 11 categories). The bottom-right panel of Figure 1 shows how the distribution of the national patenting output has evolved since 1870.¹¹ The share of patents in “Human Necessities”—that includes innovation related to both agriculture and medical sciences—declined in the first part of the century, as agriculture lost its centrality to classes complementary to the


⁸The relevant measure is a summary index that includes several indicators of the local density of human capital. The specific indicators we use change over time depending on the availability of information in the historical census. In the early decades, the measure focuses on indicators of literacy and schooling, and in the later decades, it emphasizes the local density of workers with high educational attainment.

⁹As a reference point, this rule requires that cities had a population of at least 10,711 people in 1890 and 60,387 people in 2010.

¹⁰Patents listing multiple 3-digit IPC classes are assigned fractionally to class-groups, in proportion to the frequency of appearance of each class-group in the list of 3-digit IPC classes.

¹¹For sake of clarity, we are reporting the seven main IPC classes (that correspond to the first letter in the IPC). This has the drawback of bundling together, among others, innovations related to agriculture and medicine. Appendix Figure B.1 shows the corresponding distribution across the 11 IPC class-groups described in Appendix Table A.1 that we use in our analysis, which separates, among others, agriculture and medicine. Class names are abbreviated for clarity. The full description of each class can be found at <https://www.wipo.int/classifications/ipc/en/>. Kelly et al. (2021) provide an alternative measure of technological importance by constructing technology indices based on textual analysis of patent data.

Figure 1: Composition of the technological output



Notes: Composition of patenting output across the seven main IPC classes. Patent count for year t is constructed as the sum of patents filed between $t - 10$ and $t + 9$. Class names are abbreviated. The full description for each class is available at <https://www.wipo.int/classifications/ipc/en/>.

heavy manufacturing industry, “Transporting” and “Mechanical Engineering.” Patents in “Human Necessities” rebounded in recent decades as innovation in medicine gained prominence. In the second part of the century, patents in “Physics” and “Electricity” classes became more central in the national shares, making up more than 50% of the overall innovation output in 2010.¹²

The composition of patenting not only has changed significantly over time, but also varies considerably across cities at any point in time. The top panels of Figure 1 depict two archetypal examples of this heterogeneity. Detroit (top-left) has been specialized in the production of patents related to “Transporting” and “Mechanical Engineering” since the early 1900s. In 1930, these two classes made up about 70% of its patenting portfolio. This pattern has remained broadly unchanged throughout the century, with a modest shift towards patents of classes “Physics” and “Electricity” since the 1990s. Austin (top-right) exhibits fairly diversified innovation activities until the 1970s, when the share of patents in classes “Physics” and “Electricity” started expanding, reaching 90% of the portfolio by 2010. By contrast, Boston (bottom-left) displays a diversified patenting output that, throughout the decades, has closely tracked the national trends.

¹²Classes “Physics” and “Electricity” include the bulk of innovation related to computers, electronics, and information and communication technology.

In this paper, we argue that the heterogeneity in the composition of local knowledge as proxied by city patenting makes cities unevenly positioned to take advantage of new innovation opportunities. The key hypothesis underlying our argument is that knowledge, to a large extent, is localized and diffuses slowly. This makes cities' trajectories sensitive to changes in the technological landscape because their current knowledge portfolio determines to which extent they can seize new technological opportunities. As a result, cities experience heterogeneous productivity improvements from common technological shocks. Ultimately, this contributes to explaining the heterogeneous historical dynamics of U.S. urban and regional growth.

The population growth experiences of Detroit, Austin, and Boston since the late 1800s exemplify this point. Figure 2 shows the 20-year population growth of these three commuting zones since 1890, residualized with respect to Census Division-time fixed effects, which control for systematic regional differences in population growth over time. Detroit displays the most striking growth rates in the decades after the rise of the automobile industry around 1910, followed by a long-lasting decline that resulted in a steady loss of population since the 1980s. The commuting zone of Austin experienced a specular trajectory. The city declined in relative terms in the first half of the twentieth century, as the Texas Oil Boom favored areas of the state that were rich in oil, making Austin slip from the 4th to the 10th place among Texas's largest cities.¹³ However, in recent decades Austin has emerged as one of the leading innovation hubs in the country, leveraging its large number of science-based firms and sizable college-educated population. Finally, the commuting zone of Boston presents yet a different experience. Throughout the last century, it has retained a considerably less volatile path, characterized by moderate relative growth interrupted by occasional periods of modest relative decline. The persistent diversification of Boston's patenting output could have made the city less sensitive to changes in the technological landscape, explaining the stability of its growth path.¹⁴


2.3 Technological waves and the growth and decline of cities

Taken together, Figures 1 and 2 suggest that changes in the importance of technological fields might differentially affect the growth trajectory of cities because of their pre-existing specialization across different technology fields. We now show that this pattern indeed holds systematically over

¹³<https://tshaonline.org/handbook/online/articles/hda03>

¹⁴Glaeser (2005) provides an overview of the causes of the slow decline of Boston between 1920 and 1980, and the subsequent re-emergence of the city. The high density of human capital is proposed as the major factor behind its resilience.

Figure 2: City dynamics



Notes: Residuals of a regression of 20-year growth rate of population on Census Division-time fixed effects, 1890-2010.

the long time period covered by our data. Moreover, we also show that this pattern is robust to controlling for possible confounders such as the local density of human capital and the local industrial composition.


As noted previously, throughout this paper, we refer to changes in the technological landscape, here captured by shifts in the composition of national patenting by class-group, as *technological waves*. As we have discussed, we explore the hypothesis that on account of the localized nature of knowledge transmission, cities whose patenting portfolio are concentrated in expanding fields are in a better position to take advantage of new innovation possibilities. As a result, these cities will experience higher productivity and population growth. To capture this idea in a simple setting, in the spirit of Bartik (1991), we construct a shift-share measure of local exposure to technological waves:

$$Exp_{n,t} \equiv \sum_{s \in S} Share_{n,s,t-1} \times g_{s,t}, \quad (1)$$

where $Share_{n,s,t-1}$ is the share of patents filed in commuting zone n belonging to class-group s at time $t - 1$ and $g_{s,t}$ is the growth rate in the national share of patents of class-group s between $t - 1$ and t . A city whose portfolio of patents is concentrated in expanding (declining) class-groups will record a positive (negative) value of $Exp_{n,t}$, reflecting a favorable (adverse) exposure to the current technological wave.¹⁵

¹⁵To gain intuition, cities with a high share of patenting in expanding fields will display a positive and large

Figure 3: Technological waves and city growth



Notes: Bin-scatter plot of exposure to the technological wave, as defined in Equation (1), and 20-year population growth, 1910-2010. The bin-scatter plot is residualized with respect to Census Division-time fixed effects and two lags of log-population.

Figure 3 shows a bin-scatter plot of the relationship between the measure of exposure, $Exp_{n,t}$, and the 20-year growth rate of local population, between 1910 and 2010. Both measures are residualized with respect to two lags of log-population and Census Division-time fixed effects in order to account for size, convergence, and persistence effects, and for the differential growth rates of commuting zones across space explained by factors such as the westward expansion or the Great Northward Migration.¹⁶ The scatter plot reveals a strong positive correlation, implying that, over the period considered, cities with a more favorable exposure to the technological wave have experienced systematically higher population growth compared with other cities in the same Census Division.

Table 1 reports the corresponding regression results. The estimates in column 2 (which controls for Census Division-time fixed effects) imply that an increase in the measure of exposure of one residual standard deviation is associated with an increase of 12.5% of a residual standard deviation in population growth. In column 3, we further control for the historically consistent measure of human capital. As documented by Glaeser and Saiz (2003), this indicator is correlated with population growth. Yet, it has a negligible effect on the estimated coefficient of the exposure

exposure measure. By contrast, cities with a high share of patenting in declining fields will display a large in absolute value but negative exposure measure.

¹⁶The earliest period corresponds to population growth between 1890 and 1910, and controls for two lags of log-population (1870 and 1890). The latest period corresponds to population growth between 1990 and 2010, and controls for two lags of log-population (1970 and 1990).

measure, suggesting that the measure of exposure to technological waves is not simply capturing the availability of human capital in the city.

In our analysis, we do not take a stance on what drives changes in the national patenting shares by class-group. Technological waves could result from genuine scientific and technological developments, such as advances in computing and bio-technology. Alternatively, they could be triggered by political and environmental factors, such as regulation, trade agreements, or changes in consumer preferences. What is critical for our analysis is that, whatever their origin, technological waves *differentially affect the returns to innovation* in different fields and, as a result, they affect the evolution of patenting shares across class-groups. In other words, our view is one of profit-driven innovation, as emphasized in [Schmookler \(1966\)](#); the intensity of innovation across fields responds to changes in their returns, whether that be from changes in its costs (e.g., due to scientific breakthroughs) or market size (e.g., due to changes in demand).

An adversarial view to what we propose would be one in which innovation occurs as a byproduct of production. More broadly, one could state that factors that differentially impact patenting across fields might be correlated with other industry-level shocks that drive differences in population growth across cities, confounding our interpretation of the estimates. To address this concern, we directly control for the sectoral composition at the city level.¹⁷ In column 4 of [Table 1](#), we directly control for a shift-share variable built using employment by industry.¹⁸ This Bartik variable is a strong predictor of contemporaneous population growth, and its inclusion slightly reduces the estimated coefficient on the exposure measure, that nevertheless remains large and significant at the 1% level. The estimate in this specification implies that an increase in the measure of exposure of one residual standard deviation is associated with an increase of 9.1% of a residual standard deviation in population growth. Finally, in [Appendix Table A.3](#), we also show that results are consistent when splitting the sample into early (1910-1950) and late (1970-2010) sub-samples, confirming that this correlation is a stable regularity throughout history.

¹⁷Note that, if our previous interpretation is correct, introducing this control would purge part of potentially valid variation.

¹⁸This shift-share variable is analogous to the one in [Equation \(1\)](#) but uses employment by industry in place of patenting by class-group. See [Appendix D](#) for details on the construction of the data on employment by industry at the commuting zone level. Industries correspond to the 12 main industries in the 1950 Census Bureau industrial classification system.

Table 1: **Technological waves and city growth**

| | Growth rate of population | | | |
|-------------------------|---------------------------|----------------------|----------------------|----------------------|
| | (1) | (2) | (3) | (4) |
| Exposure to tech. wave | 0.428*** (0.082) | 0.400*** (0.067) | 0.370*** (0.071) | 0.279*** (0.071) |
| Log-population (lag 1) | 0.278*** (0.056) | 0.264*** (0.045) | 0.255*** (0.045) | 0.203*** (0.048) |
| Log-population (lag 2) | -0.301*** (0.053) | -0.272*** (0.038) | -0.269*** (0.038) | -0.242*** (0.039) |
| Human capital (ranking) | | | 0.082* (0.047) | 0.034 (0.047) |
| Industry composition | | | | 0.660*** (0.116) |
| Fixed effects | T | CD×T | CD×T | CD×T |
| # Obs. | 2,238 | 2,238 | 2,238 | 2,228 |
| R^2 | 0.39 | 0.50 | 0.51 | 0.53 |

Notes: Commuting zone (CZ)-level regression, 1910-2010. Dependent variable defined as growth rate of population over 20 years. “T” denotes time fixed effects, and “CD×T” denotes Census Division-time fixed effects. Standard errors clustered at the CZ level in parenthesis. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

2.4 Evidence on frictions to knowledge diffusion: cities and fields of knowledge

In this paper, we propose that the evidence in Table 1 is explained at least in part by the existence of frictions to the diffusion of ideas across space and fields of knowledge. These frictions prevent cities from optimally reallocating resources to take advantage of technological waves. Cities whose innovation portfolio is skewed towards expanding fields are better positioned to embrace new innovation opportunities and will become more attractive for workers and firms. In the following section, we formalize this hypothesis by embedding frictional knowledge diffusion in a spatial equilibrium model of endogenous growth.

The fact that knowledge diffusion is highly localized has been widely documented in the literature on the geography of innovation. Within this literature, a rich body of work, starting with [Jaffe et al. \(1993\)](#), has provided evidence of this localization by studying the spatial patterns of patent citations ([Murata et al., 2014](#); [Kerr and Kominers, 2015](#)). This evidence of localization is

Table 2: **Localized patent citations**

| Share of citations to local patents | | |
|--|-----------|-----------|
| 1940-1979 | 1980-2015 | 1940-2015 |
| 14.5% | 18.1% | 16.8% |

Notes: Probability that patents from any commuting zone cite other patents from the same commuting zone. Probabilities are computed using patents filed since 1940 whose first inventor is in one of the 373 commuting zones in the main sample. We weight each citation by the inverse of the total number of citations given by the citing patent.

confirmed in our citation data and it does not appear to weaken over time. As we show in Table 2, for patents filed since 1940, citations to the same commuting zone account for 16.8% of all citations.¹⁹ When we split the sample between an early (1940-1979) and a late sub-sample (1980 onwards) we find that, if anything, the evidence of localization becomes stronger over time. The share of citations to the same commuting zone of the citing patent is 14.5% in the early sub-sample, and increases to 18.1% in the late sub-sample.


Analogously, we find strong evidence of localization of patent citations in the technological space. The heatmap in Figure 4 displays the probability that a citation from each technological class-group on the vertical axis is directed toward each of the class-groups on the horizontal axis. The heatmap shows that citations are strongly concentrated along the diagonal, suggesting a high degree of technological localization in the diffusion of ideas. Appendix Figure B.2 displays the corresponding heatmaps separately for the early sub-sample (1940-1979) and late sub-sample (1980 onwards), showing that this evidence of technological localization remains strong over time.

3 Model

In this section, we develop a model that embeds an endogenous growth component into a spatial equilibrium framework. In our model, growth occurs through innovations that improve the current stock of ideas, and there are frictions to the diffusion of ideas across space and fields of knowledge.

¹⁹Since patent citations are not consistently available in the earlier decades, when considering citations we restrict the sample to all the patents filed since 1940. A separate section containing referenced patents was formally introduced in patent documents only in 1947. In constructing these statistics, we only consider citations to and from commuting zones included in our sample, and weight each citation by the inverse of the total number of citations given by the citing patent. By doing so, each citing patent has a weight of one in our sample.

Figure 4: Patent citations across fields



Notes: Probability that patents from the class-group on the vertical axis cite patents from the class-group on the horizontal axis. Probabilities are computed using patents filed since 1940 whose first inventor is in one of the 373 commuting zones in the main sample. We weight each citation by the inverse of the total number of citations given by the citing patent. Class-groups are described in Appendix Table A.1.

The theory formalizes the feedback between changes in the innovation landscape and the evolution of the economic geography outlined in the previous sections. Thus, we provide a theory rationalizing the reduced-form relationship between population growth and exposure to technological waves documented in Section 2. We quantify our model in Section 4 and evaluate its ability to replicate the reduced form relationship that we document in Section 2.

Before proceeding, we want to emphasize a conscious choice that we have made in developing our baseline model. We have attempted to present a parsimonious framework, so that we can focus on the novelty of our proposed mechanism and provide a transparent illustration of its mechanics. Appendix F shows how the baseline model can be extended, without a prohibitive loss in tractability, to incorporate overlapping generations, costly migration, trade, and local agglomeration or congestion forces. While important from a welfare perspective, these additional economic forces are not necessary to present our mechanism, so we choose to abstract from them in our baseline model.

3.1 Environment

We consider an economy comprising a finite set N of locations (cities) and a finite set S of sectors. Time is discrete and indexed by t . At each point in time, the economy is populated by a mass L_t of individuals. To simplify the notation, in what follows, we denote by N and S both the sets of locations and sectors and their cardinality and by L_t both the set of individuals at time t and its mass.

3.1.1 Preferences, endowments, and demographics

In each period, a new generation of individuals is born in the location of their parents and makes migration and occupational decisions. Individuals live for one period and at the end of the period have f_t children. There are no moving costs across locations. Given that we take a time period to be 20 years throughout our empirical and quantitative analysis, the assumptions about the demographic structure should be interpreted with this time horizon in mind.

Agents make migration and occupational choices to maximize expected utility, subject to idiosyncratic utility draws that affect the agent-specific desirability of each location-sector. Specifically, at the beginning of the period, each individual $i \in L_t$ receives a full set of stochastic utility draws, one for each location-sector in the economy:

$$\mathbf{x}_i = \{x_{n,s,i}\}_{(n,s) \in N \times S}.$$

Each value $x_{n,s,i}$ is a random draw from a Fréchet distribution with shape parameter $\zeta > 1$. If individual i chooses location-sector (n, s) , they obtain utility

$$U_{n,s,t}(\mathbf{x}_i) = u_n x_{n,s,i} c_{n,s,i,t}, \quad (2)$$

where u_n is the level of time-invariant amenities in city n and $c_{n,s,i,t}$ denotes consumption of the final good by individual i in location-sector (n, s) at time t . As we explain below, agents will face uncertainty about their consumption $c_{n,s,i,t}$ when choosing a location-sector (i.e., after the idiosyncratic utility draws \mathbf{x}_i are realized) because their labor productivity is going to be stochastic. Accordingly, they will choose the location-sector (n, s) that provides them with the highest expected utility. We return to this in Section 3.2.2.

3.1.2 Production and innovation technologies

Each agent i is endowed with one unit of labor that they supply inelastically with productivity q_i . Total output in the economy is given by a linear aggregator over individual productivity across all locations and sectors:

$$Y_t = \sum_{n \in N} \sum_{s \in S} L_{n,s,t} \mathbb{E}[q_{n,s,t}], \quad (3)$$

where $L_{n,s,t}$ denotes the mass of agents in the location-sector (n, s) and $\mathbb{E}[q_{n,s,t}]$ denotes their average productivity.²⁰

Individual productivity is determined endogenously by a choice to either imitate or innovate. The quality of ideas in the choice set of each agent is stochastic, and its distribution varies by location-sector.²¹ At the beginning of each period, every agent i in the new generation receives a full set of idiosyncratic, independently distributed draws:

$$\mathbf{z}_{n,s,i} = \left\{ z_{n,s,i}^l, \{z_{m,r,i}^x\}_{m,r \in N \times S} \right\}. \quad (4)$$

The first term, $z_{n,s,i}^l$, represents a random draw from the distribution of productivity among agents employed in location-sector (n, s) in the previous generation, whose cumulative distribution is denoted by $F_{n,s,t-1}(q)$. This draw can be interpreted as knowledge that individual i learns from their teacher, mentor, or manager, and can be imitated and adopted directly in production.²² If the agent chooses to adopt this idea in production, their lifetime productivity is

$$q_{n,s,i,t} = z_{n,s,i}^l.$$

The second set of terms, $\{z_{m,r,i}^x\}_{m,r \in N \times S}$, represents a full vector of random draws from each productivity distribution in all locations and sectors in the previous generation, with corresponding cumulative distributions $\{F_{m,r,t-1}(q)\}_{m,r \in N \times S}$. Note that this full set of draws includes local ones (i.e., $m = n$ and $r = s$). These draws can be interpreted as knowledge that the agent acquires by various channels of transmission, such as books, radio, television, and the internet, or even via casual interactions with local or non-local individuals. Although these ideas cannot be imitated and adopted directly in production, they can be used as an input for innovation. In particular, an agent employed in (n, s) can use an idea $z_{m,r,i}^x$ to innovate and achieve productivity

$$q_{n,s,i,t} = \frac{\epsilon_{n,s,t} \alpha_{r,t} z_{m,r,i}^x}{d_{(m,r) \rightarrow (n,s)}}. \quad (5)$$

²⁰Equation (3) can be obtained as an equilibrium representation between aggregate output and total labor in a setting where the fundamental production function involves a constant returns to scale production function between labor and location-sector specific intermediates (see, for example, Chapter 14.1 in [Acemoglu, 2009](#)).

²¹Our description of the process of innovation and knowledge diffusion builds on the model developed by [Buera and Oberfield \(2020\)](#), who study an environment in which the distribution of ideas endogenously converges to a Fréchet distribution.

²²[De la Croix et al. \(2018\)](#) develop a model of knowledge diffusion in which the institutions controlling the effectiveness of knowledge transmission between journeymen and apprentices contribute to explain differences across countries in long-run growth.

The term $d_{(m,r) \rightarrow (n,s)} \geq 1$ captures geographical and technological frictions that discount the effectiveness of knowledge transmission between the idea origin (m, r) and the idea destination (n, s) . The term $\alpha_{r,t}$ represents the importance of sector r in the innovation landscape. The higher the value of $\alpha_{r,t}$, the more effectively can knowledge in sector r be developed into innovation for any sector. We refer to changes in $\alpha_{r,t}$ as *technological wave* shocks. Note that this term is independent of the location-sector (n, s) of agent i . Rather, it is an intrinsic characteristic of the sector of origin r at time t . Finally, the term $\epsilon_{n,s,t}$ is a structural residual that captures the current effectiveness of innovation in (n, s) and is common to all innovators in that location-sector. It accounts for all residual factors that affect the productivity of the local sector but are not otherwise included in (5), such as the opening of production facilities, universities, and research centers.

3.2 Equilibrium

We normalize the price of the final good in each period to one. Thus, the wage of an agent i is simply their productivity $q_{n,s,i,t}$. Since agents live for one period, their consumption of final good is given by their own production:

$$c_{n,s,i,t} = q_{n,s,i,t}.$$

3.2.1 Diffusion of knowledge

Agent i in location-sector (n, s) chooses whether to imitate or innovate to maximize their productivity given their vector of idiosyncratic idea draws $\mathbf{z}_{n,s,i}$:

$$q_{n,s,i,t} = \max \left\{ z_{n,s,i}^l, \max_{m,r \in N \times S} \left\{ \frac{\epsilon_{n,s,t} \alpha_{r,t} z_{m,r,i}^x}{d_{(m,r) \rightarrow (n,s)}} \right\} \right\} \quad (6)$$

Equation (6) shows how this process can be divided into two steps. First, the agent chooses the best innovative idea available to them. Then they compare this idea with their imitation draw, and picks the one that yields higher productivity.

The following assumption, which we maintain throughout the paper, plays an important role in keeping the theory tractable:

Assumption A1. *The initial productivity distribution $F_{n,s,0}(q)$ in all location-sectors (n, s) is Fréchet with shape parameter $\theta > 1$ and scale parameter $\lambda_{n,s,0} > 0$:*

$$F_{n,s,0}(q) = e^{-\lambda_{n,s,0} q^{-\theta}}. \quad (7)$$

A multivariate Fréchet distribution with common shape parameter θ is max-stable. This implies that, under Assumption [A1](#), the resulting distribution over the max of Fréchet draws is also Fréchet with the same shape parameter.²³ Combining (6) with (7), we find that individual productivity at any time $t \geq 0$ is distributed Fréchet with shape parameter $\theta > 1$ and with scale parameter evolving according to the following law of motion:

$$\lambda_{n,s,t} = \underbrace{\lambda_{n,s,t-1}}_{\text{Imitation}} + \underbrace{\sum_{m \in N} \sum_{r \in S} \lambda_{m,r,t-1} \left(\frac{\epsilon_{n,s,t} \alpha_{r,t}}{d_{(m,r) \rightarrow (n,s)}} \right)^\theta}_{\text{Innovation}}. \quad (8)$$

Equation (8) plays a central role in our theory, since it describes the dynamics of the productivity distributions across location-sectors. The scale parameter $\lambda_{n,s,t}$ governing the distribution of the new generation in location-sector (n, s) is equal to the scale parameter of the previous generation augmented by a second term that captures inventive activities. This second term is composed by the sum of scale parameters across all location-sectors weighted by their applicability to location-sector (n, s) . The applicability term includes the importance of each field of knowledge, $\alpha_{r,t}$, as well as the local effectiveness of innovation, $\epsilon_{n,s,t}$, and is discounted by technological and physical distance between location-sector pairs, $d_{(m,r) \rightarrow (n,s)}$.

The scale parameter of the productivity distribution summarizes the stock of knowledge in each location-sector. In particular, using the result from Equation (8) that the productivity distribution in (n, s) at time t is Fréchet with shape parameter θ and scale parameter $\lambda_{n,s,t}$, we obtain a one-to-one mapping between local average productivity and the scale parameter:

$$\mathbb{E}[q_{n,s,t}] = \Gamma\left(1 - \frac{1}{\theta}\right) \lambda_{n,s,t}^{\frac{1}{\theta}}, \quad (9)$$

where $\Gamma(\cdot)$ denotes the gamma function. In other words, from Equation (8) we can also easily compute the dynamics of the average productivity of each location-sector.

The process of knowledge diffusion described by Equation (6) combined with the Fréchet assumption [A1](#) also implies that, conditional on innovating, the probability that an inventor in location-sector (n, s) builds upon an idea from any location-sector (m, r) at time t can be ex-

²³The same degree of tractability can be achieved without assuming independence, as in [Lind and Ramondo \(2019\)](#).

pressed as

$$\eta_{(m,r) \rightarrow (n,s)}^t = \frac{\lambda_{m,r,t-1} \left(\frac{\alpha_{r,t}}{d_{(m,r) \rightarrow (n,s)}} \right)^\theta}{\sum_{l \in N} \sum_{p \in S} \lambda_{l,p,t-1} \left(\frac{\alpha_{p,t}}{d_{(l,p) \rightarrow (n,s)}} \right)^\theta}. \quad (10)$$

3.2.2 Migration and occupational choice

At the beginning of period t , agents in the new generation observe sectoral and local shocks ($\alpha_{r,t}$ and $\epsilon_{n,s,t}$) but do not know their idiosyncratic idea draws. They have to form expectations about their future wages (determined by the idea draws) before making their migration and occupational decisions. Agent i moving to location-sector (n, s) has expected utility equal to

$$\mathbb{E}[U_{n,s,t}(\mathbf{x}_i)] = u_n x_{n,s,i} \mathbb{E}[q_{n,s,t}]. \quad (11)$$

Combining Equations (9) and (11), the probability that any newborn individual selects location-sector (n, s) is

$$\pi_{n,s,t} = \frac{\left(u_n \lambda_{n,s,t}^{\frac{1}{\theta}} \right)^\zeta}{\sum_{m \in N} \sum_{r \in S} \left(u_m \lambda_{m,r,t}^{\frac{1}{\theta}} \right)^\zeta}. \quad (12)$$

This expression is intuitive: the probability of choosing location-sector (n, s) is increasing in its expected productivity (which is proportional to $\lambda_{n,s,t}^{\frac{1}{\theta}}$) and its appeal due to amenities, u_n , relative to the average across location-sectors appearing in the denominator. The mass of agents in location-sector (n, s) at time t corresponds to

$$L_{n,s,t} \equiv \pi_{n,s,t} L_{t-1} f_t. \quad (13)$$

3.2.3 Equilibrium Definition

We now have all the ingredients to define an equilibrium of the model.

Definition 1. *Given*

$$L_0, \{ \lambda_{n,s,0} \}_{n,s \in N \times S}, \{ u_n \}_{n \in N}, \{ d_{(m,r) \rightarrow (n,s)} \}_{(m,r),(n,s) \in (N \times S)^2},$$

and a path of exogenous variables

$$\{ f_t \}_{t \geq 0}, \{ \alpha_{r,t} \}_{r \in S, t \geq 0}, \{ \epsilon_{n,s,t} \}_{n,s \in N \times S, t \geq 0},$$

an equilibrium is a path for the endogenous variables

$$\{\lambda_{n,s,t}, \pi_{n,s,t}, L_{n,s,t}\}_{n,s \in N \times S, t \geq 0}$$

that satisfies the following conditions:

1. Migration and occupational probabilities $\{\pi_{n,s,t}\}_{n,s \in N \times S, t \geq 0}$ satisfy Equation (12).
2. The path for $\{\lambda_{n,s,t}\}_{n,s \in N \times S, t \geq 0}$ satisfies the law of motion of Equation (8).
3. Population by location-sector $\{L_{n,s,t}\}_{n,s \in N \times S, t \geq 0}$ satisfies the transition identity (13).

It is readily verified that all equilibrium conditions listed in the definition above have a unique (and explicit) solution. Hence, a unique equilibrium exists and can be written in closed form for any given set of initial conditions and any given path of exogenous variables.

3.2.4 Existence and uniqueness of a balanced growth path

We define a balanced growth path as an equilibrium in which sectoral importance $\alpha_{r,t}$ and structural residuals $\epsilon_{n,s,t}$ are constant over time and in which scale parameters $\lambda_{n,s,t}$ grow at the same rate for all location-sectors (n, s) . Using Equation (12), these conditions also imply that migration and occupational choices (and, as a result, the distribution of people across locations and sectors) are constant over time.

Notice that Equation (8) can be rewritten in matrix form as

$$\vec{\lambda}_{t+1} = A_t \vec{\lambda}_t, \tag{14}$$

where $\vec{\lambda}_t$ is a $N \times S$ vector of all scale parameters $\lambda_{n,s,t}$ and A_t is the $(N \times S)^2$ diffusion matrix implied by Equation (8). In BGP, the matrix A_t is constant, and we denote it by A^* (in what follows, we use stars to denote variables at their BGP value).

From Equation (14), it is immediately evident that in BGP $\vec{\lambda}_t$ must be an eigenvector of A^* , with the associated eigenvalue equal to its gross growth rate $1 + g_\lambda^*$. The Perron-Frobenius theorem states that A^* has a unique positive eigenvector (and corresponding eigenvalue), provided that all entries in A^* are positive. A sufficient condition for A^* to have only positive entries is that frictions to knowledge diffusion $d_{(m,r) \rightarrow (n,s)}$ are strictly positive and finite for each combination of location-sector pairs. This proves the following:

Proposition 1. *Let $1 \leq d_{(m,r) \rightarrow (n,s)} < +\infty$ for all $(m,r), (n,s) \in (N \times S)^2$. Then, for each set of constant sectoral importance $\{\alpha_r^*\}_{r \in S}$ and structural residuals $\{\epsilon_{n,s}^*\}_{(n,s) \in N \times S}$, there exists a unique balanced growth path in which $\{\lambda_{n,s,t}\}_{(n,s) \in N \times S, t \geq 0}$ grow at constant rate, g_λ^* , with $g_\lambda^* > 0$. The gross growth rate $(1 + g_\lambda^*)$ is given by the unique largest eigenvalue of A^* (the Perron-Frobenius eigenvalue), and the normalized scale parameters $\{\lambda_{n,s,t}/(1 + g_\lambda^*)^t\}_{(n,s) \in N \times S, t \geq 0}$ correspond to the associated right eigenvector of A^* .*

As stated in the proposition, along the BGP, scale parameters grow at the same rate across location-sectors. However, different location-sectors have different scale parameters along the BGP because the entries of the eigenvector associated with the Perron-Frobenius eigenvalue are generically different from each other. This result is also apparent from the fact that, along the BGP, the following relationship holds for each location-sector (n, s) :

$$g_\lambda^* = (\epsilon_{n,s}^*)^\theta \sum_{m \in N} \sum_{r \in S} \left(\frac{\lambda_{m,r}}{\lambda_{n,s}} \right)^* \left(\frac{\alpha_r^*}{d_{(m,r) \rightarrow (n,s)}} \right)^\theta. \quad (15)$$

This equation illustrates that, conditional on a growth rate g_λ^* , the stationary value of the scale parameters (which can be obtained by inverting the system of equations implied by 15) depends on the matrix of diffusion frictions across location-sectors, $d_{(m,r) \rightarrow (n,s)}$, in addition to local and sectoral characteristics, α_r^* and $\epsilon_{n,s}^*$. Conversely, Equation (15) can be interpreted as stating that, conditional on relative productivities $\left(\frac{\lambda_{m,r}}{\lambda_{n,s}} \right)^*$, higher BGP growth can result from higher local and sectoral characteristics (α_r^* and $\epsilon_{n,s}^*$) or from lower frictions to knowledge diffusion ($d_{(m,r) \rightarrow (n,s)}$).²⁴

3.3 Log-linearized model dynamics

The central question we want to address in this paper is how technological waves affect the relative growth of cities. Thus, while the BGP is a useful benchmark, we are ultimately interested in the heterogeneous response of cities to technological waves. Indeed, our BGP analysis precluded technological waves because we held them constant by definition, $\alpha_{r,t} \equiv \alpha_r^*$. As a result, the relative size of cities does not change along the BGP. We next turn to characterize the model dynamics once we allow for technological wave shocks.

We study the dynamics of the model by log-linearizing the equilibrium conditions around the

²⁴Huang and Zenou (2020) is another paper that studies the BGP properties of an endogenous growth model with spillovers across multiple sectors. While the setting for idea diffusion in Huang and Zenou (2020) is different from ours, in both models the Perron-Frobenius theorem is central to establishing the existence and uniqueness of a BGP equilibrium.

BGP. By log-linearizing our model, we are able to obtain sharp and intuitive characterizations of what drives city growth as a response to technological waves.²⁵ For our characterization of the log-linearized dynamics, we assume that at time $t - 1$ the economy is in a BGP (i.e., the average productivity in each location-sector grows at the same rate, and as a result, the distribution of population across locations is constant). At time t , the economy is hit by technological wave shocks $\{\hat{\alpha}_{r,t}\}_{r \in S}$, where hats denote log-deviations from BGP values.

First, we consider the dynamics of the scale parameter of the local distribution of productivity, $\lambda_{n,s,t}$. Log-linearizing Equation (8) yields

$$\hat{\lambda}_{n,s,t} = \frac{\theta(\epsilon_{n,s}^*)^\theta}{1 + g_\lambda^*} \sum_{m,r} \left(\frac{\lambda_{m,r}}{\lambda_{n,s}} \right)^* \left(\frac{\alpha_r^*}{d_{(m,r) \rightarrow (n,s)}} \right)^\theta \hat{\alpha}_{r,t}. \quad (16)$$

Multiplying and dividing the right-hand side of (16) by g_λ^* , and using (10) and (15), we derive the following proposition that links changes in local sectoral productivity to technological wave shocks via the strength of the knowledge diffusion links between the perturbed sectors and the receiving location-sector.

Proposition 2. *The log deviation of the scale parameter of the productivity distribution of each location-sector (n, s) from the BGP, $\hat{\lambda}_{n,s,t}$, is equal to the sum over all sectors $r \in S$ of the sectoral shock to r , $\hat{\alpha}_{r,t}$, weighted by the reliance of innovation in (n, s) on ideas from sector r , as measured by the probability of building on ideas from sector r when innovating, $\eta_{r \rightarrow (n,s)}^* \equiv \sum_{m \in N} \eta_{(m,r) \rightarrow (n,s)}^*$ (as defined in Equation 10); i.e.:*

$$\hat{\lambda}_{n,s,t} = \frac{\theta g_\lambda^*}{1 + g_\lambda^*} \sum_{r \in S} \eta_{r \rightarrow (n,s)}^* \hat{\alpha}_{r,t}. \quad (17)$$

Proposition 2 implies that, other things being equal, the *sensitivity* of local productivity $\hat{\lambda}_{n,s,t}$ to shocks to any given sector, $\hat{\alpha}_{r,t}$, is increasing in the probability of drawing knowledge from that sector to innovate, $\eta_{r \rightarrow (n,s)}^*$. The existence of frictions to knowledge diffusion implies that this reliance on ideas from sector r , $\eta_{r \rightarrow (n,s)}^*$, depends on how “close” sector r is to (n, s) in the geographical and technological space. Quantitatively, given the evidence discussed in Section 2.4 on the localized nature of knowledge diffusion, the local stock of knowledge in the same sector, $\lambda_{n,r}$, will play a decisive role in determining this reliance term. From Equation (12) it is also immediate to see that this stock of knowledge is tightly related to the local share of population employed in

²⁵Moreover, in our calibrated model (which fully accounts for the nonlinear dynamics), we find that the log-linear model captures most of the variation generated by the full model in the relevant time horizon.

the same sector.²⁶ This implies that the sensitivity of local productivity to shocks to any given sector is strongly linked to the prevalence of the sector in the local economy.

Consider now the dynamics of population in location n . Combining Equation (12) with the definition $\pi_{n,t} \equiv \sum_{s \in S} \pi_{n,s,t}$ and log-linearizing the resulting expression for any arbitrary deviation of $\{\lambda_{m,s,t}\}_{m,s \in N \times S}$ from their BGP values yields

$$\hat{\pi}_{n,t} = \frac{\zeta}{\theta} \sum_{s \in S} \left\{ (1 - \pi_n^*) \pi_{s|n}^* \hat{\lambda}_{n,s,t} - \sum_{m \neq n} \pi_{m,s}^* \hat{\lambda}_{m,s,t} \right\}, \quad (18)$$

where $\pi_{s|n}^*$ denotes the probability of being employed in sector s conditional on living in location n and where $\sum_{m \neq n}$ is the sum across all elements of N except n . Equation (18) implies an intuitive condition that controls whether a city grows or shrinks relative to the rest of the economy. A location grows if and only if changes in local sectoral productivities, weighted by the incidence of each sector in the city, are larger than the average corresponding changes for the rest of the economy:

$$\hat{\pi}_{n,t} > 0 \iff \sum_{s \in S} \pi_{s|n}^* \hat{\lambda}_{n,s,t} > \sum_{s \in S} \sum_{m \neq n} \frac{\pi_{m,s}^*}{1 - \pi_n^*} \hat{\lambda}_{m,s,t}.$$

Finally, we characterize changes in population shares in terms of pre-shock values and fundamental model parameters, rather than stating them as a function of endogenous changes in the shape parameters, $\hat{\lambda}_{n,s,t}$ (as derived in Equation 18). By combining Equations (17) and (18), we derive the following proposition, which summarizes the population dynamics implied by the model in response to technological wave shocks directly.

Proposition 3. *The log-change in the population shares of location n , $\hat{\pi}_{n,t}$, depends on technological waves shocks as follows:*

$$\hat{\pi}_{n,t} = \frac{\zeta g_\lambda^*}{1 + g_\lambda^*} \sum_{r \in S} \sum_{s \in S} \left\{ (1 - \pi_n^*) \pi_{s|n}^* \eta_{r \rightarrow (n,s)}^* - \sum_{m \neq n} \pi_{m,s}^* \eta_{r \rightarrow (m,s)}^* \right\} \hat{\alpha}_{r,t}. \quad (19)$$

To interpret Equation (19) and better illustrate the economic mechanism at play, we first consider a simplified version of the model in which knowledge flows across sectors are of second-order importance relative to flows within sectors. In particular, we impose the following assumptions:

Assumption A2. *(for illustration purposes only)*

²⁶To see this, note that in the limit case of $\theta = \zeta$, $\alpha_r^* = \alpha_s^*$, and $d_{(n,r) \rightarrow (n,s)} = \bar{d}$ for all $r, s \in S$, the reliance of (n, s) on ideas from r , $\eta_{r \rightarrow (n,s)}^*$, is exactly equal to the local sectoral share, $\pi_{r|n}^*$, i.e., the employment share in r in location n .

1. Frictions to knowledge diffusion across sectors are large enough so that, effectively, knowledge flows are only within sector, that is, $\eta_{s \rightarrow (n,s)}^* \approx 1$ for all $s \in S$
2. The size of any given city is negligible with respect to the overall economy, that is, $\sum_{m \neq n} \pi_m^* \approx 1$, for all $n \in N$.

Under Assumption A2, we can combine Equations (17) and (18) to obtain the following expression for the change in population shares:

$$\hat{\pi}_{n,t} \stackrel{A2}{=} \frac{\zeta g_\lambda^*}{1 + g_\lambda^*} \left(\sum_{s \in S} \pi_{s|n}^* \hat{\alpha}_{s,t} - \bar{\hat{\alpha}}_t \right), \quad (20)$$

where $\bar{\hat{\alpha}}_t \equiv \sum_{s \in S} \pi_{\cdot,s}^* \hat{\alpha}_{s,t}$, with $\pi_{\cdot,s}^*$ denoting the share of the national population employed in sector s , is a common term across all locations. Under these simplifying assumptions, cities' differential patterns of population growth depend on a weighted average of aggregate sectoral shocks, with the weights corresponding to the city's pre-existing pattern of specialization across sectors. Equation (20) thus provides a rationale for the reduced-form relationship between exposure to technological waves and population growth that we have documented in Section 2.3.²⁷

Consider now the general case in which knowledge flows across fields are non-negligible (i.e., $\eta_{s \rightarrow (n,s)}^* < 1$) and each city has a non-trivial size. In this case, Equation (20) captures only part of the total effect of technological waves on population growth described in Proposition 3.²⁸ In the general case (stated in Equation 19), because of geographical frictions to idea diffusion, cities display different degrees of reliance of local innovation on ideas from each sector in the economy (as captured by $\eta_{r \rightarrow (n,s)}^*$). This implies that productivity growth in *all* sectors will be larger (smaller) in cities where expanding (shrinking) fields are more prominent, thus amplifying the “shift-share” effect in Equation (20). In other words, for the general case, in response to technological wave shocks, localized knowledge flows *across* fields amplify fluctuations in productivity growth and,

²⁷Equation (20) also implies that under Assumption A2, n grows (shrinks) if and only if the average local exposure to the technological wave is larger (smaller) than the average exposure for the economy:

$$\hat{\pi}_{n,t} > 0 \stackrel{A2}{\iff} \sum_{s \in S} \pi_{s|n}^* \hat{\alpha}_{s,t} > \sum_{s \in S} \pi_{\cdot,s}^* \hat{\alpha}_{s,t}. \quad (21)$$

²⁸If we dispense with the second part of our simplifying assumption of all cities having negligible size, we obtain that

$$\hat{\pi}_{n,t} \stackrel{A2,1}{=} \frac{\zeta g_\lambda^*}{1 + g_\lambda^*} \sum_{s \in S} \left\{ (1 - \pi_n^*) \pi_{s|n}^* - \sum_{m \neq n} \pi_{m,s}^* \right\} \hat{\alpha}_{s,t}. \quad (22)$$

via Equation (19), fluctuations in population dynamics.²⁹

Taking stock and looking ahead Propositions 2 and 3 show that frictions to knowledge diffusion across geographical areas and technological fields imply rich and heterogeneous effects of technological waves on the evolution of local productivity and on the distribution of population across cities. These results leverage the log-linearization around a BGP to provide a simple characterization of the city dynamics generated by technological waves. In our model calibration and quantification exercises, we do not rely on any log-linearization of the model. Instead, we solve for the entire non-linear dynamics. We find, however, that the log-linearization results provide useful insights to understand the mechanics of the model and they approximate well the response of the dynamic model over the relevant time horizon. In the next section, we show how we bring the model to the data to infer the key parameters and unobserved variables. We then use the calibrated model to obtain the quantitative results in Section 5.

4 Model calibration

The remainder of the paper is devoted to performing a quantitative analysis of the impact of technological waves on city dynamics through the lens of the model developed in the previous section. In this section, we describe our calibration strategy. In the next section, we report our quantification results and also provide some counterfactual analysis.

We present two quantification exercises. First, we start calibrating and analyzing the predictions of our model for the second half of the twentieth century. The technological landscape during this time period is broadly characterized by the decline of manufacturing and the rise of IT and medicine-related fields. Since there are more data available for this later time period, we use this exercise to calibrate some of the deep parameters governing the dispersion of Fréchet shocks—which we then maintain for our second quantification exercise. In this second exercise, we extend the analysis to the first half of the twentieth century, capturing the rise of manufacturing-related fields and the decline of agriculture that we have documented in Section 2.

As we elaborate in Appendix E, we find that the city dynamics generated by the calibrated process of knowledge diffusion are very protracted and exhibit substantial inertia.³⁰ In fact, the

²⁹It is also possible to characterize the speed of convergence (e.g., as measured by the half life) of productivity or population using Equations (16) and (20). We leave this analysis for future work.

³⁰Kleinman et al. (2021) consider a model where capital accumulation and frictional mobility also give rise to slow convergence toward the steady state. In our model, slow convergence toward the steady state is generated by

half life of a typical technological wave shock is over a hundred years. To simplify the exposition and distill the impact of technological waves from the effect of the model’s inertia, we focus our baseline analysis on the *contemporaneous* effect of technological waves on city growth. That is, we separately study the effect of technological waves in the first and second half of the twentieth century on contemporaneous city dynamics. We do so by assuming that the economy is in BGP in the first period of each of the two intervals that we study, so that in the absence of shocks (such as technological waves), all cities would grow at the same rate. However, in Appendix E we show that our results are robust when allowing for the model’s inertial dynamics. In particular, we demonstrate that when we calibrate the model for the entire century, the contribution of technological waves in the late period, after normalizing city dynamics by the model’s inertia generated in the early period, is very similar to the effect we find in our baseline exercise. Intuitively, this result follows from the fact that the model dynamics are initially well approximated by the log-linear response discussed in Section 3. As a result, the compounded effect of technological waves are essentially additively separable in the relevant time horizon.

4.1 Overview of the calibration strategy

The model has a recursive structure that allows us to calibrate parameters and unobserved variables sequentially by making a limited set of transparent assumptions on how to map the models objects to data on population, income, and patenting. As in the empirical analysis of Section 2, throughout the model calibration and quantification, we set the model period to 20 years, we let N be the set of 1990 commuting zones that accounted for at least 0.02% of the total population for each decade since 1890, and we define sectors as the 11 class-groups detailed in Appendix Table A.1.

As we have already mentioned, we run our quantitative analysis starting from two sets of initial conditions—each corresponding to the outset of two of the most striking episodes of technological and geographical transformation in the last century. In particular, we split the 1910-2010 period into two long intervals: an early interval (\mathcal{E})—which starts in the 20-year period around 1910 and ends in the 20-year period around 1950 and is characterized by the rise of manufacturing-related fields and the decline of agriculture—and a late interval (\mathcal{L})—which starts in the 20-year period around 1970 and ends in the 20-year period around 2010 and is characterized by the decline of manufacturing and the rise of IT and medicine-related fields. We analyze these two intervals

a distinct and, possibly, complementary mechanism, that is, the slow diffusion of ideas across locations and fields of knowledge.

separately by resetting the initial conditions in the first period of each interval. As we show in Appendix E, the model calibrated for the early interval economy \mathcal{E} still exhibits substantial inertia going into the second half of the century. However, once we normalize by this inertia, the effect of the technological waves of the late period \mathcal{L} in the model with inertia are very similar to that of our baseline exercise.

For each quantification exercise, our calibration proceeds in three steps. In the first step, we infer exogenous amenities $u_{n,\mathcal{L}}$,³¹ the path of local productivities $\lambda_{n,s,t}$, and aggregate fertility f_t , and simultaneously pin down the time-invariant parameters ζ and θ by matching moments on the dispersion of income and population across cities. The time-invariant parameters, ζ and θ , are only calibrated in the first calibration exercise (i.e., for the late interval, 1970-2010). In particular, to pin down ζ and θ we use cross-sectional empirical moments from 1990, for which we have the most recent and complete data on population, income, and patenting.³² We then show that the model accurately reproduces the relationship between city size and income despite not being directly targeted. In the second step, we infer the costs of knowledge transmission $d_{(m,r)\rightarrow(n,s)}$ by deriving and estimating a gravity equation for idea flows using patent citations data. In the third step, we recover technological wave shocks $\alpha_{s,t}$ and structural residuals $\epsilon_{n,s,t}$ via the law of motion for local productivities.

4.2 First Step: Amenities and productivity

In the first step of our calibration, we jointly calibrate the shape parameters of the Fréchet distributions of utility draws, ζ , and the initial distribution of productivity, θ . Here, we also recover the values of local amenities $u_{n,\mathcal{L}}$, as well as the full path of scale parameters $\lambda_{n,s,t}$ and aggregate fertility f_t .

4.2.1 Productivity distribution

Consider first the scale parameters $\{\lambda_{n,s,t}\}$. They summarize the stock of knowledge in each location-sector, and they are at the core of the quantitative analysis: Higher values of $\lambda_{n,s,t}$ imply higher local income, higher ability to attract population, and higher potential to innovate and grow in the future. We draw on Schumpeterian endogenous growth theory to postulate (and

³¹We assume exogenous amenities to be time-invariant within each interval, but to vary between the two intervals. We denote them accordingly as $u_{n,\mathcal{E}}$ for the early interval and as $u_{n,\mathcal{L}}$ for the late interval.

³²Since we assign patents according to their filing year, patents data and citations in the most recent observation (2010) might suffer from truncation issues. Similarly, data on income and population for the 2010 observation are only available from the ACS, which offers a less complete picture than the 1990 census.

later validate) a direct mapping between the stock of knowledge and the stock of patents in each location-sector. We then exploit our model structure to link the stock of knowledge in a location-sector to its scale parameter $\lambda_{n,s,t}$, and use this mapping to recover the full path of scale parameters $\lambda_{n,s,t}$.

We establish a connection with Schumpeterian growth theory by interpreting each idea draw $q_{n,s,i,t}$ as a point in a quality ladder with $h_{n,s,i,t}$ steps, so that

$$\log(q_{n,s,i,t}) = ah_{n,s,i,t}, \quad (23)$$

where $a > 0$ is the step size of each quality improvement. We define the local stock of knowledge in location-sector (n, s) at time t as the average number of steps in the local productivity distribution of innovators, $\mathbb{E}[h_{n,s,i,t}]$. In other words, the knowledge stock captures the average number of rungs that have been climbed in the “knowledge ladder” by innovators. Using that ideas $q_{n,s,i,t}$ are distributed Fréchet with scale $\lambda_{n,s,t}$ and shape θ , we have that the stock of knowledge satisfies

$$\mathbb{E}[h_{n,s,i,t}] = \frac{\gamma + \log(\lambda_{n,s,t})}{a\theta}, \quad (24)$$

where γ denotes the Euler-Mascheroni constant.

We then assume the following parametric relationship between the stock of knowledge in location-sector (n, s) at time t and the corresponding stock of patents:

$$\mathbb{E}[h_{n,s,i,t}] = \log \left[\tilde{G}_t \times \left(1 + \sum_{\tau=0}^{\tau_{max}} \xi^\tau Pat_{n,s,t-\tau} \right) \right], \quad (25)$$

where $Pat_{n,s,t}$ denotes the total number of patents filed at time t in location-sector (n, s) , $0 < \xi < 1$, and \tilde{G}_t is a time-varying factor.³³ That is, we assume a concave (logarithmic) relationship between the accumulated flow of all patents in a location-sector discounted at rate ξ and the stock of knowledge in this location-sector. The time-varying factor \tilde{G}_t captures changes in the innovative content of patents that are independent from a location-sector (e.g., patents becoming more defensive over time). We discuss the calibration of these parameters below.

³³We add one to the stock of patents in each location-sector to assign a meaningful value to cases in which patenting is zero.

Combining Equations (24) and (25), we obtain that

$$\lambda_{n,s,t} = G_t \times \left(1 + \sum_{\tau=0}^{\tau_{max}} \xi^\tau Pat_{n,s,t-\tau} \right)^\sigma, \quad (26)$$

where $\sigma \equiv a \times \theta$ represents the elasticity of $\lambda_{n,s,t}$ with respect to the observed stock of patents and $G_t \equiv e^{-\gamma} \times \tilde{G}_t^\sigma$. The elasticity σ converts the variation in the local stock of patents into meaningful variation in the average productivity across location-sectors. Thus, Equation (26) provides a mapping between the stock of patents in a location-sector at a given point in time (which we can observe) and a central element of our theory, the scale parameter of the Fréchet distribution in that location-sector, $\lambda_{n,s,t}$ (which is unobservable).³⁴ To assess the plausibility of our calibration strategy, we show in Section 4.2.3 that the resulting calibrated mapping between the scale parameters and patent stocks generates a correlation between city size and income that is very close to the one observed in the data despite not being targeted in our calibration.

We calibrate σ and θ to jointly match the standard deviation of log-income per capita across cities (in the sample of 373 CZs) and in the overall population in 1990, which are equal to 0.19 and 0.67, respectively.³⁵ The constant G_t is set to induce an aggregate growth in income per capita of 2% per year.³⁶ Finally, we set $\xi = 0.5$ and $\tau_{max} = 2$, reflecting an assumption that the contribution of past patents to variation in the local stock of knowledge halves every 20 years and vanishes after 60 years. This depreciation rate is in line with the 4% obsolescence found in Caballero and Jaffe (1993).³⁷

4.2.2 Amenities, preference draws, and fertility

Consider now local amenities $u_{n,\mathcal{L}}$ and the shape parameter of the distribution of utility draws ζ . Given values for ζ , θ , and $\lambda_{n,s,t}$, we calibrate local amenities to exactly match population by city in the first period of the interval (1970).³⁸ The value of ζ is then calibrated to match the standard

³⁴Equation (25) defines the stock of knowledge in a reduced-form way following the strategy in Hall et al. (2001). Alternatively, it is possible to define the probability of patenting as the probability of improving existing ideas, following the approach in Kortum (1997). This approach delivers a direct mapping between the scale parameter and the number of patents without relying on the reduced-form definition of the stock of knowledge, yielding an expression similar to Equation (26).

³⁵The standard deviation of log-income in the overall population is taken from Krueger and Perri (2006).

³⁶We choose units of the final good so that the geometric average of $\lambda_{n,s}^{\frac{1}{\sigma}}$ is equal to one in the first time period.

³⁷This obsolescence rate is inferred from the rate of decline in patent citations. Our results also go through with a higher obsolescence rate of around 8% inferred from patent renewal rates; see Anzoategui et al. (2019) and the references therein.

³⁸We normalize amenities to have a geometric mean of one.

Table 3: **Parameter values and targets**

| Parameter | Value | Target | Model | Data |
|-----------|-------|--|-------|------|
| σ | 0.21 | s.d. log-income (across CZs), 1990 | 0.19 | 0.19 |
| θ | 2.00 | s.d. log-income (overall), 1990 | 0.67 | 0.67 |
| ζ | 5.50 | s.d. log-population (across CZs), 1990 | 1.07 | 1.07 |

Notes: Standard deviation of log-income for the overall population is taken from [Krueger and Perri \(2006\)](#). Standard deviation of log-income and log-population across CZs are author’s calculations from the NHGIS.

deviation of log-population across cities in 1990. The intuition for the identification is that a higher value of ζ implies lower dispersion in the utility draws among newborn agents, so that differences in the desirability of locations, given by amenities and productivity, are more strongly reflected in migration choices.³⁹ While we assume time-invariant amenities (within each interval) and do not match population by city in each period, we show in the discussion below that the joint calibration of θ , ζ , and σ generates a realistic dispersion of income and population across cities.

Finally, we calibrate the path of fertility, f_t , to match total population by period in our sample. Notice that, in the absence of moving costs, this is equivalent to assuming that the aggregate increase in population occurs through migration from abroad, fertility, or a combination of the two.

4.2.3 Discussion


Table 3 shows the values of θ , ζ , and σ calibrated through this procedure. The corresponding data moments are matched exactly by construction.⁴⁰

There are two key aspects of this calibration strategy that are worth further discussion. First, the mapping of $\lambda_{n,s,t}$ to the stock of patenting (Equation 26) includes a size effect in which larger cities have, other things being equal, higher average productivity. The existence of a correlation between size and productivity is a well-known empirical regularity (see, e.g., [Glaeser and Gottlieb, 2009](#)) that can emerge as a result of a range of theoretical mechanisms (e.g., sorting, variety, local learning productivity spillovers, and higher availability of productive inputs). While the model is silent on the underlying mechanism behind this correlation (besides the fact that more productive cities will *attract* more population), what is crucial for the quantitative performance

³⁹This identification of the dispersion of idiosyncratic preference draws follows a similar intuition as [Peters \(2019\)](#).

⁴⁰In Appendix Figure B.3, we show computationally that there are unique values of the three parameters that jointly match those data moments.

Figure 5: **Population and Income: Data vs. Model (untargeted)**



Notes: Bin-scatter plot of the relationship between log-population and log-income per capita in the data (blue) and the model (red) in 1990. All variables are displayed as deviations from the mean.

of the model is that the resulting elasticity of population with respect to income per capita is empirically accurate. Figure 5 shows a bin-scatter plot of the relationship between log-population and log-income in 1990, both in the model and in the data. Although this correlation is not directly targeted in the calibration, the model captures it closely.⁴¹

Second, in quantifying the model we assume that residential amenities are time-invariant. This assumption is crucial for the identification of the shape parameter ζ but comes at the cost of not matching population by city exactly after the first period. As we show in Section 5, even without time-varying amenities, the model goes a long way in accounting for population growth by city over the last century.⁴²

4.3 Second Step: Gravity equation for knowledge flows

In the second step of the calibration, we recover the parameters controlling knowledge transmission costs, $d_{(m,r) \rightarrow (n,s)}$. To this end, we derive a gravity representation for knowledge flows that we estimate using data on patent citations. We parametrize frictions to knowledge diffusion as

⁴¹The slope of the regression line is equal to 15.9 in the model and 12.7 in the data.

⁴²It would be possible to pursue alternative modeling and calibration strategies. For example, we show in Appendix F that, given a value for ζ , allowing for time-varying residential amenities would be an immediate extension of the model.

multiplicatively separable between a geographical component and a technological one:

$$d_{(m,r)\rightarrow(n,s)} = e^{\delta^G \mathbf{1}_{m \neq n} + \delta_{r \rightarrow s}^K}, \quad (27)$$

where δ^G controls the effectiveness of knowledge flows *across* locations relative to flows *within* locations and where $\delta_{r \rightarrow s}^K$ controls the applicability of ideas from sector r for innovation in sector s .

Combining Equations (10) and (27) and taking logs on both sides yields

$$\log(\eta_{(m,r)\rightarrow(n,s)}^t) = \phi_{m,r,t}^0 + \phi_{n,s,t}^1 - \theta \delta^G \mathbf{1}_{m \neq n} - \theta \delta_{r \rightarrow s}^K, \quad (28)$$

where ϕ^0 and ϕ^1 represent idea origin and idea destination-time fixed effects, respectively.

Equation (28) illustrates that bilateral citation probabilities $\eta_{(m,r)\rightarrow(n,s)}^t$ depend on the composite parameters $\theta \delta^G$ and $\theta \delta_{r \rightarrow s}^K$. To recover those composite parameters, we estimate Equation (28) using data on patent citations across location-sector pairs from the 1990 period (i.e., using patents filed between 1980 and 1999). We compute $\eta_{(m,r)\rightarrow(n,s)}^t$ as the share of citations given by patents in (n, s) and directed to patents in (m, r) .⁴³

Table 4 shows OLS estimates of the composite parameter $\theta \delta^G$. In our baseline specification, reported in column (1), we replace zero-valued outcomes with the minimum among positive values.⁴⁴ Using our estimates from column (1), in combination with the estimate of θ , we obtain $\delta^G = 4.34$. The coefficient implies highly localized knowledge flows, with the effectiveness of transmission across locations, defined as $e^{-\delta^G}$, estimated at around 1.31% of the effectiveness of transmission within locations. Notice that, despite the apparent low effectiveness of transmission, the overall weight of ideas from outside locations may still be large in determining innovation in n , since transmission can happen from *all* the other locations $m \neq n$. Column 2 shows the same regression when only positive values of $\eta_{(m,r)\rightarrow(n,s)}^t$ are used in the estimation. The estimate still reveals highly localized knowledge flows, but the coefficient declines in absolute value, suggesting that, as

⁴³Note that the direction of the arrow from (m, r) to (n, s) denotes knowledge flows going from the *cited* patent to the *citing* patent. Every citing patent in our regression has a total weight of one. In other words, every observation is weighted by the inverse of the total number of citations given by (n, s) . To account for the fact that *local* knowledge transmission is more likely to be tacit and less likely to be captured by citations, we further assume that every grant cites one patent from its own location and sector. This also guarantees that all patents, including the ones with no backward citations, are included in the estimation.

⁴⁴Notice that because of the fractional nature of the assignment of patents to class-group, $\eta_{(m,r)\rightarrow(n,s)}^t$ cannot be written as a count divided by an exposure variable, which implies that Equation (28) cannot be estimated via Poisson Pseudo-Maximum Likelihood (PPML). However, different mappings of $\eta_{(m,r)\rightarrow(n,s)}^t$ into a count variable lead to similar estimates when using PPML.

Table 4: Gravity equation for knowledge flows

| | Log share of citations | |
|---------------------------------|------------------------|---------------------|
| | (1) | (2) |
| Origin CZ \neq Destination CZ | -8.677*** (.046) | -2.892*** (.022) |
| Origin location-sector FE | yes | yes |
| Destination location-sector FE | yes | yes |
| Origin-Destination sector FE | yes | yes |
| # Obs. | 16,834,609 | 1,267,834 |
| R^2 | 0.32 | 0.69 |
| Zero values | Set to min | No |

Notes: OLS estimates. The sample includes patents filed between 1980 and 1999. Observations are all the combinations of pairs of location-sectors. The dependent variable is the logarithm of the share of citations given by each destination location-sector to each origin location-sector, where each citing patent is given a weight of one. We assume that every grant cites one patent from its own location and sector. Standard errors clustered at the destination location-sector level in parenthesis. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

expected, zero values are concentrated among pairs of different locations. Yet the effectiveness of transmission across locations is estimated at around 6% of the effectiveness within locations.


The same regression also delivers a full set of bilateral transmission costs across sectors ($\delta_{r \rightarrow s}^K$), which we show in a heatmap in Figure 6. As expected, these costs are estimated to be lower within sectors (along the diagonal of the heatmap), although all pairs of sectors display some degree of knowledge exchange that, in some cases, is far from negligible, such as in the cluster of class-groups G1 (“Physics”) and H1 (“Electricity”).

4.4 Third Step: Technological waves and structural residuals

In the third step of the calibration, we use the estimates of θ and $\delta_{r \rightarrow s}^K$ and the values of $\lambda_{n,s,t}$ in combination with the law of motion (8) to recover technological wave shocks ($\alpha_{s,t}$) and structural residuals ($\epsilon_{n,s,t}$).

For all periods t , we first guess the full vector of technological wave shocks $\{\alpha_{s,t}\}_{s \in S}$. Given this guess, we use Equation (8) to recover the full set of structural residuals. By construction, this step rationalizes the path of $\lambda_{n,s,t}$ for any initial guess of $\{\alpha_{s,t}\}_{s \in S}$. Hence, to complete the

Figure 6: **Knowledge transmission costs across sectors**



Notes: Estimated (OLS) coefficients $\delta_{r \rightarrow s}^K$, from regression of Table 4, Column 1. The sample includes patents filed between 1980 and 1999. Observations are all the combinations of pairs of location-sectors. The dependent variable is the logarithm of the share of citations given by each destination location-sector to each origin location-sector, where each citing patent is given a weight of one. We assume that every grant cites one patent from its own location and sector. Rows correspond to citing (idea destination) sectors. Columns correspond to cited (idea origin) sectors. Number of observations: 16,834,609. R^2 : 0.32. Class-groups are described in Appendix Table A.1.

identification, we need to impose S additional conditions. We proceed by making the following identifying assumption: in each period, variation across sectors in *average* productivity growth (and thus, scale parameters) is explained by technological waves and their interaction with the endogenous process of knowledge creation and diffusion implied by Equation (8). As a result, structural residuals, $\epsilon_{n,s,t}$, explain the remaining variation in productivity growth across locations for any given sector (i.e., the variation that is not explained by the common growth component in productivity across locations). Specifically, to implement our identification, we make the following assumption:

Assumption A3. *The sets of adjusted structural residuals, $\{\epsilon_{n,s,t}^\theta\}_{n,s \in N \times S, t \geq 0}$, have a common average for each sector and time period that we normalize to one:*

$$\mathbb{E}[\epsilon_{n,s,t}^\theta] = 1, \quad \forall s \in S, t \geq 0. \quad (29)$$

Combining Equation (29) with the law of motion (8), we can then recover a unique set of technological wave shocks, $\{\alpha_{s,t}\}_{s \in S}$ and structural residuals $\{\epsilon_{n,s,t}\}$. In practice, given the mapping we established in the first step between knowledge stocks and scale parameters, our identification assumption implies that the common component of shifts in the knowledge stock in a given sector across locations is attributed to the technological wave. The rest of the sector-location specific variation is attributed to the structural residuals.

Our identification of technological waves is thus linked to the co-movement in knowledge stocks across locations in a given sector and, ultimately, to the common movements in patenting across locations (since this is how knowledge stocks are inferred, see Equation 26). For this reason, because we infer technological waves from patenting behavior, our model is silent on what are the potential causes driving them. For example, some technological waves may be driven by scientific breakthroughs in particular fields that spur innovation. Other factors can also drive our technological waves. One prominent example is firms forecasting demand to rise (or fall) in the future and then directing their innovation efforts to certain sectors and away from others (as suggested by Comin et al. (2019) when studying long-term patterns of sectoral productivity growth in the United States). More broadly, any expected common sectoral shock that re-directs innovation across sectors is going to be picked up as a technological wave by our identification strategy. Even though our model is silent on the origins of technological waves, it allows us to trace their effects over time across locations and sectors.

It is also important to emphasize that we do not make any assumption on the nature and properties of the structural residuals, including whether they are stochastic or deterministic, whether they are spatially or temporally correlated, and whether they are systematically correlated with the other terms in Equation (8). As we have pointed out in Section 3, these structural residuals capture a host of factors that we are not directly modeling, such as the opening of production facilities, universities, and research centers. Importantly, *both* contemporaneous demand shocks (e.g., due to trade) or supply shocks outside of the model (e.g., changes in transportation costs across pairs of locations) are potentially captured in reduced-form by the structural residuals. Thus, these residual terms capture both elements that are related to innovation and elements that are orthogonal to them. For this reason, to quantifying our mechanism, we perform counterfactuals in which we keep these structural terms constant.⁴⁵

We find substantial heterogeneity in the technological waves that we uncover. In the late interval, we find that the largest technological wave shocks (measured as percent changes, as suggested by our theory, $\hat{\alpha}_{s,t}$) occur in the "Health" and "Physics" classes, with an increase of around 20% in each over the 1970 to 2010 period, and the "Electricity" class, with an increase of around 10%. By contrast, innovations related to "Engineering in General" and "Lightning and Heating" register substantial declines of around 15%.

⁴⁵We leave for future work explicit modeling of different elements that are now subsumed in the structural residuals. A first step in this direction is the extension presented in Appendix F, in which we enrich the baseline model to incorporate elements of standard quantitative economic geography frameworks.

4.5 Calibration of the early interval (1910-1950)

To calibrate the model in the early interval (1910-1950), we start from the estimates of the time-invariant parameters that we have obtained (σ , θ , ζ , and $d_{(m,r)\rightarrow(n,s)}$). We then apply the same three-step procedure to recover the corresponding values of local amenities, sectoral productivities, technological waves, and structural residuals.

In the early interval, we find technological waves of similar magnitude to that of such waves in the late interval, but they are concentrated in different fields of knowledge. The most negative technological shocks occur in the "Agriculture and Foodstuffs" and "Building, Drilling, and Mining" classes, which are directly associated with two of the largest sectors of the U.S. economy in 1910; the largest positive technological wave shocks occur in "Electricity" and "Chemistry" classes.⁴⁶

5 Quantitative results

We now use the calibrated model to quantify the role of technological waves, interacted with the endogenous mechanism of knowledge creation and diffusion, in explaining the variation in population growth across U.S. cities throughout the last century. We also study how local diversification mediates the impact of technological waves, and speculate on the future evolution of the United States' economic geography under different plausible scenarios of technological trends.

We report two sets of results, one for the early interval (1910-1950) and one for the late interval (1970-2010). We show that across the two intervals, the model yields comparable results in terms of its ability to reproduce the empirical patterns and its implications for the role of technological waves and frictional knowledge diffusion.

In our baseline experiments, we assume that the economy is in BGP in the first period of each interval, and simulate the model forward under different assumptions about the evolution of the exogenous shocks and the nature of knowledge flows. In Appendix E, we then show that the calibrated model dynamics exhibit substantial inertia. Despite this fact, we show that once we properly normalize outcomes to account for the model's inertial dynamics, the quantitative findings on the role of technological waves are very similar regardless of whether we calibrate the model using separate BGPs for the early and late intervals or we perform a single calibration for

⁴⁶The sectors associated with these last two fields of knowledge employ a relatively small share of the population in 1910. As a result, these positive technological waves, although large, have a limited aggregate effect.

the entire century.

Finally, with the calibrated model, we document how local diversification mediates the impact of technological waves and also run counterfactual scenarios under alternative future technological trends.


5.1 City growth in the model and the data

To start, we consider the model with the full set of shocks (technological waves and structural residuals) and show that it captures a substantial part of the overall variation in city growth observed in the data. Figure 7 shows bin-scatter plots of population growth across cities in the data (horizontal axis) and the model (vertical axis) for both intervals. The model reproduces the data closely, with the R^2 of the underlying regressions equal to 0.48 and 0.52 for the early and late intervals, respectively. Since the model abstracts from other sources of variation in city growth, such as time-varying amenities and endogenous housing supply, population growth in Figure 7 only reflects the evolution of local sectoral productivity, $\lambda_{n,s,t}$, as inferred by the mapping in Equation (26). Figure 7 illustrates that even in the absence of time-varying amenities, the full model goes a long way in accounting for the empirical variation in population growth across cities.

We next consider a counterfactual experiment in which, starting from the BGP, we only feed the path of calibrated technological wave shocks, $\alpha_{s,t}$, while keeping the structural residuals constant at their initial BGP values. In this version of the model, the evolution in local productivity is dictated entirely by the interaction between the initial stock of ideas in each city and the path of technological wave shocks, via the law of motion in Equation (8). Figure 8 shows for both intervals bin-scatter plots of population growth in the data (horizontal axis) and the counterfactual experiment (vertical axis). A comparison between the ranges on the vertical axes in Figures 7 and 8 reveals that the variation in population growth generated by the counterfactual experiments is significantly smaller than the one generated by the full model. The standard deviation across cities declines from 0.28 to 0.07 in the early interval, and it declines from 0.34 to 0.07 in the late interval. This is not surprising, since the counterfactual abstracts not only from time-varying amenities, but also from other determinants of local productivity captured by the structural residuals and the interaction between the structural residuals and technological wave shocks. However, population growth in the counterfactual still displays a strong correlation with the data.

Notably, this counterfactual is successful in reproducing historical movements in relative population growth across prominent U.S. cities. For example, in the early interval, the counterfactual

Figure 7: **Population growth: Data vs. Full Model**



Notes: Bin-scatter plots of the 1910-1950 (left panel) and 1970-2010 (right panel) population growth across CZs in the data (horizontal axis) and the model with the full set of shocks (vertical axis).


records positive growth (relative to the U.S. average) in Detroit (+2.2%), Cleveland (+8.3%), and Pittsburgh (+18.6%), the three main urban centers in the Rust Belt. In the late interval, both Detroit (-8.5%) and Cleveland (-6.3%) receive a negative impact on city population. Pittsburgh experiences a smaller decline (-1.1%) which can be explained by the fact that, in 1970, the city had already planted the initial seed for what would later transform it into a fast-growing center for robotics and artificial intelligence. At the same time, the commuting zones of Austin (+5.7%) and San Jose (+15.6%) experience positive relative growth as an effect of the technological wave. These movements are in line with the observed historical patterns and the narrative that accompanied them (Klepper, 2010; Moretti, 2012).

5.2 City growth, technological waves, and knowledge diffusion

We now provide a quantification of our mechanism by linking it to our motivating empirical exercise. We show that the endogenous mechanism of innovation and knowledge diffusion driven by technological waves only (i.e., keeping the structural residuals at their BGP values) can account for most of the reduced-form empirical relationship between exposure to technological waves and population growth documented in Section 2. We then exploit the structure of our model to decompose this relationship into the contribution from frictions to idea diffusion across cities and across fields of knowledge.

To define exposure to technological waves in the quantitative model, we adapt the empirical measure in Equation (1) to account for patenting growth by class-group throughout each of the

Figure 8: Population growth: Data vs. Model with only technological waves



Notes: Bin-scatter plots of the 1910-1950 (left panel) and 1970-2010 (right panel) population growth across CZs in the data (horizontal axis) and a counterfactual in which we feed the path of technological wave shocks, $\alpha_{s,t}$, and keep the structural residuals constant at their initial BGP values. (vertical axis).

two long intervals. In particular, we define measures of exposure separately for the two intervals $\mathcal{I} \in \{\mathcal{E}, \mathcal{L}\}$, as

$$Exp_{n,\mathcal{I}} \equiv \sum_{s \in S} Share_{n,s,\mathcal{I}} \times g_{s,\mathcal{I}}, \quad (30)$$

where $Share_{n,s,\mathcal{I}}$ is the share of patents filed in commuting zone n belonging to class-group s in the first period of interval \mathcal{I} and where $g_{s,\mathcal{I}}$ is the growth rate in the national share of patents of class-group s between the first and the last period of interval \mathcal{I} .⁴⁷

Column 1 of Table 5 reports regressions of actual population growth on the exposure measure in Equation (30) in the early (Panel A) and late (Panel B) intervals. In line with the reduced-form estimates of Section 2.3, the coefficients are statistically significant and economically large. An increase of one standard deviation in the measure of exposure is associated with an increase in population growth equal to 22.5% of a standard deviation in the early interval and to 20.7% of a standard deviation in the late interval. Column 2 shows the corresponding coefficients in the model with the full set of shocks. The coefficients in columns 1 and 2 are similar, suggesting that the sources of variation in city growth from which the full model abstracts (e.g., time-varying

⁴⁷In Appendix Table A.4 we show results of this section when we define exposure to technological waves using only calibrated model objects, that is, local sectoral shares $\pi_{s|n}^*$ interacted with calibrated technological wave shocks, $\hat{\alpha}_{s,\mathcal{I}}$. Formally:

$$\tilde{Exp}_{n,\mathcal{I}} \equiv \sum_{s \in S} \pi_{s|n}^* \times \hat{\alpha}_{s,\mathcal{I}}. \quad (31)$$

This measure of exposure and the one in Equation (30) yield comparable results in terms of their ability to explain the empirical variation in population growth across cities in the early and late intervals, and in terms of the decomposition of the role of frictions to diffusion in the geographical and technological spaces.

amenities) do not correlate systematically with cities' exposure to technological waves.

Column 3 of Table 5 reports the corresponding estimates for the counterfactual in which we only feed the path of technological wave shocks, and keep structural residuals constant at their BGP values. In the absence of shocks to the structural residuals, Proposition 3 implies that city growth is strongly correlated with exposure to technological waves, as suggested by the R^2 of the regressions in Table 3. Crucially, in both intervals, the slope of the relationship stays remarkably stable between columns 1 and 3 (in fact, coefficients are statistically indistinguishable), even though this slope is not targeted in the calibration. The coefficients imply that an increase of one standard deviation in the measure of exposure is associated with an increase in population growth equal to 15.8% of a standard deviation in the early interval and to 20.0% of a standard deviation in the late interval. This implies that the endogenous mechanism of innovation and frictional knowledge diffusion accounts for most of the empirical relationship between exposure to the technological wave and city growth.

5.2.1 Separating the effect of diffusion frictions across cities and fields of knowledge

Our theory embeds two separate channels behind the impact of a city's exposure to technological waves on local population growth, as summarized by Equation (19). First, frictions to knowledge diffusion *across fields* imply that productivity will increase more in sectors that receive favorable technological wave shocks. As a result, cities in which expanding fields are more prominent will experience higher productivity and population growth. This channel is emphasized by Equation (20), which is derived under the approximation $\eta_{s \rightarrow (n,s)}^* \approx 1$ (Assumption A2). Second, frictions to knowledge diffusion *across locations* imply that cities where expanding fields are more prominent will experience higher productivity growth in *all* sectors because of localized knowledge spillovers between fields.

To decompose the relative importance of these two channels, we re-calibrate technological wave shocks and structural residuals under Assumption A2, that is, by imposing that knowledge flows only happen *within* fields. For both intervals, we then run counterfactual experiments in which we predict city growth by feeding the path of (re-calibrated) technological wave shocks while keeping structural residuals constant at their BGP values. Column 4 of Table 5 reports estimates of the resulting relationship between exposure to the technological wave and population growth. The magnitude of the coefficient declines by 50% and 42% for the early and late intervals, respectively. This suggests that around half of the overall impact of technological waves on city growth is driven

Table 5: **Population growth and technological wave shocks**

| | Population growth | | | |
|---|-------------------|-------------------|-------------------|-------------------|
| | Data | Model | | |
| | | Full | T.Waves | Within |
| | (1) | (2) | (3) | (4) |
| <i>Panel A: Early interval, 1910-1950</i> | | | | |
| Exposure to tech. wave | .897*** (.201) | .946*** (.156) | .631*** (.022) | .316*** (.009) |
| # Obs. | 373 | 373 | 373 | 373 |
| R^2 | 0.05 | 0.09 | 0.69 | 0.75 |
| <i>Panel B: Late interval, 1970-2010</i> | | | | |
| Exposure to tech. wave | .400*** (.098) | .367*** (.107) | .385*** (.009) | .222*** (.006) |
| # Obs. | 373 | 373 | 373 | 373 |
| R^2 | 0.04 | 0.03 | 0.83 | 0.80 |
| Idea flows across fields | - | Yes | Yes | No |
| Structural residuals | - | Yes | No | No |

Notes: OLS estimates. Exposure to the technological wave is defined as $Exp_{n,\mathcal{I}} \equiv \sum_{s \in \mathcal{S}} Share_{n,s,\mathcal{I}} \times g_{s,\mathcal{I}}$ for $\mathcal{I} \in \{\mathcal{E}, \mathcal{L}\}$. The dependent variable is defined as population growth in the data (column 1), in the full model (column 2), in the model with technological wave shocks and constant structural residuals (column 3), and in the model with technological wave shocks, constant structural residuals, and knowledge flows restricted to within-field flows only (column 4). *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

by the existence of localized knowledge flows across fields that amplify the direct impact of shocks to sectoral productivity.

5.3 Diversification and resilience to technological waves

The process of innovation through frictional idea diffusion implies a role for local diversification in making cities resilient to changes in the technological environment. In particular, the same two channels that control local population dynamics in response to technological wave shocks, reflecting respectively frictions to knowledge diffusion across fields and across locations, make the growth trajectory of diversified cities less volatile than the one of specialized cities. First, frictions to

knowledge diffusion *across fields* imply that the path of productivity of any given sector is mainly driven by technological wave shocks to the sector itself. As a consequence, diversified cities, whose sectoral composition is dispersed across multiple sectors, experience a less volatile path of average productivity. Second, frictions to knowledge diffusion *across locations* imply that the reliance of each location-sector on ideas from any given field is an increasing function of the *local* availability of ideas from that field. For this reason, innovators in diversified cities rely on ideas from a broader set of fields, and the local path of productivity is less sensitive to shocks to individual sectors. Overall, the less volatile path for average productivity in more diversified cities also implies a less volatile trajectory in population growth.

Exploring this link formally requires us to define the correct measure of local specialization. Proposition C.1 in the Appendix shows that under Assumption A2 and intuitive conditions on the distribution of shocks, the variance of local population growth is approximately equal to the Euclidean distance between the local and nationwide vectors of sectoral shares. For both intervals $\mathcal{I} \in \{\mathcal{E}, \mathcal{L}\}$, we use this distance as our measure of local specialization. Formally,

$$Spec_{n,\mathcal{I}} \equiv \sum_{s \in S} (\pi_{s|n,\mathcal{I}}^* - \pi_{\cdot,s,\mathcal{I}}^*)^2, \quad (32)$$

where $\pi_{\cdot,s,\mathcal{I}}^*$ is the share of the national population employed in sector s . According to this measure, cities are perfectly diversified if their local sectoral shares are exactly equal to the national ones.

To quantify the effect of diversification on the volatility of city growth, we perform simulations in which we randomly draw shocks $\{\hat{\alpha}_{r,t}\}_{r \in S}$ and compute the corresponding counterfactual equilibria for the economy. We then correlate the standard deviation of population growth across all the simulations with the measure of local specialization in Equation (32).

Columns 1 and 3 of Table 6 report results for the early and late intervals, respectively. The simulations are obtained by keeping structural residuals at their BGP values and randomly drawing technological wave shocks from a normal distribution with mean zero and standard deviation equal to the one of the calibrated shocks.⁴⁸ We run 10,000 simulations for each interval. The results reveal that, as predicted by the theory, specialized cities have higher volatility of population growth across simulations. The magnitude of the coefficients implies that the standard deviation for cities at the 90th percentile of the specialization distribution is 2.7 percentage points higher than for cities at the 10th percentile of the distribution in the early interval and 3.6 percentage points

⁴⁸This standard deviation is equal to 0.126 for the early interval and to 0.116 for the late interval.

Table 6: **Specialization and volatility of population growth**

| | Standard deviation across simulations | | | |
|------------------------------------|---------------------------------------|-------------------|-------------------|-------------------|
| | 1910-1950 | | 1970-2010 | |
| | (1) | (2) | (3) | (4) |
| $Spec_{n,\mathcal{I}}$ | 3.46*** (0.06) | 1.88*** (0.03) | 1.87*** (0.03) | 1.59*** (0.03) |
| # Obs. | 373 | 373 | 373 | 373 |
| R^2 | 0.91 | 0.92 | 0.92 | 0.91 |
| Idea flows across fields | Yes | No | Yes | No |
| Structural residuals | No | No | No | No |
| $Spec_{n,\mathcal{I}}$ 90-10 perc. | 0.008 | 0.008 | 0.019 | 0.019 |

Notes: OLS estimates. Specialization is defined as in Equation (32). The dependent variable is defined as the city-level standard deviation of population growth across 10,000 simulations. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

higher in the late interval.

To disentangle the importance of the first channel (frictions across fields) and the second channel (frictions across locations) in explaining this relationship, we run the same experiment restricting knowledge diffusion to within-field flows only. Results for the two intervals are reported in columns 2 and 4 of Table 6. The coefficient on the measure of specialization drops by about 46% in the early interval and 15% in the late interval. This implies that the direct effect of technological wave shocks interacting with the local sectoral shares explains most of the impact of specialization on local volatility. However, a smaller but still significant portion of the impact is accounted for by the channel of localized knowledge flows across fields, which further attenuate the fluctuations in productivity growth in more diversified cities.

5.4 Impact of future technological waves on the U.S. geography

The quantitative model can be used to predict the evolution of the United States' economic geography in the coming decades in response to transformations in the technological environment. In this section, we propose plausible scenarios for future technological waves and look at which commuting zones will be most positively and negatively affected by those changes. In particular, we project population growth across cities until 2050 under different assumptions about the evolution

of the importance of different sectors ($\alpha_{s,t}$), and compare the outcome with a baseline in which the importance of all sectors is kept constant at their 2010 values.


In the first scenario, we assume that class-group B2 (“Transporting”) experiences a technological wave shock of magnitude equal to twice the standard deviation of technological wave shocks throughout the late interval (+23.2%). This scenario could emerge as new advances in transit technologies and autonomous vehicles induce innovation in transportation to return to a pivotal role. The left map in Figure 9 visually illustrates the results. Commuting zones in blue (red) experience a net gain (loss) of population compared to the baseline. Cities in the Rust Belt are the areas that are best positioned to take advantage from this transformation. According to our experiment, Detroit would experience a 10.0% increase in population relative to the baseline. Other centers of manufacturing related to (but not specialized in) transportation, would also benefit, albeit to a lesser extent. For example, population in Cleveland and Gary would increase by 0.7% and 5.2%, respectively. A relative loss of population would be experienced by the three knowledge hubs of Austin (-4.8%), San Jose (-5.7%), and Seattle (-2.8%).

An alternative way of modeling transportation-related technologies regaining prominence is to assume that ideas from B2 (“Transporting”) become more relevant for innovation in either G1 (“Physics”) or H1 (“Electricity”) and vice versa. An example of the increasing interdependence of those sectors is the gradual integration of IT components in electric and autonomous vehicles. We model this strengthening connection as a drop in the cost of knowledge transmission ($\delta_{s \rightarrow r}$) by assuming a 20% decline in composite knowledge frictions ($d_{(m,r) \rightarrow (n,s)}^\theta$) from (to) B2 to (from) both G1 and H1.⁴⁹ In this case, we keep sectoral importance ($\alpha_{s,t}$) at its 2010 value. The right map in Figure 9 displays the results. Also in this case, Detroit (+4.2%) gains population, while the knowledge hubs of Austin (-0.6%) and San Jose (-1.1%) experience a net loss of population. The reason is that while the economy of Detroit has, to some extent, recently diversified towards fields G1 and H1, Austin and San Jose have increasingly specialized, preventing them from leveraging cross-field spillovers. That said, Seattle (+1.4%) experiences a relative gain in population, leveraging its more diversified base in IT and transportation.

In the second scenario, we simulate a large positive technological wave shock to class-group A3 (“Health; Life-Saving; Amusement,” which includes the bulk of innovation related to pharmaceuticals and medical sciences) possibly in response to new challenges in global health such as the

⁴⁹While the assumption of a proportional 20% decline is arbitrary, this choice only affects the magnitude of the results but does not alter the qualitative patterns.

Figure 9: **Future scenarios: Autonomous vehicles**



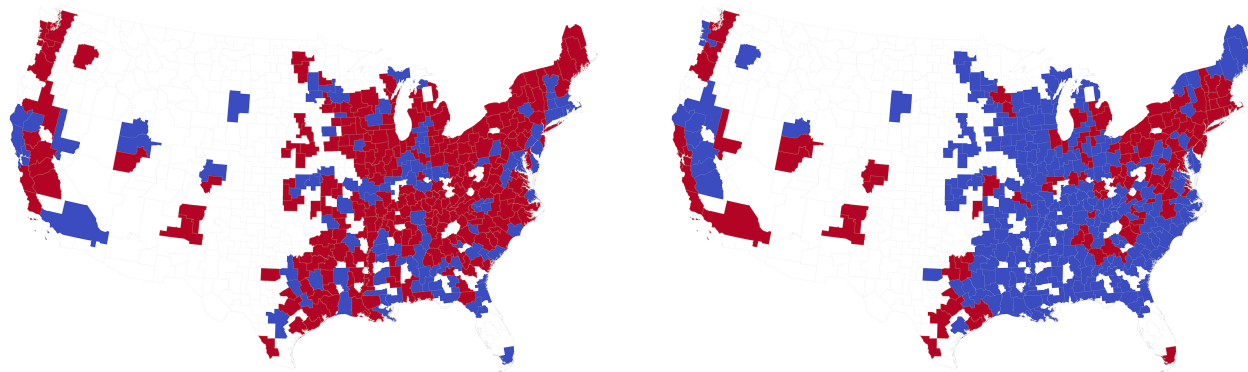
Notes: The maps show log-population in 2050 after technological wave shocks of magnitude +23.2% to B2 (left map), as well as a 20% decline in composite knowledge frictions ($d_{(m,r) \rightarrow (n,s)}^\theta$) from (to) B2 to (from) both G1 and H1 (right map), in deviation from a status quo in which $\alpha_{s,t}$ are kept at their 2010 values. Blue CZs correspond to a net population gain, while red CZs correspond to a net population loss.

COVID-19 pandemic.⁵⁰ The results are depicted in the left map of Figure 10. The experiment suggests that major commuting zones in the North-East, such as Boston (+5.4%) and Providence (+15.0%), and in California, such as Los Angeles (+4.2%) and San Francisco-Oakland (+2.2%), would experience a net inflow of population at the expense of IT clusters such of Austin (-7.6%), San Jose (-3.5%), and Seattle (-3.9%).

In the third scenario, we assume that class-group A1 (“Agriculture”) regains centrality by experiencing an analogous large technological wave shock. This scenario could emerge as a result of tightening regulatory constraints and shifting demand towards sustainable farming, possibly in response to global challenges such as climate change. Results are in the right map of Figure 10. Under this scenario, the economic geography of the United States experiences a pronounced shift away from the East and West coast and the Rust Belt, toward the Central States. Among the major commuting zones, Des Moines (IA) receives the highest net gain (+18.3%). This scenario would represent a significant convergence force in relative population across commuting zones. A regression of log-population in 2010 on the predicted growth rate between 2010 and 2050 delivers a coefficient of -1.1%, implying that population would mostly relocate away from larger commuting zones and towards less-populated ones.

⁵⁰Also in the second and third scenario, we input a shock of magnitude 23.2%, equal to twice the standard deviation of technological wave shocks throughout the late interval.

Figure 10: **Future scenarios: Pharmaceuticals and Agriculture**



Notes: The maps show log-population in 2050 after technological wave shocks of magnitude +23.2% to A3 (left map) and to A1 (right map) in deviation from a status quo in which $\alpha_{s,t}$ are kept at their 2010 values. Blue CZs correspond to a net population gain, while red CZs correspond to a net population loss.

6 Conclusions

The economic geography of countries is characterized by rich and heterogeneous dynamics, alternating between periods of relative stability and periods marked by reversal of fortunes. Some cities remain large and important throughout long time spans, while others experience episodes of sharp growth and decline. In this paper, we explore and quantify the hypothesis that these rich dynamics are partly driven from cities' patterns of specialization across fields of knowledge, coupled with frictions to idea diffusion and a continuously evolving technological landscape.

We develop a parsimonious framework that combines elements from quantitative spatial equilibrium models and theories of endogenous growth through innovation and idea diffusion. The model remains tractable for any arbitrary number of sectors, locations, and time periods—which makes it suitable for quantitative analysis—and delivers a wide range of predictions on how the economic geography of countries responds to changes in the technological environment. The quantitative results support the idea that the interaction of frictional knowledge diffusion with technological waves played a substantial role in shaping the evolution of the U.S. economic geography in the last century. We use the model to speculate on future transformations of the U.S. economic geography under different technological scenarios, such as a comeback of transportation and agriculture and a further rise in the centrality of medical sciences.

Given the novelty of the framework, we have chosen to use a parsimonious model to obtain a sharp characterization of the mechanism. However, the current model can be embedded without

a prohibitive loss of tractability into “standard” quantitative spatial equilibrium models. For example, we show in Appendix F how our model can be extended to include trade and migration frictions across locations, overlapping generations, local externalities in productivity, and endogenous residential amenities. A model extended along those dimensions can be used for policy analysis. For example, our results emphasize a novel channel through which policies that shape the local degree of diversification can affect cities’ long-run success or decline, and an extended model could be used to evaluate the effect of those policies on welfare.

Finally, our quantitative results also suggest that residual factors contribute significantly to the dynamics of local innovation and to the variation in city growth. The framework allows us to isolate the direct effect of technological waves via innovation and knowledge diffusion, and does not require us to make specific assumptions about the nature of this residual. A possible way to endogenize this residual term is to allow innovators to exert effort to improve their ideas, in the spirit of an endogenous growth theory with expanding varieties (as in Jones, 2005). An alternative route to unpack the residual term is to account for the granularity of the location choices of individual firms. Events such as Microsoft’s relocation to the Seattle area or Amazon’s selection of a site for its second headquarters, can have a major impact in shaping the destiny of cities. In addition to the mechanisms described so far, other endogenous forces that enter the residual term include congestion, pecuniary externalities on local assets, and the response of policy to local shocks. Understanding how these factors contribute to amplifying or dampening the effects of technological waves is the next step in our research agenda.

References

- ACEMOGLU, D. (2009): *Introduction to modern economic growth*, Princeton, NJ [u.a.]: Princeton Univ. Press.
- ADAO, R., C. ARKOLAKIS, AND F. ESPOSITO (2020): “General Equilibrium Effects in Space: Theory and Measurement,” Tech. rep., Department of Economics, Tufts University.
- AHLFELDT, G. M., S. J. REDDING, D. M. STURM, AND N. WOLF (2015): “The economics of density: Evidence from the Berlin Wall,” *Econometrica*, 83, 2127–2189.
- AKCIGIT, U. AND W. R. KERR (2018): “Growth through heterogeneous innovations,” *Journal of Political Economy*, 126, 1374–1443.

- ALLEN, T. AND C. ARKOLAKIS (2014): “Trade and the Topography of the Spatial Economy,” *The Quarterly Journal of Economics*, 129, 1085–1140.
- ALLEN, T. AND D. DONALDSON (2018): “Geography and path dependence,” Tech. rep.
- ANDREWS, M. J. (2021): “Historical patent data: A practitioner’s guide,” *Journal of Economics & Management Strategy*, 30, 368–397.
- ANZOATEGUI, D., D. COMIN, M. GERTLER, AND J. MARTINEZ (2019): “Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence,” *American Economic Journal: Macroeconomics*, 11, 67–110.
- AUDRETSCH, D. B. AND M. P. FELDMAN (1996): “R&D spillovers and the geography of innovation and production,” *The American economic review*, 86, 630–640.
- BABINA, T., A. BERNSTEIN, AND F. MEZZANOTTI (2020): “Crisis Innovation,” *Available at SSRN*.
- BALLAND, P.-A., D. RIGBY, AND R. BOSCHMA (2015): “The technological resilience of US cities,” *Cambridge Journal of Regions, Economy and Society*, 8, 167–184.
- BARTIK, T. J. (1991): *Who benefits from state and local economic development policies?*, WE Upjohn Institute for Employment Research.
- BERKES, E. (2018): “Comprehensive Universe of U.S. Patents (CUSP): Data and Facts,” mimeo, available at: <https://sites.google.com/view/enricoberkes/work-in-progress>.
- BERKES, E. AND P. NENCKA (2020): “Knowledge Access: The Effects of Carnegie Libraries on Innovation,” Available at SSRN: <https://ssrn.com/abstract=3629299>.
- BLEAKLEY, H. AND J. LIN (2012): “Portage and path dependence,” *The quarterly journal of economics*, 127, 587–644.
- BOSTIC, R. W., J. S. GANS, AND S. STERN (1997): “Urban productivity and factor growth in the late nineteenth century,” *Journal of urban economics*, 41, 38–55.
- BREZIS, E. S. AND P. R. KRUGMAN (1997): “Technology and the life cycle of cities,” *Journal of Economic Growth*, 2, 369–383.
- BUERA, F. J. AND R. E. LUCAS (2018): “Idea flows and economic growth,” *Annual Review of Economics*, 10, 315–345.
- BUERA, F. J. AND E. OBERFIELD (2020): “The global diffusion of ideas,” *Econometrica*, 88, 83–114.

- CABALLERO, R. J. AND A. B. JAFFE (1993): “How High Are the Giants’ Shoulders: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth,” in *NBER Macroeconomics Annual 1993, Volume 8*, National Bureau of Economic Research, Inc, NBER Chapters, 15–86.
- CAI, J. AND N. LI (2019): “Growth through inter-sectoral knowledge linkages,” *The Review of Economic Studies*, 86, 1827–1866.
- CALIENDO, L., F. PARRO, E. ROSSI-HANSBERG, AND P.-D. SARTE (2018): “The impact of regional and sectoral productivity changes on the US economy,” *The Review of economic studies*, 85, 2042–2096.
- CLEMENS, J. AND P. ROGERS (2020): “Demand Shocks, Procurement Policies, and the Nature of Medical Innovation: Evidence from Wartime Prosthetic Device Patents,” Tech. rep., National Bureau of Economic Research.
- COMIN, D., D. LASHKARI, AND M. MESTIERI (2019): “Structural Change in Innovation,” Tech. rep.
- DAVIS, D. R. AND J. I. DINGEL (2019): “A spatial knowledge economy,” *American Economic Review*, 109, 153–70.
- DAVIS, D. R. AND D. E. WEINSTEIN (2002): “Bones, bombs, and break points: the geography of economic activity,” *American Economic Review*, 92, 1269–1289.
- DE LA CROIX, D., M. DOEPKE, AND J. MOKYR (2018): “Clans, guilds, and markets: Apprenticeship institutions and growth in the preindustrial economy,” *The Quarterly Journal of Economics*, 133, 1–70.
- DESMET, K., R. E. KOPP, S. A. KULP, D. K. NAGY, M. OPPENHEIMER, E. ROSSI-HANSBERG, AND B. H. STRAUSS (2018a): “Evaluating the economic cost of coastal flooding,” Tech. rep., National Bureau of Economic Research.
- DESMET, K., D. K. NAGY, AND E. ROSSI-HANSBERG (2018b): “The geography of development,” *Journal of Political Economy*, 126, 903–983.
- DESMET, K. AND J. RAPPAPORT (2017): “The settlement of the United States, 1800–2000: the long transition towards Gibrats law,” *Journal of Urban Economics*, 98, 50–68.
- DESMET, K. AND E. ROSSI-HANSBERG (2014): “Spatial development,” *American Economic Review*, 104, 1211–43.
- DURANTON, G. (2007): “Urban evolutions: The fast, the slow, and the still,” *American Economic Review*, 97, 197–221.

- DURANTON, G. AND D. PUGA (2001): “Nursery cities: Urban diversity, process innovation, and the life cycle of products,” *American Economic Review*, 91, 1454–1477.
- ECKERT, F. AND M. PETERS (2019): “Spatial Structural Change,” mimeo, available at: <https://mipeters.weebly.com/research.html>.
- GLAESER, E. L. (2005): “Reinventing Boston: 1630–2003,” *Journal of Economic Geography*, 5, 119–153.
- GLAESER, E. L. AND J. D. GOTTLIEB (2009): “The wealth of cities: Agglomeration economies and spatial equilibrium in the United States,” *Journal of economic literature*, 47, 983–1028.
- GLAESER, E. L., H. D. KALLAL, J. A. SCHEINKMAN, AND A. SHLEIFER (1992): “Growth in cities,” *Journal of political economy*, 100, 1126–1152.
- GLAESER, E. L. AND A. SAIZ (2003): “The rise of the skilled city,” .
- GREENSTONE, M., R. HORNBECK, AND E. MORETTI (2010): “Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings,” *Journal of Political Economy*, 118, 536–598.
- HALL, B., A. JAFFE, AND M. TRAJTENBERG (2001): “The NBER patent citation data file: Lessons, insights and methodological tools,” *NBER working paper 8498*.
- HEBLICH, S., S. J. REDDING, AND D. M. STURM (2020): “The making of the modern metropolis: evidence from London,” *The Quarterly Journal of Economics*, 135, 2059–2133.
- HOLMES, T. J. AND J. J. STEVENS (2004): “Spatial distribution of economic activities in North America,” in *Handbook of regional and urban economics*, Elsevier, vol. 4, 2797–2843.
- HORNBECK, R. AND E. MORETTI (2018): “Who benefits from productivity growth? Direct and indirect effects of local TFP growth on wages, rents, and inequality,” Tech. rep., National Bureau of Economic Research.
- HUANG, J. AND Y. ZENOU (2020): “Key Sectors in Endogenous Growth,” .
- JACOBS, J. (1969): *The economy of cities*, Vintage international, Random House.
- JAFFE, A. B., M. TRAJTENBERG, AND R. HENDERSON (1993): “Geographic localization of knowledge spillovers as evidenced by patent citations,” *the Quarterly journal of Economics*, 108, 577–598.
- JONES, C. I. (2005): “Growth and ideas,” in *Handbook of economic growth*, Elsevier, vol. 1, 1063–1111.
- KELLY, B., D. PAPANIKOLAOU, A. SERU, AND M. TADDY (2021): “Measuring Technological Innovation over the Long Run,” *American Economic Review: Insights*, 3, 303–20.

- KERR, W. R. AND S. D. KOMINERS (2015): “Agglomerative forces and cluster shapes,” *Review of Economics and Statistics*, 97, 877–899.
- KLEINMAN, B., E. LIU, AND S. J. REDDING (2021): “Sufficient Statistics for Dynamic Spatial Economics,” .
- KLEPPER, S. (2010): “The origin and growth of industry clusters: The making of Silicon Valley and Detroit,” *Journal of Urban Economics*, 67, 15–32.
- KLINE, P. AND E. MORETTI (2014): “Local economic development, agglomeration economies, and the big push: 100 years of evidence from the Tennessee Valley Authority,” *The Quarterly journal of economics*, 129, 275–331.
- KOGAN, L., D. PAPANIKOLAOU, A. SERU, AND N. STOFFMAN (2017): “Technological Innovation, Resource Allocation, and Growth*,” *The Quarterly Journal of Economics*, 132, 665–712.
- KORTUM, S. S. (1997): “Research, Patenting, and Technological Change,” *Econometrica*, 65, 1389–1420.
- KRUEGER, D. AND F. PERRI (2006): “Does income inequality lead to consumption inequality? Evidence and theory,” *The Review of Economic Studies*, 73, 163–193.
- LIND, N. AND N. RAMONDO (2019): “The Economics of Innovation, Knowledge Diffusion, and Globalization,” in *Oxford Research Encyclopedia of Economics and Finance*.
- LUCAS, R. E. AND B. MOLL (2014): “Knowledge growth and the allocation of time,” *Journal of Political Economy*, 122, 1–51.
- MANSON, S., J. SCHROEDER, D. VAN RIPER, T. KUGLER, AND S. RUGGLES (2021): “IPUMS National Historical Geographic Information System: Version 16.0 [dataset]. Minneapolis, MN: IPUMS,” <http://doi.org/10.18128/D050.V16.0>.
- MARSHALL, A. (1890): *The Principles of Economics*, McMaster University Archive for the History of Economic Thought.
- MICHAELS, G., F. RAUCH, AND S. J. REDDING (2012): “Urbanization and structural transformation,” *The Quarterly Journal of Economics*, 127, 535–586.
- MORETTI, E. (2012): *The new geography of jobs*, Houghton Mifflin Harcourt.
- MURATA, Y., R. NAKAJIMA, R. OKAMOTO, AND R. TAMURA (2014): “Localized Knowledge Spillovers and Patent Citations: A Distance-Based Approach,” *Review of Economics and Statistics*, 96, 967–985.

- NAGY, D. K. (2017): “City Location and Economic Development,” mimeo, available at: <https://sites.google.com/site/davidknagy/research>.
- PERLA, J. AND C. TONETTI (2014): “Equilibrium imitation and growth,” *Journal of Political Economy*, 122, 52–76.
- PETERS, M. (2019): “Market Size and Spatial Growth—Evidence from Germanys Post-War Population Expulsions,” *Unpublished manuscript*.
- REDDING, S. J. AND E. ROSSI-HANSBERG (2017): “Quantitative spatial economics,” *Annual Review of Economics*, 9, 21–58.
- RUGGLES, S., S. FLOOD, S. FOSTER, R. GOEKEN, J. PACAS, M. SCHOUWEILER, AND M. SOBEK (2021): “IPUMS USA: Version 11.0 [dataset]. Minneapolis, MN: IPUMS, 2021,” <https://doi.org/10.18128/D010.V11.0>.
- SCHMOOKLER, J. (1966): *Invention and Economic Growth*, Cambridge: Harvard University Press.
- SIMON, C. J. AND C. NARDINELLI (2002): “Human capital and the rise of American cities, 1900–1990,” *Regional Science and Urban Economics*, 32, 59–96.