

Innovation Networks in the Industrial Revolution *

Lukas Rosenberger W. Walker Hanlon Carl Hallmann
LMU Munich, CESifo Northwestern, NBER, CEPR Jane Street

April 2026

Abstract

We offer a new approach to understanding technological progress and economic growth during the Industrial Revolution using modern growth theory combined with detailed micro-data. We develop a multi-sector endogenous growth model that we use to study one leading theory of British advantage during the Industrial Revolution: knowledge access costs. We apply our model to patent data from Britain and France in order to estimate key parameters, including vectors of country and technology-specific knowledge access parameters. We validate our estimates and then use the model to study their implications during the transition to modern economic growth. We show that, relative to France, British inventors faced lower knowledge access barriers, a difference that generated meaningful growth rate differences on the transition path, even when accounting for cross-country knowledge and technology flows.

*We thank Isabelle Hauf-Pisoni for outstanding research assistance. We thank Enrico Berkes, Davide Cantoni, David de la Croix, Quoc-Ahn Do, Martin Fiszbein, Carola Frydman, Leander Heldring, Rick Hornbeck, Naomi Lamoreaux, Joel Mokyr, Bang Nguyen, Alessandro Nuvolari, Sebastian Ottinger, Monika Schnitzer, Uwe Sunde, Claudia Steinwender, Alex Trew, Fabian Waldinger, and seminar participants at Northwestern, NBER Summer Institute, Chicago Fed, Mannheim, LEAP, UW Milwaukee, Peking University, “Productivity Revolutions” workshop Manchester, EEA Barcelona, EHES Vienna, Louvain, EHS Newcastle, Bayreuth, and LMU Munich for helpful comments. This paper subsumes an earlier draft titled “Why Britain? The Right Place (in the Technology Space) at the Right Time.” Rosenberger gratefully acknowledges financial support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) Project 491578970. Contact: Hallmann, carlhallmann2022@u.northwestern.edu; Hanlon, whanlon@northwestern.edu; Rosenberger, lukas.rosenberger@econ.lmu.edu.

1 Introduction

One of the enduring questions of the Industrial Revolution is: why was Britain able to achieve more rapid economic growth than other European countries? There is now a substantial list of potential British advantages, including the country's uniquely practical Enlightenment tradition (Mokyr, 2009; Almelhem, Iyigun, Kennedy, and Rubin, 2026), its well-developed apprenticeship systems (Kelly, Mokyr, and Ó Gráda, 2014, 2023), the stable institutions established in the wake of the Glorious Revolution of 1688 (North and Weingast, 1989; Acemoglu, Johnson, and Robinson, 2005), higher wages (Allen, 2009a; Voth, Caprettini, and Trew, 2022), and its advantageous natural resources (Pomeranz, 2000; Fernihough and O'Rourke, 2021).¹

Despite the substantial body of ongoing research on this topic, the debate over the key factors that advantaged Britain during the Industrial Revolution remains largely unsettled. One reason for this is that existing studies rarely provide quantitative estimates of the growth contribution of the various hypothesized advantages. In the absence of quantification, it is difficult to assess how much different factors mattered. In addition, existing explanations have been subject to concerns about *post hoc, propter hoc* logic, because it is difficult to separate factors that happened to be present around the time when the Industrial Revolution occurred from those that actually contributed to its onset (Crafts, 1977, 1995).

In this paper, we offer a new approach to studying technological progress and economic growth during the Industrial Revolution that combines modern growth theory and detailed micro-data. We apply this approach to quantify the contribution one of the leading theories of British advantage during the Industrial Revolution: the idea that lower *knowledge access costs*, due to factors such as a well-functioning apprenticeship system, technical publications, and networks of scientific societies, made it easier for British inventors to emerge and produce useful new technologies (Mokyr, 2002, 2005, 2009).

We begin by developing an endogenous growth model featuring multiple technology types and multiple production sectors. In our model, growth depends on research effort as well as within and cross-technology knowledge spillovers, as in Liu and Ma (2023). Researchers operating in a decentralized market face country and technology-specific *knowledge access costs* in order to obtain the knowledge needed to undertake research in a particular technology type. The knowledge access parameters are the key primitives in our model and determine both the level of research effort in a country and how that effort is distributed across technology types.

The level and distribution of research effort in our model determines what we call a country's *latent growth advantage*. This is the growth advantage that one country would have relative to another if the two countries were operating in isolation. We describe this

¹Other explanations have also been offered. Voigtländer and Voth (2006), for example, emphasize the role of demographic change and capital deepening.

advantage as latent because economies do not operate in isolation, but instead share both ideas and technologies. Reflecting this, our model incorporates two types of cross-country effects. The first is knowledge spillovers through the innovation network, whereby new ideas developed in one country affect innovation in another. The second is technology sharing, where technologies developed in one country are put to use in another, affecting productivity and output.

We apply this theoretical framework to study innovation and economic growth during the early stages of the Industrial Revolution. To do so, we begin with detailed data on patenting in Britain from 1700 to 1849 as well as from France, the natural comparison country, from 1791–1844.² Like modern patent data, these historical patent data cover a large number of inventors and their inventions, providing a rich source of information on innovation during the Industrial Revolution.³

Taking our data to the model requires us to overcome two challenges. First, we need to measure spillovers across technology types in order to construct the innovation network. Previous studies use patent citations data, but those are not available in our historical setting. To overcome this, we introduce a new approach based on the idea that if there are spillovers between two technology types, then inventors working primarily on one type will occasionally file patents in the other. In particular, we measure the extent of spillovers from technology category j to i based on the propensity of inventors patenting in i to have previously patented in j . Since our approach is new, we validate it using modern data and show that it comes very close to replicating the innovation network obtained from patent citation data. Developing and validating this new approach, which allows the study of innovation networks much further back in history than currently possible, is one contribution of our study. A second challenge arises from the fact that, unlike most existing theories (e.g., Liu and Ma (2023)), technologies in our model do not map uniquely into production sectors. This feature makes the model more realistic, by incorporating more “general” technologies that have impacts across a broader set of production sectors. We offer a method that uses information on the industries inventors were associated with, together with the type of technology they produced, to create a mapping from technology types to productive sectors.

Using these inputs, together with the structure of the theory, we are able to estimate key parameters, including the vectors of country and technology-specific knowledge access costs. Obtaining these knowledge access parameters is one primary contribution of our paper. While influential existing work emphasizes the role that knowledge access played in the Industrial Revolution (Mokyr, 2002, 2005, 2009), no existing quantitative estimates of

²Both of these were periods during which the patent systems were largely stable. We end just before the major British patent reform of 1852 and the French patent reform of 1844.

³Of course, not every useful invention was patented, as (Moser, 2012) has shown. To account for this, we also generate results based on data from non-patented inventions exhibited in the 1851 Crystal Palace exhibition. These results confirm the patterns that we obtain with the patent data.

these costs exist. Our estimates reveal clear differences between Britain and France in both the overall level of knowledge access costs and how those costs varied across technology types.

Because our knowledge access parameter estimates are novel and dependent on the structure of our theory, in order to build confidence in them it is useful to validate them against external evidence. Our validation exercise is based on the idea that one important determinant of knowledge access was the availability of printed material related to each technology type in a particular language. To operationalize this idea, we use the patent titles to derive a set of equivalent keywords in English and French that strongly predict patents, in both countries, associated with a particular technology category. With these equivalent keywords, we query the Google Books corpus to find the number of printed sources in English and French related to each technology category during the Industrial Revolution. We then compare our estimated knowledge access parameters to the availability of printed material in each language. We find that technologies where we estimate better knowledge access in the UK relative to France have substantially more printed material in English relative to French published between 1750 and 1850, and vice versa. Thus, salient features of the economies we study appear to be captured by our estimated knowledge access parameters.

To study the growth consequences of the different knowledge access costs we estimate, we also need estimates of the parameter governing cross-country knowledge spillovers (when new inventions in one country contribute to further innovation in the other) and cross-country technology sharing (the adoption by one country of technologies developed in the other). We show how these can be estimated from patent and sectoral production data. Our estimates suggest that cross-country knowledge spillovers were small during the study period, but technology sharing was sizable.

We use the structure of the model and our estimated parameters to study the growth implications of the knowledge access costs that we estimate for Britain and France. Because we are interested in the early stages of the Industrial Revolution, we focus on simulating two economies, one representing Britain and one France, that start from common initial technology conditions but differ in their underlying knowledge access cost vectors. Our simulations suggest that the economy with Britain's knowledge access cost vector has a latent growth advantage of around 1 percentage point during the first two decades of the transition to modern economic growth, falling to around 0.3 percentage points after four decades. However, this latent advantage was partially offset by cross-country knowledge spillovers and, more importantly, technology sharing. Accounting for these, our simulations suggest that Britain enjoyed around a one-half percentage point growth advantage during the first few decades of the Industrial Revolution, falling to about 0.1 percentage points after three decades and then slowly converging toward zero over the following century. Thus, our results suggest that knowledge access costs explain a meaningful fraction, but

not the entirety, of the difference in the growth rate of industrial production in these two economies over the period we study (Britain's advantage is estimated at between 0.6 and 1.2 percentage points in the early 19th century).

Finally, we offer a set of counterfactual exercises that allow us to decompose Britain's advantage into the contribution of (i) differences in the average level of knowledge access costs between the two economies and (ii) differences in how knowledge access costs were distributed across different technology types. We find that Britain had a sizable advantage in terms of the average level of knowledge access costs across all technologies, which more than offset France's advantage from having a larger population and therefore more potential inventors. Added to this, Britain also had an advantageous distribution of knowledge access costs across technology types, which affects growth through two channels in our model. First, because some technologies make a larger contribution to consumption, either because they are used by more industries or by industries representing a larger share of consumer expenditure, and second, because some technologies generate more spillovers through the innovation network. Our results suggest that Britain's advantage was driven mainly by the first of these two channels.

This study contributes to an enormous body of work examining the sources of British advantage during the Industrial Revolution, one of the most important questions in economic history. A number of the seminal papers in this literature are cited above. In addition to those papers, which focus mainly on Britain, our study is also related to a vibrant recent literature aimed at better understanding French growth during this period (Squicciarini and Voigtländer, 2015; Juhász, 2018; Squicciarini, 2020; Juhász, Squicciarini, and Voigtländer, 2024; Ridolfi and Nuvolari, 2021; Nuvolari, Tortorici, and Vasta, 2023) as well as a broader set of work studying historical knowledge access institutions more broadly (Dittmar, 2011; Cantoni and Yuchtman, 2014; De la Croix, Doepke, and Mokyr, 2018; Cinnirella, Hornung, and Koschnick, 2025; Juhász, Sakabe, and Weinstein, 2025; Dittmar and Meisenzahl, 2024). Our paper differs from previous work in this area by offering a new approach that brings together frontier growth theory with detailed micro-data for both Britain and France in order to make quantitative statements about the sources of Britain's growth advantage during the Industrial Revolution. In addition, while a number of researchers, most prominently Joel Mokyr, have emphasized the importance of knowledge access during the Industrial Revolution, we are the first to produce theoretically-grounded economy-wide estimates of those costs, estimates that are likely to be useful in future work.

Our study is also related to existing work using aggregate growth models to study historical processes spanning the Industrial Revolution. This includes work, such as Galor and Weil (2000), Galor (2011), or Cervellati and Sunde (2005), that takes a very long-run view of the Industrial Revolution, as well as studies by Hansen and Prescott (2002) and Peretto (2015). We differ from this line of work in three main ways. First, in contrast to the long-run view taken by Galor and Weil (2000) and Galor (2011), we are more focused

on the period of the Industrial Revolution. Second, we use a more disaggregated growth model which differs substantially from the more aggregate theories used in previous work. Finally, we estimate key parameters of our theory using detailed microdata.

Our study also contributes to work by growth economists on the importance of innovation networks in modern economies. Work in this area includes Acemoglu, Akcigit, and Kerr (2016), Cai and Li (2019), Huang and Zenou (2020), and—most similarly—Liu and Ma (2023). While our framework builds on Liu and Ma (2023), we extend their model in several ways. First, rather than taking researcher allocations as given, we show how they can be rationalized as the product of decentralized decisions in response to a vector of underlying knowledge access parameters. We also show how these cost parameters can be estimated, even in the absence of direct information on research effort (i.e., R&D expenditures). Second, rather than treating our economies in isolation, we incorporate both knowledge spillovers and technology sharing, and show how the model can be used to estimate the parameters governing those forces. Third, rather than treating technologies and industries as identical, we allow technologies to be used across multiple industries, a structure that allows us to better match the model to the data, including by opening of the possibility that some technology types are more “general” than others. Finally, we offer new empirical methods that allow us to measure and study innovation networks further back in history, when standard tools such as systematic patent citations and R&D spending are unavailable. This opens up the possibility of studying the influence of innovation networks in different contexts or over longer periods.

Finally, our work builds on a long line of literature using patent data to examine innovation during the Industrial Revolution and into the nineteenth century. Early papers in this area include Sullivan (1989) and Sullivan (1990). More recent work includes MacLeod, Tann, Andrew, and Stein (2003), Khan and Sokoloff (2004), Moser (2005), Khan (2005), Brunt, Lerner, and Nicholas (2012), Nicholas (2011), Nuvolari and Tartari (2011b), Moser (2012), Bottomley (2014b), Bottomley (2014a), Burton and Nicholas (2017), Khan (2018), Bottomley (2019), Nuvolari, Tartari, and Tranchero (2021), Nuvolari et al. (2023), and Hanlon (2023). Relative to this extensive literature, we are the first to combine patent data from multiple countries with a modern economic growth model, the first to estimate knowledge access costs at the industry-technology level, and the first to quantify their impact on relative growth.

The next section of this paper summarizes key features of the historical context that we study. We then present the theoretical framework that we use, in Section 3, followed by a discussion of our data, in Section 4, and our approach to measuring the innovation network, in Section 5. Section 6 describes and compares the estimated innovation networks. In Section 7, we estimate the allocation of effective research effort for each country and technology. In Section 8 we recover the underlying knowledge access parameters and then validate our estimates using data from Google Books. Section 9 examines the growth implications of our estimates through the lens of the model.

2 Historical context

The classical Industrial Revolution is typically thought of as starting around the 1770s and lasting up to the middle of the nineteenth century. Led by Britain, this was a period in which Western European countries experienced industrialization and the initiation of sustained modern economic growth. There is broad agreement that technological progress played a central role in this process.⁴

During the period we study, invention was done by individuals, typically working alone (only about 10 percent of the patents in our dataset have multiple inventors). Developing a useful new technology required some combination of crafts skills, such as in metal or woodworking, an understanding of the best current technologies and existing production systems, and, in some cases, scientific or mathematical knowledge. Numerous factors could influence potential inventor’s ability to access these necessary skills. The structure of apprenticeship systems influenced the availability of craft skills. Access to written information allowed inventors to keep track of the frontier of knowledge and to access scientific insights. Local production exposed inventors to new problems. Proximity to other inventors, engineers, and scientists facilitated information sharing. Individuals who were geographically or socially isolated, who lacked access to written technical or scientific information, or who were far from centers of production, were at a disadvantage when it came to developing new technologies. Thus, as Mokyr argues,

“what counted for useful knowledge to play a role in generating economic growth was therefore *access costs* . . . the easier the access to existing propositional knowledge and to practices in use, the more likely inventions were to emerge and result in sustained economic growth.”⁵

Differences in knowledge access costs may help explain why some countries were leaders in certain technologies but not in others. From historical evidence, we know that Britain was a leader in many of the characteristic technologies of the first Industrial Revolution: steam engines, spinning machines, metalworking, railroads, etc.⁶ However, Continental inventors were also innovative, particularly those in France, the natural comparison country for Britain during this period.⁷ “There is,” wrote Harris (2017) (p.

⁴See, e.g., Landes (1969), Mokyr (1990), Mokyr (1999), Allen (2009a).

⁵Mokyr (2005), p. 295, 296. Italics in original.

⁶One reflection of this lead was the effort that Continental countries, particularly France, (as well as Americans) put into obtaining British technologies in these sectors (Harris, 2017). Another reflection was the leading position that Britain occupied in many of the industries reliant on these technologies. Broadberry (1994), for example, states that “Few would dissent from the view that in terms of technology, Britain was the leading manufacturing nation during the first half of the nineteenth century.”

⁷Leading comparative studies of Britain during the Industrial Revolution focus on France as the comparison country. See, e.g., Crafts (1977), Crafts (1995), Crafts (1998), Allen (2009b), and Harris (2017). After about 1850, the primary comparison typically shifts to the U.S. and Germany (see, e.g., Broadberry (1994)). Since we focus on the earlier period, we follow existing work in focusing on France.

560), “no suggestion that Frenchmen were inherently less inventive than Britons.” Lagging behind in mechanical technologies, French inventors instead made important advances in papermaking, glassmaking, fine ceramics, and chemical processes (bleaching, soda, dyes, etc.).⁸

In addition to differences in knowledge access costs, many other factors influenced growth rates in Britain and France during the period that we study, some advantaging France and others Britain. For example, the French economy suffered severe disruptions during the French Revolution and, to a lesser degree, the Napoleonic Wars.⁹ While temporarily retarding growth, the disruptions also ushered in important institutional changes, such as the adoption of a patent system and the introduction of the Code Napoleon, which have been shown to have contributed to subsequent economic growth (Acemoglu, Cantoni, Johnson, and Robinson, 2011). The Napoleonic Blockade also cut Britain off from many of its closest trade partners during this period, while providing an umbrella of protection that allowed some French industries to modernize (Juhász, 2018). Within this tumultuous historical period, our aim is to isolate and analyze the contribution of one potentially important mechanism to the observed differences in technological progress and output growth in the two countries.

The net result of these various factors, including the mechanism that we emphasize, was more rapid growth in Britain than in France during the early stages of the Industrial Revolution. The best available estimates of industrial growth suggest that industrial production in Britain grew at between 3.1 and 3.5% during the first half of the nineteenth century (Broadberry, Campbell, Klein, Overton, and van Leeuwen, 2015), while in France industrial production grew at between 1.9 and 2.5% (Crouzet, 1996; Asselain, 2007). As a result, the British industry grew 0.6 to 1.2 percentage points per year faster than the French industry during 1815–1860s.¹⁰ Recent reconstructions of GDP per capita suggest a more modest growth advantage for Britain.¹¹ However, it is important to recognize that because Britain had not begun its fertility transition by 1850, GDP per capita growth was undermined by explosive population growth (e.g., Crafts and Harley, 1992).

⁸See, e.g., Harris (2017), p. 540-543, which discusses all of these important French innovations with the notable exception of papermaking. Also p. 560.

⁹Most actual battles were fought outside of France, and a portion of the funding for them was extracted from conquered territory.

¹⁰We obtain this estimate of 0.6 to 1.2 percentage points by calculating growth rates over the exact same period. Asselain (2007) reports three best estimates for French industrial growth during the Industrial Revolution, summarizing the work by the leading French economic historians: 2.5% during 1815–1850 (Crouzet) and between 1.9% (Lévy-Leboyer) and 2.5% (Toutain) during 1820/24–1860/64 per year. For Britain, we calculated the corresponding industry growth rates using the latest data by Broadberry et al. (2015) as 3.5% during 1815–1850 and 3.1% during 1820/24–1860/64 per year.

¹¹Broadberry et al. (2015) and Ridolfi and Nuvolari (2021) put the GDP per capita growth rates in England at 0.47% per year from 1780–1850 and the comparable rate in France at 0.30% (Ridolfi and Nuvolari, 2021, Table 8, p. 22).

3 Theory

This section presents a theory of growth with innovation networks. Our model builds on the framework of Liu and Ma (2023), which introduces a matrix of spillovers across technology sectors into an endogenous growth framework. We extend their framework in four ways.

First, instead of taking differences in researcher allocations across countries as given, our model rationalizes observed differences as a product of the decentralized choices of researchers facing country- and technology-specific differences in the cost of acquiring the knowledge needed to undertake research in a particular technology.¹² These cost parameters, which we estimate, allow us to connect our results to existing work on the Industrial Revolution emphasizing how factors such as apprenticeship systems and scientific societies facilitated entry into certain research areas.

Second, we decouple technology categories from goods sectors. Rather than mapping technologies 1-to-1 onto goods, in our model goods are produced using a variety of different technologies with parameters governing the importance of each technology type in the production of each good. This structure adds realism and allows for a cleaner mapping from the model to the data.

Third, we consider two types of spillovers across countries: knowledge spillovers through the innovation network that influence innovation rates, and technology sharing where innovations from one country are put to use for production in another.¹³

Fourth, rather than assuming that all research in a particular technology is dominated by a single research firm, as in the decentralized extension of Liu and Ma (2023), we allow for free entry of independent researchers. This makes the model more realistic while also generating new insights. In particular, in our model not every country will necessarily undertake research in every sector in every period, and new research sectors may emerge as growing spillovers through the innovation network make research in an area profitable.

We formulate the model in discrete time to facilitate the mapping onto the data. We begin by describing a closed-economy model, which we later generalize to allow knowledge spillovers and technology sharing across countries.

Preferences and Production The model features a representative consumer with utility at time t that is a function of discounted log consumption Y_s in period t and all future

¹²In their appendix, Liu and Ma (2023) present an extension of their model with decentralized equilibrium allocations, but in their decentralized model the researcher shares allocated to each technology match the consumption shares associated with the products those technologies produce. Thus, their extension cannot rationalized differences in allocations across countries except as a product of differential consumption preferences.

¹³Liu and Ma (2023) incorporate the first of these spillover types, but they do not consider the second.

periods:

$$V_t = \sum_{s=0}^{\infty} e^{-\rho s} \ln Y_{t+s}$$

The aggregate consumption good in period t , Y_t , is a Cobb-Douglas combination over consumption from sectors $k = 1 \dots K$:

$$Y_t = \prod_{k=1}^K y_{kt}^{\beta_k} \quad (1)$$

where the β_k parameters give the consumption shares for each sector and $\sum_k \beta_k = 1$. We normalize the price of the aggregate consumption good to one.

Production of each final good is done by a group of perfectly competitive firms that use a set of intermediate inputs based on different technology types according to,

$$\ln y_{kt} = \sum_i \alpha_{ik} \int_0^1 \ln u_{ikt}(v) dv \quad (2)$$

where $u_i(v)$ is the amount of intermediate input of variety v produced using technology type i used in production and α_{ik} is a parameter that governs the importance of technology type i intermediates in the production of type k goods that satisfies $\sum_i \alpha_{ik} = 1$. Each intermediate input is produced using technology of quality $q_{it}(v)$ and labor:

$$u_{ikt}(v) = q_{it}(v) l_{ikt}(v) \quad (3)$$

where $l_{ikt}(v)$ is the amount of labor used to produce the intermediate used in industry k .

Innovation For each technology type i , there is a continuum of varieties $v \in [0, 1]$, each with quality level $q_{it}(v)$. We summarize the state of technology i by the aggregate quality index q_{it} , defined as the geometric average across varieties:

$$\ln q_{it} = \int_0^1 \ln q_{it}(v) dv$$

In each period, innovations in each variety arrive according to a Poisson process with rate ϕ_{it} per period, so that the number of innovations in variety v during a period is Poisson-distributed with mean ϕ_{it} . Each innovation raises the quality of the variety by the proportional step λ : if variety v receives N_v innovations in a period, its quality becomes $(1 + \lambda)^{N_v} q_{it}(v)$.

The expected number of innovations per variety, ϕ_{it} , is increasing in the level of research effort in technology type i , r_{it} , and in knowledge spillovers through the innovation network,

χ_{it} , and decreasing in the current technology level q_{it} ¹⁴:

$$\phi_{it} = \begin{cases} \ln r_{it} + \ln \chi_{it} - \ln q_{it} & \text{if } r_{it}\chi_{it} > q_{it} \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

where $\ln \chi_{it} = \sum_{j=1}^I \omega_{ij} \ln q_{jt}$. The knowledge aggregator term χ_{it} incorporates the innovation network $\Omega_{I \times I}$, whose elements ω_{ij} are governing the level of innovation spillovers from j to i with $\sum_{j=1}^I \omega_{ij} = 1$.¹⁵

Since there is a continuum of varieties, the law of large numbers implies that the aggregate quality index evolves deterministically. Each innovation adds $\ln(1 + \lambda)$ to log quality, and the expected number of innovations per variety is ϕ_{it} , so:

$$\ln q_{i,t+1} = \ln q_{it} + \phi_{it} \ln(1 + \lambda) \quad (5)$$

Resource constraints and knowledge access costs In the model, a fixed population of workers \bar{L} can choose between research and production work. Let R_{it} be the amount of labor doing research in technology i . We then have the following resource constraint:

$$\sum_{k=1}^K \sum_{i=1}^I l_{ikt} + \sum_{i=1}^I R_{it} = \bar{L} \quad l_{ikt} = \int_0^1 l_{ikt}(v) dv \quad (6)$$

If an individual chooses to become a researcher, then they must expend a fraction $(1 - \delta_i^c)$ of their time in order to develop the skills and knowledge necessary to undertake research in sector i . These *knowledge access costs* depend on the parameters $\delta_i^c \in (0, 1)$, which we label *knowledge access efficiency* and which vary across countries c and technology categories i . This reflects the idea that different technologies require different types of knowledge and that knowledge may be more accessible in some countries than in others. The different costs of accessing knowledge reflect differences in the educational institutions, apprenticeship systems, or other factors that influence the ability of researchers to access the skills and knowledge that are needed to develop new technologies in sector i , as well as the amount of knowledge needed to undertake such research. After paying the time cost to acquire the needed skills and knowledge, researchers are left with $r_{it} = R_{it}\delta_i^c$ of their time remaining as *effective research effort*.

Decentralized researcher and worker allocation Next, we describe how the market determines the allocation of researchers across sectors. When a new quality improvement is produced, the researcher that produced it receives a one-period monopoly on producing

¹⁴This last condition, which is important for generating non-explosive growth, reflects a “low hanging fruit” argument: as the level of technology rises, it becomes more difficult to produce the next advance.

¹⁵The condition $\sum_{j=1}^I \omega_{ij} = 1$ allows sustained economic growth.

that intermediate variety with that quality level.¹⁶ Any intermediate good or quality level that is not under monopoly is produced by a perfectly competitive set of intermediate good producers.

When a technology is not under monopoly protection, free entry implies that the price of the intermediate good is driven down to the marginal cost of the labor input, w_t . Note also that the Cobb-Douglas structure of the final goods aggregator (Eq. 1) ensures that the revenue expended on each final good is $\beta_k Y_t$. Using these, and solving the optimization problem for the firms that produce good k , we can derive the amount of labor employed by non-monopoly intermediate firms that use technology i to produce inputs for sector k :

$$l_{ikt}^{P.C.}(v) = \frac{\alpha_{ik}\beta_k Y_t}{w_t} = \alpha_{ik}\beta_k \bar{L}$$

using the fact that product market clearing implies $w_t \bar{L} = Y_t$.¹⁷

When a variety is under monopoly protection, limit pricing implies that the monopolist must set a price that is a markup over the perfectly competitive price of the best available alternative technology. The labor allocation to varieties under monopoly protection is,

$$l_{ikt}^{Mon}(v) = \frac{\alpha_{ik}\beta_k Y_t}{(1 + \lambda)w_t}$$

and the per-period profit of the monopoly producer of a variety is,

$$\pi_{ikt}(v) = \frac{\lambda}{1 + \lambda} \alpha_{ik}\beta_k Y_t$$

Free entry into research implies that individuals will choose to become researchers until the return to being a researcher is equal to the wage.¹⁸ A researcher's expected return, ERR_{it} , depends on the fraction of varieties that receive at least one innovation, $1 - e^{-\phi_{it}}$, divided by the number of researchers R_{it} , times the monopoly profit:

$$ERR_{it} = \frac{(1 - e^{-\phi_{it}}) \sum_k \pi_{ikt}(v)}{R_{it}}$$

This implies that the total number of individuals who decide to become researchers in

¹⁶That is, the innovator earns monopoly profits for the period in which the innovation occurs, after which the variety reverts to competitive production. When multiple innovations hit the same variety in a single period, only the latest innovator holds the monopoly.

¹⁷Note that free entry into research implies that the expected returns to researchers is driven down to the wage of production workers, so that total income in the economy is $w_t \bar{L}$.

¹⁸Note that this setup is substantially different than the assumptions used in the decentralized extension of Liu and Ma (2023). In particular, they assume that all researchers working on a particular technology type are hired by a single research that conducts all research in that technology. We believe that our formulation is more realistic, particularly in the setting that we consider. It also delivers substantially different results: under some conditions, there may be no researchers active in particular countries and technologies.

technology i must satisfy,

$$R_{it} = \frac{Y_t}{w_t} \left(1 - \frac{q_{it}}{r_{it}\chi_{it}} \right) \left(\frac{\lambda}{1+\lambda} \right) \sum_k \alpha_{ik}\beta_k \quad (7)$$

subject to $R_{it} \geq 0$. Effective research effort is:

$$r_{it} = \delta_i^c \left(1 - \frac{q_{it}}{r_{it}\chi_{it}} \right) \frac{\lambda}{1+\lambda} \sum_k \alpha_{ik}\beta_k \bar{L}. \quad (8)$$

subject to $r_{it} \geq 0$, where we have used the fact that total income $w_t \bar{L}$ must equal the total value of production Y_t .

Equation 8 can be expressed as a quadratic equation $r_{it}^2 - x \cdot r + x \cdot v = 0$ where $x = \delta_i^c \frac{\lambda}{1+\lambda} \sum_k \alpha_{ik}\beta_k \bar{L}$ and $v = q_{it}/\chi_{it}$. This equation has either zero, one, or two solutions. In cases where there is no solution, no research effort is allocated to the technology. When there are two solutions, only one is stable—the equilibrium with a larger r_{it} value—and we assume that the stable equilibrium is chosen.¹⁹

3.1 Extensions: Knowledge spillovers and technology sharing

We now extend the model to accommodate two different types of cross-country spillovers. First, we incorporate *knowledge spillovers* occurring through the innovation network, i.e., the possibility that lessons learned from research in technology i in one country might facilitate innovation in technology j in another. To allow for this, we modify our base model so that the expected number of innovations per variety in one country depends on knowledge spillovers from another. Using $'$ to denote the foreign country's knowledge stock, our modified equation for ϕ_{it} is:

$$\phi_{it} = \ln r_{it} + \iota \ln \chi_{it} + (1 - \iota) \ln \chi'_{it} - \ln q_{it} \quad (9)$$

where $\ln \chi_{it} = \sum_{j=1}^I \omega_{ij} \ln q_{jt}$, $\ln \chi'_{it} = \sum_{j=1}^I \omega_{ij} \ln q'_{jt}$, and $\iota \in [0, 1]$ is the parameter governing the importance of foreign knowledge spillovers for domestic innovation. Later, we will estimate the ι parameter and study how allowing for cross-country knowledge spillovers influences relative growth rates.

Second, countries may borrow technology from one another for use in production, which we call *technology sharing*. This channel differs from knowledge spillovers in that it impacts production efficiency but not the rate of domestic technology development. To incorporate technology sharing in a tractable way, we modify the production function so that final goods firms use intermediate inputs built with either domestic or foreign

¹⁹It is easy to prove that the higher equilibrium is stable by showing that the returns to research effort are decreasing in r_{it} at that point.

technology:

$$\ln y_{kt} = \sum_i \alpha_{ik} \int_0^1 \eta \ln(q_{it}(v)l_{kit}(v)) + (1 - \eta) \ln(q'_{it}(v)l_{kit}(v)) dv \quad (10)$$

where $q_{it}(v)$ and $q'_{it}(v)$ denote domestic and foreign technology, respectively, and $\eta \in [0.5, 1]$ determines the importance of each technology in the production process.

It is worth noting that, once one allows for cross-country spillovers, both countries must converge to the same growth rate on the balanced growth path. This fact may explain why previous papers have not examined the importance of these channels in detail. However, growth rates may still differ on the transition path, resulting in permanent level differences in output. Later, we will use simulations to examine how allowing for technology sharing influences relative growth rates and output levels in our setting.

4 Data

4.1 Patent data

Our main data for studying innovation are patent data from Britain and France. For Britain, we have about 12,500 patents covering 1700 through 1849, just before the major patent law reform of 1852. For France, we have data from the establishment of modern patent law in 1791 through 1853, though we mainly use a set of about 11,000 patents filed before a major reform took place in 1844. A useful feature for our purposes is that the British and French patent systems during this period were quite similar.²⁰ These striking similarities are unsurprising considering that the success of patents in Britain inspired legislation in France.

The British patent data used in our analysis were digitized from the *Titles of Patents of Invention, Chronologically Arranged* collected by the British Patent Office (BPO). The data cover England and Wales; for ease of exposition, we will refer to them as “British” patents throughout the paper. The data include the patent number and date, the inventor’s name and occupation, the patent title, in many cases the inventor’s address, and information on whether the patented idea originated from abroad.²¹ We add to these data technology clas-

²⁰For example, the systems were similar in terms of what could be patented (new ideas or new applications of existing ideas related to industry, broadly defined); how patents could be used (to exclude others from using the same idea); whether patents underwent an examination (neither country did so); what was required for obtaining a patent (the payment of a fee and the deposition of a technical description); and whether the priority of foreign inventors was recognized (neither country did, but it was noted whether ideas originated from abroad). One notable difference was that inventors could obtain shorter patents at a substantially lower fee in France.

²¹Most inventors were located in Britain, though a small number filed patents from a foreign address. In addition, 1,350 patent–inventor observations were “communicated from abroad.” In these cases, the location and name of the original inventor is unknown.

sifications produced by the BPO. Patents are classified into one or more of 602 technology subcategories that aggregate hierarchically into 147 technology categories.²²

The French patent data used in our analysis were digitized by the French National Patent Institute (INPI). The data include patent number, patent title, inventor name, inventor occupation, inventor address, and additional details such as the type of patent and the patent term. French patents are divided into three main types: patents of invention, the standard format for new inventions; patents of importation for inventions originating abroad; and patents of improvement for modifications of existing patents. Our analysis focuses on the first two types, as they are the categories that represent truly new inventions.²³ The French patent data also include a technology category classification for each patent. Unlike the British classifications, each French patent is classified into just one of 550 distinct technology keywords that aggregate hierarchically into 94 technology categories.²⁴

4.2 Linking inventor's patents

Our approach requires that we identify all of the patents produced by an inventor. However, doing so is not trivial given that neither the British nor the French patent data include inventor identifiers. We linked patents by individual inventors using a time-consuming careful manual linking procedure. Following Hanlon (2023), we form links using all of the available information in the patent data and in some cases additional external biographical information. In the British data, starting from 13,972 patent–inventor observations, our linking procedure identifies 8,980 distinct inventors. In the French data, starting with 14,277 patent–inventor observations based on just over 11,000 patents, this matching procedure identifies around 10,500 distinct inventors.²⁵

4.3 Concordant technology categories

Another critical challenge for our analysis is establishing concordance between the French and British technology categories. The difficulty is that the two nation's patent offices employed structurally different systems of classifying patents into technology categories during our period. To build a mapping between these two different systems, we begin by identifying 1,148 patents that were filed in both countries—and thus classified by both patent offices into their national classification system. Using these bi-national patents,

²²See Hanlon (2023) for additional details about the data.

²³Patents of improvement provided a cheap way to modify an inventor's existing patent, but they did not extend the term of the original patent.

²⁴Another difference between the French and British patent systems is that inventors could choose in France to apply for patents of different lengths: 5, 10, or 15 years. Longer patent terms required higher fees. See Hallmann, Rosenberger, and Yavuz (2021) for additional details about the data.

²⁵The French links are likely to be even more reliable than those in the British data because French inventors were less likely to have common names and many inventors had three, four, or five names.

we construct a list of how British *subcategories* correspond to French *subcategories* (keywords). Aggregating these subcategories into distinct categories, we obtain 123 concordant technology categories into which we can reliably classify patents from both countries. Appendix A provides further details.

4.4 Mapping technologies to industries

One of the novel features of our model is that we decouple technology categories from industries. This reflects the fact that only a few technologies map unambiguously to one industry, while many (e.g. valves) are used widely. To operationalize this, we need a mapping from technology categories to industrial sectors.

We start with the industries used in the 1907 British input-output table constructed by Thomas (1984). This is the earliest detailed input-output table available for either economy that we are aware of.²⁶

The key challenge is how to map these industries to patenting technology categories.²⁷ We introduce a novel approach for mapping industries and technology categories based on the inventors' occupations observed in the patent data. A substantial fraction of the occupations reported by patenting inventors can be unambiguously associated with specific industries. Given that patents are classified into technology categories, we can use unambiguous occupation-industry links to construct a probabilistic mapping from technology categories to industries, which we describe in detail in Appendix B. This mapping provides us with a measure of the importance of different technology categories for each industry, the α_{ik} parameters in the model, while Thomas' data also provides us with the data we need to calculate the β_k parameters.

4.5 Patent quality adjustments

A canonical issue when using patent data is that patents may reflect technologies of widely varying quality levels. To account for this, it is standard to adjust patent counts using some measure of patent quality, such as the payment of renewal fees (Schankerman and Pakes, 1986) or patent citations (Trajtenberg, 1990; Harhoff, Narin, Scherer, and Vopel, 1999). In our setting, we have two different measures of patent quality for the two countries we study.

²⁶The IO matrix by Thomas (1984) contains 41 sectors. We exclude the four service sectors (*Laundry, Public utility, Distributive services, Other services*), aggregate the four chemical industries into one because of difficulties in matching unique occupations (*Chemicals, Soap and candle, Oils and paint, Explosives*), and exclude the *Motor and Cycle* industry because it did not yet exist during our period. Horrel, Humphries, and Weale (1994) provide an input-output matrix for the British economy in 1841 which is much less detailed.

²⁷No such mapping exists for our historical period; even in modern settings, constructing it can be challenging (Griliches, 1990). One challenge is that it is often unclear whether a technology category should be assigned to industries that produce the technology or to industries that use it. Another challenge is that patents in some important technology categories (e.g., "Valves") may be both produced and used by several different industries.

For British patents, we have quality measures provided by Nuvolari and Tartari (2011a) and Nuvolari et al. (2021) that are similar to patent citations. For France, we have the payment of larger fees for longer patent lengths as well as other fees required to keep patents in force. In Appendix F, we describe how we construct a concordance between these two alternative measures in order to obtain a patent quality index that can be applied to patents in both countries. Briefly, this method involves mapping the distribution of British patent quality from Nuvolari et al. (2021) into the five patent quality bins identified by the different levels of renewal fees paid for patents in France. We use this quality-adjusted patent data in all of our main results, though in additional analysis we explore results using patent counts or alternative mappings between the French and British patent quality indicators.

4.6 Exhibitions data

We will also provide some supporting results based an alternative measure of innovation that does not rely on patents. Specifically, following the work of Petra Moser (Moser, 2005, 2012), we examine data covering exhibits in the 1851 Crystal Palace Exhibition. This Exhibition was the first world’s fair, a massive event that included over 17,000 exhibits from 40 countries and attracted over 6 million visitors.²⁸ We use data covering the 6,003 exhibits by British exhibitors and the 1,675 exhibits by French exhibitors. Each entry includes information such as the exhibitor name and address and a brief description of the exhibit. Each exhibit was also classified into one of 29 categories.²⁹ Following the same procedure applied to the patent data, we manually reviewed all exhibits in order to link exhibits by the same exhibitors. Additionally, we manually linked individuals who exhibited at the Crystal Palace to inventors who appear in our patent data.

5 Measuring the innovation network

This section introduces and validates a novel method for measuring the innovation network in settings where no systematic patent citation data are available.

5.1 Standard citation-based method

Standard patent datasets for modern settings come with systematic patent citation data. Modern studies on innovation networks use these citations to construct measures of the

²⁸This exhibition was followed by others, such as the Exposition Universelles in Paris (1855, 1867) or the American Centennial Exhibition in Philadelphia in 1876. We focus only on the Crystal Palace because it is the first international exhibition that occurred in close proximity to the period covered by our patent data.

²⁹The original dataset includes 30 categories, but we omit category 30, which is for fine arts exhibits. For some French exhibits, it also provides secondary or tertiary categorizations for French exhibits. But as most exhibits have only one category, we focus on each exhibit’s first category.

spillovers between different technology categories. Specifically, the strength of spillovers from technology category j to category i is typically measured as $\omega_{ij}^{cite} = Cites_{ij} / \sum_l Cites_{il}$, where $Cites_{ij}$ is the number of patents in category i citing patents in category j (e.g., Acemoglu et al., 2016; Liu and Ma, 2023). The basic assumptions in this approach are (i) that some fraction of the useful ideas generated through research in technology j that increase research productivity in technology i are reflected in citations from i to j and (ii) that this fraction is fairly stable across i - j pairs.

5.2 Novel inventor-based method

In our historical setting, inventors were not required to systematically include citations to prior work in their patent specifications. As a result, systematic citation data is not available. Thus, we need to develop an alternative to the standard citation-based approach.

To overcome this lack of data, we introduce a novel method that follows a similar intuition as the citation-based method but instead exploits information contained in the inventors' patenting sequence. The basic idea is that an inventor may learn lessons while working on research in category j that leads to a subsequent discovery in technology category i . If we found that inventors in category i had often previously patented in j but not another category g , it would signal larger knowledge spillovers from j to i relative to spillovers from g to i . Analogous to the assumptions in the citation-based approach, the key assumptions in our approach are (i) that the sequence of patenting from j to i by inventors with multiple patents is proportional to the number of ideas that are generated by researching j and useful for future research in i , and (ii) that the constant of proportionality is fairly stable across i - j pairs.

Specifically, we define ω_{ij} , the strength of spillovers from category j to i , as


$$\omega_{ij} = \frac{\sum_m P_{mij}}{\sum_m P_{mi}} \quad (11)$$

where P_{mij} is the number of patent sequence pairs by inventor m , with the first patent filed in technology category j and the next patent filed in i , and where P_{mi} is the total number of patents by inventor m in technology category i which pair with an earlier patent (which can be either in i or in another technology category g).³⁰

We can calculate ω_{ij} values using either British patents, French patents, or both. In our main analysis, we will rely on a matrix generated using both sets of patents after they have been mapped to our common set of concordant technology categories. This

³⁰One complication in our setting is that some British patents are categorized into multiple technology categories. To deal with this, we generalize P_{mi} and P_{mij} to be the weighted count of pairs of patents by inventor m , with weights corresponding to the inverse of the number of categories a patent is listed in. This is only one possible solution to the problem; another alternative is to focus on the modal technology category and throw a coin for cases without distinct mode (e.g. Jaffe, Trajtenberg, and Henderson, 1993).

Figure 1: Comparison of citation-based and inventor-based networks



Unconditional scatter plot of log edge weights constructed by the citation-based method against those by the inventor-based method. For data and method, see text.

joint matrix can be constructed in several ways; either will require judgment about what relative weight should be attached to patents from each country. As patents in the two countries are the product of different patent systems and institutional environments, it is not obvious how to determine an objectively correct weighting. Thus, we choose the simple approach: giving each patent system equal weight. Specifically, we simply average the (row-normalized) edge values across Britain and France to obtain the joint matrix.

5.3 Validating the inventor-based method

Whether the inventor-based method provides a useful measure of the innovation network is ultimately an empirical question. To provide some confidence in our new approach, we use modern patent data to compare the innovation network generated by the inventor-based method to the innovation network generated by the standard citation-based method. We use data on U.S. patents from 1970–2014 from PatStat. As described in more detail in Appendix C, we generate a citation-based innovation matrix using a standard approach taken from previous studies. Our inventor-based innovation matrix is obtained using the approach described above. Once we have the two matrices, we can compare either the edge values or the centrality of the nodes in the two matrices.

Figure 1 presents a scatterplot comparing the edges (ω_{ij} terms) obtained from the

standard citation-based method against our inventor-based method. It is evident that the two approaches give very similar results. The corresponding regression of citation-based against inventor-based log edge weights estimates the slope as .994 (95% confidence interval [.984, 1.003]) with an associated R^2 of .79. Thus, our inventor-based approach provides a very close approximation to the network generated using the citation-based approach commonly used in existing studies.

6 Innovation networks in the Industrial Revolution

Having established the validity of our method, we now apply it to British and French patent data in the Industrial Revolution. A first glimpse of the innovation network is shown in Figure 2, which provides a visualization of the innovation network based on the joint matrix constructed from British and French patents. In the figure, each technology category is a node, the size of node reflects the number of patents filed in that category, and the location of the node is determined by the strength of connections between that node and every other node in the network as determined by a multidimensional scaling algorithm.

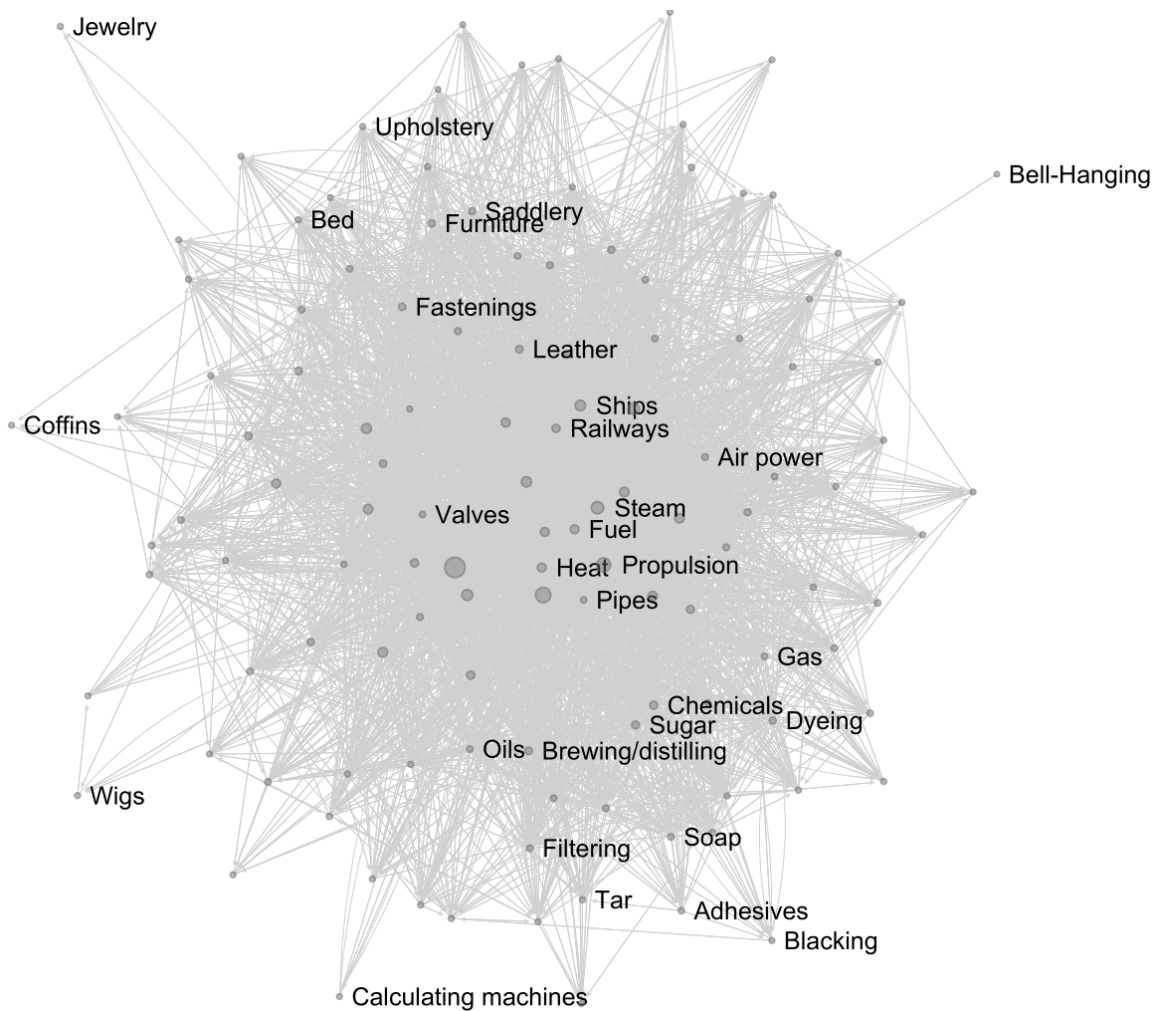
Several interesting patterns stand out in Figure 2. First, the technology space is characterized by a dense central core area. Near the center of the core area, we see categories such as Steam Engines, Railways, Heat, and Propulsion, as well as many smaller technology categories. These core technologies include a number that historians have highlighted as important for the Industrial Revolution (Landes, 1969; Mokyr, 1990; Allen, 2009a), most notably steam engines. We can also see that there are clusters of related technologies. The most visible cluster is the set of chemical and related technologies located toward the bottom of the network, which includes Chemicals, Oils, Dying, Brewing/distilling, Filtering, Soap, Tar, and Adhesives. We can also see, toward the top of the plot, a cluster of similar technologies such as Upholstery, Bedding, and Furniture. Finally, there are a number of very peripheral technologies with few connections to other categories, including Jewelry, Coffins, Wigs, Blacking, and Bell-Hanging.

How similar are the networks generated by the French and British patent data? The answer to this question matters. If the networks are similar it suggests that their structure reflects a single underlying “global” network, as assumed in the theory, rather than being determined by idiosyncratic institutional features or local economic conditions.

To assess the similarity of the British and French networks, we begin with two separate innovation networks, one constructed using only French domestic patents and another constructed using only British domestic patents, but both expressed in terms of the common technology categories. We then apply the following regression specification

$$\omega_{ij}^{UK} = \beta_0 + \beta_1 \omega_{ij}^{FR} + \epsilon_{ij}$$

Figure 2: The joint innovation network visualized



Plot generated using multidimensional scaling. The edges of the joint network are computed as $(\omega_{ij}^{UK} + \omega_{ij}^{FR})/2$. Labels for technology clusters around Steam, Chemicals, and Furniture. Additionally, some peripheral technologies are labeled (Bell-Hanging, Calculating Machines, Wigs, Coffins, Jewelry). See Figure D.1 for a fully labeled version.

Table 1: The edges of French and British networks are similar

	Dep var: British network edge ...							
	Edge indicator		Edge weight		Edge weight		Log edge weight	
	(ext. margin)		(ext.+int. margin)		(intensive margin)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
French network edge	0.332*** (0.014)	0.148*** (0.014)	0.241*** (0.035)	0.228*** (0.036)	0.493*** (0.067)	0.861*** (0.050)	0.889*** (0.050)	1.084*** (0.065)
Receiving node (<i>i</i>) FE		✓		✓		✓		✓
Sending node (<i>j</i>) FE		✓		✓		✓		✓
<i>N</i> (Obs = edges)	15129	15129	15129	15129	646	604	646	604
Estim. FE coef.		245		245		153		153
R^2	0.05	0.28	0.10	0.14	0.41	0.72	0.38	0.59
Within R^2		0.01		0.09		0.66		0.35


Observations are network edges connecting nodes (technology categories) *i* and *j* for networks constructed using only domestic inventors in each country. The full network has $123 \times 123 = 15,129$ observations. In columns (1) through (4), missing edges are replaced by zero. For the extensive margin in columns (1) and (2), edge indicators equal one if the edge is larger than zero. In columns (5) through (8), missing edges are treated as missing. Covering only the intensive margin, this smaller sample only includes edges which are larger than zero in both countries. OLS regressions, robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

where the superscripts indicate edges from either the French or UK innovation matrices. If the networks were identical, then we would estimate $\beta_1 = 1$ with an R^2 of 1. Given that the two matrices represent two different realizations of any underlying innovation network, it is unrealistic to hope that the two matrices will correspond so closely. Nevertheless, evidence of strong similarities between the two matrices would be suggestive of a common underlying network structure, as assumed by the theory.

Table 1 presents the regression results. In the first two columns the dependent variable is an indicator for whether there is any connection from node *j* to *i*. We describe this as the extensive margin. Column 1 presents the simple univariate regression results, while in Column 2 we add in a full set of receiving and sending node fixed effects. These fixed effects deal with the possibility that the shape of our network may be influenced by features of particular nodes. For example, we may worry that nodes with more patents have more connections. The inclusion of both sending and receiving node fixed effects deals with this type of concern.

In Columns 3–4 we instead use the edge weights, $\omega_{ij} \geq 0$. Note that we have a large number of edges, over 15,000, relative to the number of patents in our data. As a result, many edge weights are zero—implying no connection between two nodes—simply because the true spillover is small and we do not have enough observations to measure it with precision. In Columns 5–8 we instead limit our analysis to the set of edges where we

Figure 3: Correspondence of British and French node eigencentrality



Comparison of the left eigenvector centrality of nodes in two innovation networks, one based on British patents (y-axis), the other on French patents (x-axis). Categories at the top right are the most central in both countries. Categories at the bottom left are the least central.

observe some positive connection, $\omega_{ij} > 0$, in both the British and the French realization of the innovation network. This compares the strength of the connections at the intensive margin.

Node centrality An alternative approach to assessing matrix similarity is to focus on the centrality of the network nodes, which provides a useful way to summarize the shape of the network. In the network literature, there exist different metrics for node centrality; which is most appropriate depends on the application. In the type of model we use, Liu and Ma (2023) show that *eigenvector centrality* is the relevant measure.

Figure 3 plots the eigenvector centrality of nodes in the innovation network based on British patent data against the eigenvector centrality of nodes in the network based on French patent data, with both expressed in the concordant technology categories. The clear similarity between the centrality of nodes in the British and French innovation networks provides another reflection of their similarity. Additional results using alternative

measures of node centrality, presented in Appendix Table D.3, confirm these results. Across all centrality measures, we find strong commonalities in the network structure.

Alternative network based on exhibition data One might worry that the structure of the innovation networks above is due to the patent system. To check this, we have constructed an alternative measure of the innovation network not based on patents. Specifically, starting with data from the 1851 Crystal Palace Exhibition, we use exhibitors that exhibit in multiple technology categories to build up a matrix reflecting the connections across different technology types. This method is similar to the approach applied to the patent data (for details, see Appendix E), except that the exhibition data does not have any time dimension because all exhibits appeared in the same exhibition, so we end up with an undirected matrix.

Visual examination shows that the innovation network obtained from the exhibition data, shown in Appendix E, appears similar to the network obtained from the patent data. However, we can examine that similarity more formally by using the links that we constructed between exhibitors and the patent data. In particular, we show that, for individuals who appear as both exhibitors and patentees, those who exhibited in more central nodes within the innovation network based on exhibition data also tended to patent in nodes that were more central within the innovation network based on patent data. This provides direct evidence showing that the innovation network is similar regardless of whether it is based on patent data or on non-patented inventions.

Stability over time Our model assumes that the structure of the innovation network is stable over time. Is this reasonable? We examine this issue in Appendix D.3 and find evidence that the matrix is fairly stable over the time-frames relevant for our study, a result that is also consistent with previous findings in Liu and Ma (2023). Thus, this seems like a reasonable assumption to make over the time-frame relevant for our study, though the matrix is likely to be evolving over longer time periods.

To summarize, across all of these various approaches, we find strong evidence of similarity in the French and British innovation networks as well as evidence that the matrix is reasonably stable over time. This tells us that a substantial fraction of the observed innovation network is likely to be due to some underlying “global” factors, such as the nature of the underlying technologies, that are common to both countries and fairly stable over time, as assumed in our theory.

7 Estimation of effective research effort

The structure of the model, together with the patent data, can be used to recover estimates of effective research effort (unobserved in our setting) from observed research output (quality-adjusted patents). The starting point for this exercise is Eq. 5, which describes the evolution of technology improvements over time. Substituting for ϕ_{it} using Eq. 9, the

growth equation with country subscripts added is:

$$\ln q_{ci,t} - \ln q_{ci,t-1} = \ln(1 + \lambda) \left[\iota \ln \chi_{ci,t-1} + (1 - \iota) \ln \chi'_{ci,t-1} - \ln q_{ci,t-1} + \ln r_{ci,t-1} \right] \quad (12)$$

We map this equation to our data by setting $q_{cit} - q_{cit-1} = Patents_{cit}$ where $Patents_{cit}$ is the number of quality-adjusted patents in country c , technology i , and period t . This implies that $q_{cit-1} = \sum_{\tau=0}^{t-1} Patents_{ci\tau}$. Using this and the ω_{ij} parameters from our innovation matrix we have $\chi_{ci,t-1}$. This leaves us with two “macro” parameters to estimate, λ and ι , as well as the vector of unobserved $\ln r_{cit}$ terms. Since we are interested in average differences in researcher allocations over our study period, we estimate the $\ln(1 + \lambda) \cdot \ln r_{cit}$ terms as a set of technology-category-by-country fixed effects ξ_{ci} . Thus, our estimating equation is:

$$\begin{aligned} \ln \left(\frac{q_{cit}}{q_{cit-1}} \right) &= \ln(1 + \lambda) \iota \ln \chi_{ci,t-1} + \ln(1 + \lambda) (1 - \iota) \ln \chi'_{ci,t-1} - \ln(1 + \lambda) \ln q_{ci,t-1} + \xi_{ci} + \varepsilon_{cit} \\ &= \beta_{\chi} \ln \chi_{ci,t-1} + \beta_{\chi'} \ln \chi'_{ci,t-1} + \beta_q \ln q_{ci,t-1} + \xi_{ci} + \varepsilon_{cit} \end{aligned} \quad (13)$$

where λ , ι , and ξ_{ci} are parameters to be estimated, ε_{cit} is the error term, and everything else is observed data. Since we have three coefficients ($\beta_{\chi}, \beta_{\chi'}, \beta_q$) that jointly identify two parameters (λ, ι), we estimate this equation via constrained LS such that $\beta_{\chi} + \beta'_{\chi} = -\beta_q$. The coefficient on $\ln q_{ci,t-1}$ identifies $\ln(1 + \lambda)$, from which the quality step size can be recovered as $\hat{\lambda} = e^{-\hat{\beta}_q} - 1$. The parameter ι is identified from the ratio of coefficients β_{χ} and $\beta_{\chi'}$. We focus on the period 1800–1843 in the baseline, but also provide robustness for alternative periods.³¹

Table 2 presents our estimates of the main coefficients from Eq. 13 for several different alternative specifications. Column 1 is our baseline specification, where we pool both countries and estimate results over the full 1801–1843 period using quality-adjusted patents. At the bottom of the table, we report the values of the key parameters, λ and ι , implied by our coefficient estimates. It is useful to note that our estimate of λ is similar to the value found in previous work: Acemoglu, Akcigit, Alp, Bloom, and Kerr (2018) estimates a λ of 0.132 using a very similar model (though without cross-country knowledge spillovers), while Liu and Ma (2023) use a value of $\lambda = 0.17$ in their calibration exercise based on their estimates of the relationship between knowledge stocks and patenting activity. Our estimates indicate an ι close to 1, which suggests that cross-country knowledge spillovers were limited—though not negligible—in our setting, likely because of the delay in information flows between the two countries.

In Columns 2 and 3, we estimate the equation separately for the two countries, instead of pooling the data. This allows λ and ι to differ by country. It is reassuring to see that

³¹One decision we have to make when estimating this model is how many lags of patents to include when constructing $q_{ci,t}$. In our main specification we use patents from 1761 onward (though French patents start in 1791), though we examine the robustness of our results to other choices.

Table 2: Main estimation table

	Dep var: $\Delta \ln q_{c,i,t}$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Baseline	Britain	France	Patent counts	Own Ω	q_0^{UK} 1790	Late 1824–43
$\ln \chi_{t-1}$	0.102*** (0.011)	0.109*** (0.025)	0.094*** (0.009)	0.112*** (0.010)	0.093*** (0.011)	0.128*** (0.012)	0.160*** (0.021)
$\ln \chi'_{t-1}$	0.004 (0.002)	-0.010*** (0.003)	0.028*** (0.006)	0.001 (0.002)	0.007*** (0.002)	-0.007** (0.004)	0.006 (0.006)
$\ln q_{t-1}$	-0.106*** (0.011)	-0.099*** (0.023)	-0.122*** (0.012)	-0.113*** (0.010)	-0.100*** (0.011)	-0.121*** (0.012)	-0.166*** (0.023)
N	9325	5085	4240	9329	9163	9214	4698
Num. FE (ξ_i)	240	120	120	240	235	240	240
F -stat	26.8	32.9	29.9	29.7	26.8	23.4	26.5
implied λ	0.112	0.104	0.129	0.120	0.105	0.128	0.180
implied ι	0.963	1.000	0.772	0.988	0.931	1.000	0.965

The table presents estimates of Eq. 13 using constrained LS such that $\beta_\chi + \beta'_\chi = -\beta_q$. The top panel lists estimated coefficients, with robust standard errors in brackets, while the bottom panel gives the implied values for λ and ι . In columns 2 and 6, ι is capped at its upper bound of one. Unconstrained coefficients (not reported) are very similar. All specifications include a full set of technology-category-by-country fixed effects. The first column gives the baseline specification, which focuses on outcomes in the period 1801–43, and uses British quality-adjusted patents from 1761 and French quality-adjusted patents from 1791 together with the joint network to compute knowledge stocks q and spillovers χ . The second and third column does not pool the British or French data but instead estimates the equation separately for each country, allowing λ and ι to differ across countries (but still using the common patent quality adjustment). The fourth column uses raw patent counts instead of quality adjusting patents. The fifth column allows the innovation network to differ by country. The sixth column uses British patents from 1791 instead of 1761. The seventh column restricts the sample to the late period 1824–43 (last 20 years).

we obtain fairly similar results, in particular for λ . Interestingly, we find an ι parameter of 1 for the UK but a lower estimate of 0.77 for France.³² This suggests that knowledge spillovers from the UK to France were meaningful, but not the other way around.

The remaining columns present additional robustness results. In Column 4, we do not adjust for patent quality but instead use raw patent counts for calculating q and χ . The implied λ and ι values are somewhat larger, indicating that quality adjusting yields somewhat more conservative estimates. In Column 5, we allow the countries to have different innovation networks, i.e., networks based only on their own domestic patents. In Column 6, we show robustness to using only patents after 1791, when both patent systems were in operation, to generate our measures of the knowledge stock in each country and technology. Finally, in Column 7, we estimate results only on the later period when the growth take-off was in full swing. This gives the largest estimate for λ , though still within the range of prior studies.

Table 3: Technologies with the highest share of effective research effort

Top British category	$\frac{r_{Bi}}{\sum r_{Bi}}$	Top French category	$\frac{r_{Fi}}{\sum r_{Fi}}$
Steam Engines and Boilers	0.034	Road Vehicles	0.037
Shipbuilding and Operation	0.031	Lighting and Matches	0.033
Power and Propulsion	0.029	Power and Propulsion	0.031
Road Vehicles	0.028	Shipbuilding and Operation	0.027
Metallurgy	0.027	Railways and Locomotives	0.025
Pipes and Tubes	0.023	Heating and Evaporating	0.023
Pumps and Drainage	0.021	Stoves, Ovens, Kilns	0.022
Fabric Manufacture	0.021	Sugar Manufacture	0.021
Fuel	0.02	Fabric Manufacture	0.02
Brewing and Distilling	0.02	Spinning and Preparation	0.02

Estimates of $\hat{r}_{c,i} = \exp(\hat{\xi}_{c,i} / \ln(1 + \hat{\lambda}))$ corresponding the baseline estimates in Table 2, Column 1.

We have also estimated a vector of country-by-technology fixed effects, ξ_{ci} , which we can use to calculate effective research effort according to $\hat{r}_{c,i} = \exp(\hat{\xi}_{c,i} / \ln(1 + \hat{\lambda}))$. To give a sense of these effective research effort estimates, Table 3 presents the ten technology categories accounting for the highest share of effective research effort in each country. Several technologies appear in both countries' top ten (ship building, motive power and propulsion, road vehicles, and fabric manufacture), indicating overlap in the areas on which researchers in both countries focused their efforts. Beyond these similarities, British research effort appears concentrated in mechanical and metal technologies (steam engines, metallurgy, pipes and tubes), while French research effort was more concentrated on consumer goods (lighting and matches, sugar manufacture).

³²The regression coefficient on the foreign spillover term $\ln \chi'_{t-1}$ in Column 2 of Table 2 is small and negative (-0.010), which would imply $\iota \approx 1.1$. Since values above one are not admissible under the model, we cap the parameter at its upper bound of one.

8 Knowledge access costs

In this section, we use our estimates and the structure of the model to back out differences in the costs of accessing knowledge. Specifically, we will recover a set of knowledge access efficiency parameters δ_i^c for each country–technology pair. We then empirically evaluate our estimated parameters.

Inverting Eq. 8, we can express δ_i^c as a function of r_{cit} and observed data and parameters:³³

$$\delta_i^c = \frac{(1 + \lambda) r_{cit}^2}{\lambda \sum_k \alpha_{ik} \beta_k \bar{L}_c (r_{cit} - \frac{q_{cit}}{\chi_{cit}})} \quad (14)$$

This equation holds for any period t . The estimation of Eq. 13 gives us $\hat{\lambda}$ and an average \hat{r}_{ci} over the estimation period (in the baseline 1800–43). To match this, we calculate average \hat{q}_{ci} and $\hat{\chi}_{ci}$ over the same estimation period, using the patent data and the innovation network.³⁴ We set \bar{L}_c equal to total population in each country to their mid-period values; at that time, the ratio of British to French population was close to 0.5.³⁵

This procedure allows us to back out a full set of knowledge access efficiency parameters for both countries. To provide a sense of what these parameters look like, Table 4 presents the ten technology categories in which Britain had the greatest relative advantage on knowledge access (left panel) and the ten where France had the greatest advantage (right panel). We can see that Britain’s advantage was particularly pronounced in several categories related to metalworking and engine and machine parts. Examples include cylinders and pistons, grinding and polishing, and joints and pulleys. Many of these were relatively central within the innovation network. In contrast, France offered easier access to knowledge related to aerial conveyances, several types of consumer and luxury goods (beds and bedding, games and amusements) and food preparation (cooking and baking, preserving and curing).

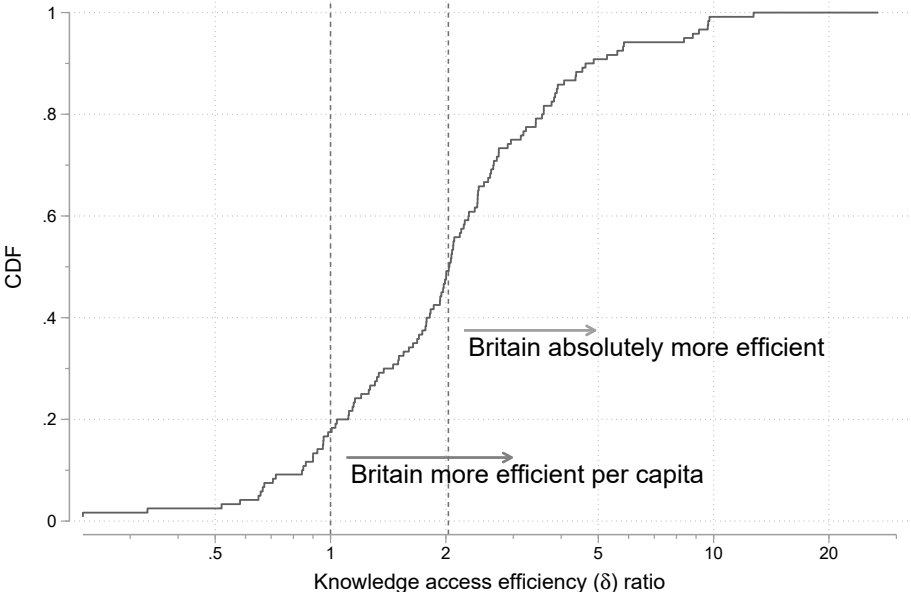
Figure 4 plots the cumulative distribution of relative knowledge access efficiencies, $\delta_i^{UK} / \delta_i^{FR}$. This graphs shows that Britain had lower knowledge access costs for around

³³Note that here we are using the expression derived from the model without technology sharing. In a model with technology sharing as we have introduced it, the expression will be the same but the interpretation of the quantity on the left-hand side will be different. Instead of being equal to δ_i^c , the value obtained from this expression will be equal to $\delta_i^c \eta$. However, this will have no practical effect on any of our results.

³⁴Note that, were the economies on the BGP, these would be the correct values to use alongside \hat{r}_{ci} in calculating δ_i^c . If we are near the BGP, they should be approximately correct. Our estimation results are based on data starting in 1800, roughly 30 years after the onset of the Industrial Revolution. Our simulation results suggest that after three decades the economy should be fairly close to the BGP, so we view this as a reasonable approximation.

³⁵We use population data from O’Brien and Keyder (2011), Table 3.2., which are based on British Historical Statistics and French census data. For the mid-period decade 1825–34, $\bar{L}_{UK} = 16.0\text{m}$ and $\bar{L}_{FR} = 32.6\text{m}$.

Figure 4: Cumulative distribution of knowledge access ratios ($\delta_i^{UK} / \delta_i^{FR}$)



The figure shows the cumulative distribution of knowledge access efficiency in the UK relative to France ($\delta_i^{UK} / \delta_i^{FR}$). In technology categories with a ratio greater 1, it was easier for people to access the knowledge needed to become an inventor in Britain than in France. These are about 80% of all categories. This advantage is partially offset by the fact that France had a population roughly twice as large as Britain's, meaning twice as many potential inventors. After accounting for differences in population, Britain still retained an absolute advantage in training workers into researchers in about 50% of technology categories.

Table 4: Technologies with the largest advantage in knowledge access

Top British category	δ ratio	Top French category	δ ratio
Agricultural Produce	26.9	Aerial Conveyances	0.226
Bleaching, Washing, Scouring	12.7	Tar, Resin, Oil Distillation	0.333
Preventing Accidents	9.8	Cooking and Baking	0.52
Mouldings and Ornaments	9.7	Coffee, Tea and Cocoa	0.581
Pipes and Tubes	9.7	Bed and bedding	0.648
Joints and Pulleys	9.2	Tunnels, Excavations, Embankments	0.657
Funeral Equipment	8.8	Games, Exercises, Amusements	0.665
Rope Manufacture	8.4	Baths and Bathing-Machines	0.672
Taps and Valves	5.8	Blacking	0.707
Mining and Quarrying	5.8	Electricity and Magnetism	0.72

δ is calculated according to equation 14, using data and the baseline estimates in Table 2, Column 1.

80 percent of all technology categories, i.e., $\delta_i^{UK}/\delta_i^{FR} > 1$. However, this advantage is partially offset by the fact that France had roughly twice the population of Britain at this time, meaning twice as many potential inventors.³⁶ After accounting for these differences, Britain still had an advantage in about one half of technology categories.

8.1 Validating the estimated knowledge access costs

Since our knowledge access parameters are novel, it is important to provide some empirical validation that they are picking up features of the actual economy we study. An ideal evaluation would compare our estimated parameters to some measure of the availability (or cost of acquiring) knowledge relevant for each technology type in each country. To approximate such a measure, we look at the availability of published knowledge. Previous work has highlighted the key role the printed knowledge, which became widely available following the introduction of Gutenberg’s moveable type printing press in the 15th century, played in knowledge sharing and economic growth in Europe (Dittmar, 2011).

Specifically, we examine the availability of technology-specific published knowledge in (British) English vs. French from the mid-18th century up to 1850 using the Google Books corpus (Michel et al., 2011).³⁷ The assumption here is that the availability of written

³⁶Though France had a population roughly twice as large as England & Wales, it does not appear to have produced as many new technologies. The larger French population was substantially less urbanized than the population of England & Wales and faced additional barriers to knowledge access, including the fact that many spoke different dialects (Blanc and Kubo, 2026). Such features likely posed important barriers to innovation. These are among the types of knowledge access barriers that are picked up by our knowledge access parameter estimates.

³⁷It is worth noting that both English and French are spoken by countries other than the U.K. and France. However, during the period we study it is reasonable to expect that the vast majority of the work produced in each of these languages was produced in these countries. We focus on British English, but this makes little difference during the time period that we study.

knowledge relevant for each technology type in a particular language is one important element of country-specific knowledge access.³⁸

To begin, we need to identify a set of equivalent keywords, in English and in French, that are strongly associated with each technology category in our data. We start with the patent titles for each patent in either the British or French patent datasets and then parse the titles into one and two-word grams. Each patent is then associated with a vector of gram indicators. We then run patent-level regressions, separately for British and French patents, to identify grams that are strong predictors for specific technology categories. This gives us a set of candidate keywords in English and in French.³⁹ We then translate these candidate keywords and verify that the translated terms are strong predictors for the same technology category in the other country.⁴⁰ Further details of this procedure and some examples of the produced equivalent keyword pairs are available in Appendix G.

Once we have pairs of equivalent keywords, one in English and one in French, we feed them into Google's N-gram Viewer and recover the number of documents in which the keyword appears.⁴¹ We focus on documents published between 1750 and 1850 in either the French or British English Google Books corpus.

This procedure gives us a measure of the availability of material in each language during the study period that relates to each technology category. We then compare this measure to our estimated knowledge access parameters. We can do this either at the technology category level, summing all keyword pairs for the category together, or at the level of keyword pairs. Our main analysis, in Table 5, is conducted at the level of technology categories. An analysis at the level of keyword pairs can be found in Appendix Table G.10.

Column 1 of Table 5 compares the log ratio of documents associated with each technology category in the Google Books corpus to the log ratio of our estimated knowledge access costs, focusing on the period 1750-1800. We find a positive and statistically significant association between our knowledge access cost ratio and the measure of knowledge contained in published sources in the two languages. In Column 2, we separate out documents in English and in French. We find that relatively better knowledge access in the UK compared to France is positively related to the availability of documents in English, and negatively related to the availability of documents in French. Columns 3-6 document similar results for later time periods.

³⁸This validation exercise is inspired in part by Juhász et al. (2025).

³⁹For each gram, we retain the technology category that received the highest t-statistic. Furthermore, we use a cutoff t-statistic of 5 for unigrams and a cutoff of 3 for bigrams to identify candidate keywords.

⁴⁰One reason that this translation procedure is useful is that it helps us identify complex keywords, particularly in French. For example, *railway* in English translates to *chemin de fer* in French and *steam engine* translates to *moteur à vapeur*. It is worth noting that a single keyword in one language may be part of more than one matched pair of keywords.

⁴¹N-grams also tells us the number of times the keyword appears across all documents. We focus on the number of documents as our primary outcome because we think this is the relevant statistic for our exercise.

Table 5: Comparing estimated knowledge access costs to Google Books documents

	Dep var: $\text{Ln } \delta^B / \delta^F$ ratio					
	1750–1800		1770–1830		1800–1850	
	(1)	(2)	(3)	(4)	(5)	(6)
Ln document ratio	0.306** (0.138)		0.320** (0.150)		0.340** (0.164)	
Ln English documents		0.272* (0.153)		0.281* (0.163)		0.297* (0.172)
Ln French documents		-0.297** (0.140)		-0.309** (0.151)		-0.332** (0.164)
N (Obs = category)	113	113	113	113	113	113
R^2	0.040	0.043	0.032	0.035	0.031	0.034

The three periods are the periods over which we aggregate document frequencies from Google Books N-gram Viewer. Document ratio is the ratio of English to French documents, where English refers to the British English corpus and French to the French corpus. p -val English = French is from a Wald test of coefficient equality. OLS regressions, robust standard errors in parentheses.

In Appendix G, we provide additional robustness check on these results. We show that similar results are obtained if we run the analysis at the level of keyword pairs, rather than aggregating to the technology category level. In addition, Appendix Figure G.6 provides a scatterplot, corresponding to the specification in Columns 1 of Table 5, which shows that our results are not driven by just one or two outliers.

Overall, these results indicate that our estimated knowledge access costs are correlated with the availability of printed material in each country related to each technology type. This provides some external validation for our model-derived parameter estimates.

The Google Books N-gram data can also be useful for evaluating whether accounting for the innovation network is important for our results. Suppose that we just used the level of patents in the two countries as an indicator of knowledge access costs, instead of using the structure of our model to back out the knowledge access parameters. Would we find similar patterns? The answer is, no. In Appendix Table G.11, we show that our estimated knowledge access parameters provide a much better prediction of the availability of knowledge in the two countries, as reflected by the documents in Google Books, than the ratio of patents in a technology in the two countries. This tells us that accounting for the influence of the innovation network using the structure of the model is important for obtaining useful estimates of underlying knowledge access costs.

9 Growth implications on the transition path

We now use the model to examine the growth implications of the differences in the knowledge access costs that we have estimated. While previous work has typically focused on growth effects in the balanced growth path (Liu and Ma, 2023, e.g.), such results are unlikely to be informative in our setting for two reasons. First, during the early stages of the Industrial Revolution it was unlikely that either of the economies we consider were near the balanced growth path. Second, our empirical results suggest that there were positive knowledge spillovers between Britain and France, which implies that their growth rates must converge in the long-run.

Instead, we focus on the growth differences along the transition path for two economies starting from similar initial conditions but with differences in their underlying knowledge access parameters. Specifically, we will run simulations for two economies, representing Britain and France, that start from the same arbitrary set of initial knowledge stock levels. These economies are characterized by the set of $\tilde{\delta}_{ci}$ parameters that we have estimated for the two countries, their populations, and the λ and ι parameters estimated in Section 7.⁴²

The only parameter we are missing, which enters when we allow for technology sharing across the economies, is η . Because η is related to production rather than innovation, it must be estimated separately from the other parameters using data on sectoral output. We do this by applying the structure in Eq. 10 to sector-specific output data from France in the early 19th century, as described in more detail in Appendix H. We focus on France because no similarly detailed industry output data are available for the UK in the early 19th century. Our results reveal a high level of technology sharing, with an estimated η of about 0.7 (0.5 would indicate perfect sharing, 1 indicates complete reliance on domestic technology).⁴³

Figure 5 plots our baseline simulation of the difference in the growth rate of output for our hypothetical British and French economies, starting from the same initial knowledge stock but differing in their knowledge access costs and populations. Since we do not have an accurate observation of the technology stock in these economies in the 18th century, our baseline simulation uses a common uniform initial knowledge stock ($q_{ic0} = 10 \forall i, c$).⁴⁴

The top line in this graph is for the closed economy case. This line suggests that, without knowledge spillovers or technology sharing, differences in the knowledge access costs parameters between the two economies would have given Britain a latent growth advantage of around one percentage-point over the first 2–3 decades of the Industrial Revolution.

⁴²It is important that we use the same population data here that we used when calculating the $\tilde{\delta}_{ci}$ parameters. Thus, we use estimated populations for the two countries in the decade 1825–34.


⁴³We will use this η value for both countries, but this means that our results allowing for technology sharing will likely represent a lower bound on growth differences between the two countries, because it requires that Britain is as reliant on French technology as France is on British technology.

⁴⁴The initial level of technology has no impact on the results, though allowing for variation in the initial technology level across technology categories does matter. We examine the impact of allowing variation in initial technology levels across technology categories in our robustness exercises.

The growth difference jumps around in the first few decades of the simulation, reflecting the complex influence of the knowledge access costs operating through the innovation network, before settling down after roughly four decades, at around 0.3 percentage points, as the economies converge to their balanced growth paths.

The middle dashed line in Figure 5 allows cross-country knowledge spillovers through the innovation network. Allowing cross-country knowledge spillovers has relatively little impact on relative growth initially, but it ensures that over the long run the difference in growth rates will slowly disappear. This is because, as Britain pulls ahead, France benefits more from cross-country knowledge spillovers, which pulls the growth rates in the two economies together.

Figure 5: Implied growth rate differences starting from a common initial technology level



The figure plots the difference in growth rates in two economies, one with the British knowledge access costs and population and the other with the French knowledge access costs and population, starting from an arbitrary knowledge stock of $q_0 = 10$. The closed economy case does not allow knowledge spillovers or technology sharing. The knowledge spillover case allows knowledge spillovers using our estimate of $\iota = 0.963$. The technology sharing case additionally allows technology sharing using our estimate of $\eta = 0.7$.

The lower dotted line in Figure 5 allows both cross-country knowledge spillovers and technology sharing. Allowing technology sharing influences the growth rate difference immediately, because output in France benefits from technology improvements in Britain. This line suggests a growth difference that starts out at around one-half percentage point and falls to 0.1 percentage point after three decades.

We view the results with cross-country knowledge spillovers and technology sharing as the relevant case for explaining the actual growth outcomes observed during the Indus-

trial Revolution. These results suggest that Britain’s knowledge access costs conferred a meaningful growth advantage relative to France in the first few decades of the transition to modern economic growth. Much of this advantage was offset by French use of technologies developed in Britain. However, even with substantial technology sharing our model suggests that British growth was several tenths of a percentage point higher in the first few decades of the transition to modern economic growth. Note that the technology sharing scenario should be thought of as a lower bound on Britain’s growth advantage, because we have forced Britain to be as dependent on French technology as France is on British technology.

Appendix I.1 presents additional figures describing the behavior of our two simulated economies. One pattern these show is that the level of research and the arrival rate of innovations increased in the first 2–3 decades of the simulation and then leveled off, as the economies approach their balanced growth paths. Output was higher in France, due in part to a larger population and in part to a greater share of workers allocated to production rather than research, but Britain experienced more rapid growth. It is interesting to note that our simulation predicts growth rates that are higher than those actually observed. This reflects the fact that our model economies include none of the other frictions, such as credit constraints or threats to the security of property rights, that are likely to slow down the implementation of new technology in the real world. This highlights that our quantitative results should be thought of as growth differences without accounting for any other frictions not present in our theory.

We also study how the choice of initial technology levels q_0 affects our results. Note that the model’s functional form implies that, with a common uniform initial knowledge stock q_0 , the level of q_0 does not influence results.⁴⁵ In Appendix I.2, we thus consider the impact of allowing variation in initial knowledge levels across technology categories.⁴⁶ We find that our results are also robust to allowing realistic variation in initial technology levels.

9.1 Counterfactuals

Finally, we study the contribution of different channels to the growth rate differences in Figure 5 by turning off different elements of the model and re-running the simulations. A summary of these results, focusing on the difference in latent growth advantages (i.e., closed economy differences) is presented in Figure 6.

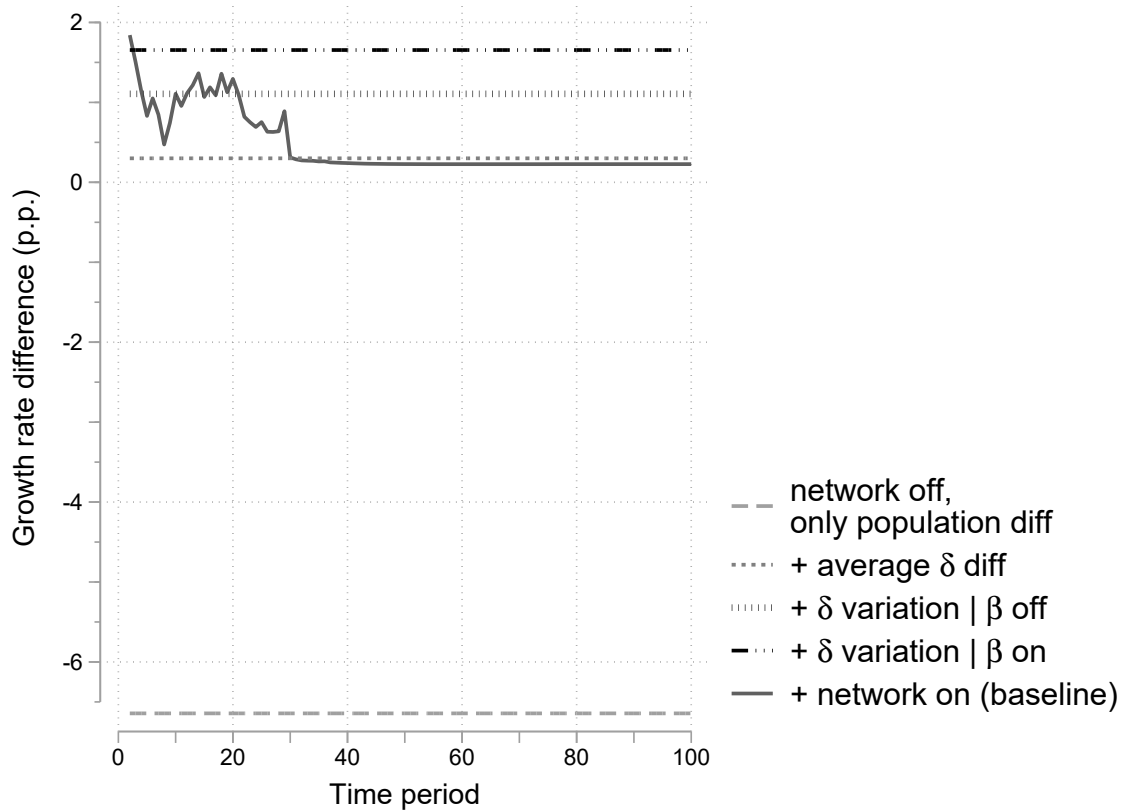
To understand this graph, it is useful to begin with the horizontal line at the bottom.

⁴⁵To see this, note that a uniform q_0 factors out of the ratio $q_{it}/c_{hi_{it}}$ in the first period. As a result, the effective research effort r_{it} and the innovation arrival rate ϕ_{it} are completely independent of q_0 .

⁴⁶While we can let the initial knowledge level vary across technologies, we do not want to let the level vary across countries, since the point of our exercise is to study the implications of the different knowledge access costs in two otherwise identical economies.

In that counterfactual case, we have turned off cross-technology spillovers through the innovation matrix by setting it to the identity matrix. We have also eliminated differences in knowledge access costs between the two economies by setting all of those parameters to the average French level. The only difference between the two economies is therefore their populations. We can see that, all else equal, a country with France’s population would have growth several percentage points faster than one with the population of England and Wales, simply by virtue of having a larger population. This is because with the same knowledge access costs, more people means more potential inventors.

Figure 6: Counterfactual latent growth advantages



The figure plots the difference in growth rates in two closed economies, starting from an arbitrary knowledge stock of $q_0 = 10$. The *network off, only population diff* line shows how this advantages changes if we set the innovation network Ω to the identity matrix and shut down all country differences except for population. The *+ average δ diff* line adds country differences in the weighted average knowledge access, calculated as $\bar{\delta}^c = \sum_i (\delta_i^c \sum_k \alpha_{ik} \beta_k)$. The *+ δ variation | β off* line adds the variation from the estimated country-by-technology knowledge access parameters δ_i^c , while keeping constant the production sector shares $\beta_k = 1/K$, where K is the number of production sectors. The *+ δ variation | β on* line adds variation in the production sector shares β_k . Finally, the *+ network on (baseline)* line reintroduces the innovation network and thus reconstitutes the full closed economy model, as depicted previously in Figure 5.

Next, we allow for differences in the average level of knowledge access costs between the two economies. We do this by setting the knowledge access costs in each technology

and country equal to a weighted average of our knowledge access parameters for all technologies in the country.⁴⁷ The dotted line shows that, once we allow for both population differences and differences in average knowledge access costs, Britain has a mild growth advantage of around 0.3 p.p. Thus, Britain’s lower knowledge access costs on average more than compensate for France’s larger population.

We then incorporate the contribution of variation in knowledge access costs across technology categories. As a first step, we allow this variation to affect growth only through the generality of each technology, represented by the α_{ik} parameters; we set the β_k parameters to be equal for all production sectors, so that results are not driven by the weight of different production sectors. This leads us to the horizontal hashed line located at around 1 p.p. in the graph. This tells us that Britain benefited from an advantageous allocation of knowledge access costs in more general technology sectors, such as steam engines, which were used more intensively or across a wider set of production sectors.

The top horizontal line adds in the influence of differences in the weight of production sectors k in consumption, using the β_k values measured in the data. This provides an additional growth advantage to Britain, which tells us that Britain also had better knowledge access in technologies that contributed to production sectors with a greater weight in consumption.

Finally, the solid line in the graph incorporates the effect of the innovation network to bring us back to the full model. This makes growth a more dynamic problem, since knowledge access cost advantages in one sector are eventually eroded as the technology level in the sector becomes relatively higher, slowing growth through the “low hanging fruit” effect. We can see that Britain’s initial growth advantage is similar to what we would find without the innovation network, but that over time it falls as the economy converges toward a balanced growth path with a lower growth difference.⁴⁸

To summarize, these results show us that Britain’s growth benefit was a product of both lower knowledge access costs on average, which more than offset its lower population relative to France, and how those costs were distributed across technologies and production sectors.

10 Conclusions

We have studied the source of Britain’s innovation advantage during the Industrial Revolution using a combination of modern growth theory and detailed micro-data. This allows us to provide the first quantitative assessment of *one source* of British advantage during

⁴⁷To specific, we calculate the weighted average of knowledge access parameters in each country weighted by the contribution of each technology in consumption. So, for each country we have $\bar{\delta}^c = \sum_i (\delta_i^c \sum_k \alpha_{ik} \beta_k)$.

⁴⁸Note that it appears in the graph that the long-run growth difference converges to the difference observed when we allow only for average differences in knowledge access costs. This is a coincidence of the simulation, not a necessary result of the model.

this crucial period of economic history. Our results reveal that Britain had a persistent *latent growth advantage* during the first few decades of the Industrial Revolution due to a combination of lower knowledge access costs on average and how those costs were distributed across technology types. However, the impact of this advantage on actual growth rates was substantially reduced by technology sharing, i.e., the use of British technology by French producers.

Our results provide quantitative empirical support for the argument, put forth most forcefully by Joel Mokyr, that knowledge access costs played an important role in determining innovation patterns during the Industrial Revolution (Mokyr, 2002, 2005, 2016). Of course, as Mokyr points out, there are many types of knowledge and many different factors that influence knowledge access costs. Our knowledge access costs estimates will reflect only a subset of costs, those that influence the production of patented knowledge at the country-by-technology level. This measure is likely to miss other important types of knowledge, such as access to basic scientific insights. Nevertheless, by providing the first quantitative country-by-technology measure of knowledge access costs, our results take a first step toward quantifying and measuring the role of such costs in the transition to modern economic growth.

In order to produce these results, we have extended previous theories in order to incorporate rich variation in knowledge access costs and decentralized choices by potential researchers into an endogenous growth model featuring an innovation network. One interesting feature of our framework, which should be explored in future work, is that not all technology sectors are active in all periods. Instead, over time the evolution of knowledge stocks, acting through the innovation network, can trigger the emergence of research in new types of technologies. Thus, our framework provides a potential starting point for examining how the observed innovation network evolves over time as new technology types emerge.

Finally, we have developed a set of new methodological tools that open up the possibility of investigating the role of innovation networks over longer periods of time. Specifically, we have shown how we can use historical patent data to measure the structure of the innovation network in periods where citation data are missing, and how we can use the structure of the model to back out a measure of research effort in settings where direct measures of R&D effort are not available. This opens up a range of new possibilities for studying long-run innovation patterns.

References

- ACEMOGLU, D., U. AKCIGIT, H. ALP, N. BLOOM, AND W. KERR (2018): “Innovation, Reallocation, and Growth,” *American Economic Review*, 108, 3450–3491. 24
- ACEMOGLU, D., U. AKCIGIT, AND W. R. KERR (2016): “Innovation Network,” *Proceedings of*

- the National Academy of Sciences*, 113, 11483–11488. 5, 17
- ACEMOGLU, D., D. CANTONI, S. JOHNSON, AND J. A. ROBINSON (2011): “The Consequences of Radical Reform: The French Revolution.” *American Economic Review*, 101, 3286–3307. 7
- ACEMOGLU, D., S. JOHNSON, AND J. ROBINSON (2005): “The Rise of Europe: Atlantic Trade, Institutional Change, and Economic Growth.” *American Economic Review*, 95, 546–579. 1
- ALLEN, R. C. (2009a): *The British Industrial Revolution in Global Perspective*, Cambridge University Press. 1, 6, 19
- (2009b): “The Industrial Revolution in Miniature: The Spinning Jenny in Britain, France, and India,” *The Journal of Economic History*, 69, 901–927. 6
- ALMELHEM, A., M. IYIGUN, A. KENNEDY, AND J. RUBIN (2026): “Enlightenment Ideals and Belief in Progress in the Run-up to the Industrial Revolution: A Textual Analysis,” *The Quarterly Journal of Economics*, 141, 263–314. 1
- ASSELAIN, J.-C. (2007): “Le Projet Française d’Histoire Économique Quantitative: Ambitions et Résultats,” *Economies et Sociétés (Serie ‘Histoire Economique Quantitative’)*, *Association Française de Cliométrie*, 36, 567–609. 7
- BLANC, G. AND M. KUBO (2026): “The Making of France,” . 29
- BOTTOMLEY, S. (2014a): *The British Patent System during the Industrial Revolution 1700–1852: From Privilege to Property*, Cambridge University Press. 5
- (2014b): “Patenting in England, Scotland and Ireland during the Industrial Revolution, 1700–1852,” *Explorations in Economic History*, 54, 48–63. 5
- (2019): “The Returns to Invention during the British Industrial Revolution,” *The Economic History Review*, 72, 510–530. 5
- BROADBERRY, S., B. M. S. CAMPBELL, A. KLEIN, M. OVERTON, AND B. VAN LEEUWEN (2015): *British Economic Growth, 1270–1870*, Cambridge: Cambridge University Press. 7
- BROADBERRY, S. N. (1994): “Technological Leadership and Productivity Leadership in Manufacturing since the Industrial Revolution: Implications for the Convergence Debate,” *The Economic Journal*, 104, 291–302. 6
- BRUNT, L., J. LERNER, AND T. NICHOLAS (2012): “Inducement Prizes and Innovation,” *Journal of Industrial Economics*, 60, 657–696. 5
- BURTON, M. D. AND T. NICHOLAS (2017): “Prizes, Patents and the Search for Longitude,” *Explorations in Economic History*, 64, 21–36. 5
- CAI, J. AND N. LI (2019): “Growth through Inter-Sectoral Knowledge Linkages,” *The Review of Economic Studies*, 86, 1827–1866. 5
- CANTONI, D. AND N. YUCHTMAN (2014): “Medieval Universities, Legal Institutions, and the Commercial Revolution,” *The Quarterly Journal of Economics*, 129, 823–887. 4
- CERVELLATI, M. AND U. SUNDE (2005): “Human Capital Formation, Life Expectancy, and the Process of Development,” *American Economic Review*, 95, 1653–1672. 4
- CINNIRELLA, F., E. HORNUNG, AND J. KOSCHNICK (2025): “Flow of Ideas: Economic Societies

- and the Rise of Useful Knowledge,” *The Economic Journal*, 135, 1496–1535. 4
- CRAFTS, N. (1998): “Forging Ahead and Falling behind: The Rise and Relative Decline of the First Industrial Nation,” *Journal of Economic Perspectives*, 12, 193–210. 6
- CRAFTS, N. AND C. K. HARLEY (1992): “Output Growth and the British Industrial Revolution: A Restatement of the Crafts-Harley View.” *Economic History Review*, 45, 703–730. 7
- CRAFTS, N. F. (1977): “Industrial Revolution in England and France: Some Thoughts on the Question, “Why Was England First?”,” *The Economic History Review*, 30, 429–441. 1, 6
- (1995): “Macroinventions, Economic Growth, and Industrial Revolution’ in Britain and France,” *The Economic History Review*, 48, 591–598. 1, 6
- CROUZET, F. (1996): “La Première Révolution Industrielle,” in *Histoire de La France Industrielle*, ed. by M. Lévy-Leboyer, Larousse. 7
- 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. 4
- DITTMAR, J. E. (2011): “Information Technology and Economic Change: The Impact of the Printing Press,” *Quarterly Journal of Economics*, 126, 1133–1172. 4, 29
- DITTMAR, J. E. AND R. R. MEISENZAHN (2024): “The Research University, Invention, and Industry: Evidence from German History,” . 4
- DORNER, M. AND D. HARHOFF (2018): “A Novel Technology-Industry Concordance Table Based on Linked Inventor-Establishment Data,” *Research Policy*, 47, 768–781. 45
- EVENSON, R. E., S. KORTUM, AND J. PUTNAM (1988): “Estimating Patents by Industry Using the Yale-Canada Patent Concordance,” *Yale University, New Haven, mimeo*. 45
- FERNIHOUGH, A. AND K. O’ROURKE (2021): “Coal and the European Industrial Revolution,” *Economic Journal*, 131, 1135–1149. 1
- GALOR, O. (2011): *Unified Growth Theory*, Princeton University Press. 4
- GALOR, O. AND D. N. WEIL (2000): “Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and Beyond,” *American Economic Review*, 90, 806–828. 4
- GRILICHES, Z. (1990): “Patent Statistics as Economic Indicators: A Survey,” *Journal of Economic Literature*, 28, 1661–1707. 15
- HALLMANN, C., L. ROSENBERGER, AND E. E. YAVUZ (2021): “Invention and Technological Leadership during the Industrial Revolution,” *mimeo*. 14
- HANLON, W. W. (2023): “The Rise of the Engineer: Inventing the Professional Inventor during the Industrial Revolution,” . 5, 14, 54
- HANSEN, G. D. AND E. C. PRESCOTT (2002): “Malthus to Solow,” *American Economic Review*, 92, 1205–1217. 4
- HARHOFF, D., F. NARIN, F. M. SCHERER, AND K. VOPEL (1999): “Citation Frequency and the Value of Patented Inventions,” *The Review of Economics and Statistics*, 81, 511–515. 15, 57

- HARRIS, J. (2017): *Industrial Espionage and Technology Transfer*, London: Routledge. 6, 7
- HORREL, S., J. HUMPHRIES, AND M. WEALE (1994): “An Input-Output Table for 1841,” *The Economic History Review*, 47, 545–566. 15
- HUANG, J. AND Y. ZENOU (2020): “Key Sectors in Endogenous Growth,” . 5
- 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. 17
- JUHÁSZ, R. (2018): “Temporary Protection and Technology Adoption: Evidence from the Napoleonic Blockade,” *American Economic Review*, 108, 3339–76. 4, 7
- JUHÁSZ, R., S. SAKABE, AND D. WEINSTEIN (2025): “Codification, Technology Absorption, and the Globalization of the Industrial Revolution,” Tech. Rep. w32667, National Bureau of Economic Research. 4, 30
- JUHÁSZ, R., M. P. SQUICCIARINI, AND N. VOIGTLÄNDER (2024): “Technology Adoption and Productivity Growth: Evidence from Industrialization in France,” *Journal of Political Economy*, 132, 3215–3259. 4
- KELLY, M., J. MOKYR, AND C. Ó GRÁDA (2014): “Precocious Albion: A New Interpretation of the British Industrial Revolution,” *Annual Review of Economics*, 6, 363–389. 1
- (2023): “The Mechanics of the Industrial Revolution,” *Journal of Political Economy*, 131, 59–94. 1
- KHAN, B. Z. (2018): “Human Capital, Knowledge and Economic Development: Evidence from the British Industrial Revolution, 1750–1930,” *Cliometrica*, 12, 313–341. 5
- KHAN, B. Z. AND K. L. SOKOLOFF (2004): “Institutions and Democratic Invention in 19th-Century America: Evidence from “Great Inventors,” 1790-1930,” *The American Economic Review*, 94, pp. 395–401. 5
- KHAN, Z. (2005): *The Democratization of Invention: Patents and Copyrights in American Economic Development, 1790-1920*, Cambridge: Cambridge University Press. 5
- KORTUM, S. AND J. PUTNAM (1997): “Assigning Patents to Industries: Tests of the Yale Technology Concordance,” *Economic Systems Research*, 9, 161–176. 45
- LANDES, D. S. (1969): *The Unbound Prometheus*, London: Cambridge University Press. 6, 19
- LIU, E. AND S. MA (2023): “Innovation Networks and R&D Allocation,” NBER Working Paper 29607. 1, 2, 5, 8, 11, 17, 22, 23, 32, 49
- LYBBERT, T. J. AND N. J. ZOLAS (2014): “Getting Patents and Economic Data to Speak to Each Other: An ‘Algorithmic Links with Probabilities’ Approach for Joint Analyses of Patenting and Economic Activity,” *Research Policy*, 43, 530–542. 45
- MACLEOD, C., J. TANN, J. ANDREW, AND J. STEIN (2003): “Evaluating Inventive Activity: The Cost of Nineteenth-Century UK Patents and the Fallibility of Renewal Data,” *The Economic History Review*, 56, pp. 537–562. 5
- MARKOVITCH, T. J. (1965): “L’industrie Française de 1789 à 1964. Sources et Méthodes,” *Cahiers de L’I.S.E.A.*, AF4. 70

- MICHEL, J.-B., Y. K. SHEN, A. P. AIDEN, A. VERES, M. K. GRAY, THE GOOGLE BOOKS TEAM, J. P. PICKETT, D. HOIBERG, D. CLANCY, P. NORVIG, J. ORWANT, S. PINKER, M. A. NOWAK, AND E. L. AIDEN (2011): "Quantitative Analysis of Culture Using Millions of Digitized Books," *Science*, 331, 176–182. 29
- MIURA, H. (2012): "Stata Graph Library for Network Analysis," *Stata Journal*, 12, 94–129. 51
- MOKYR, J. (1990): *The Lever of Riches: Technological Creativity and Economic Progress*, Oxford University Press. 6, 19
- (1999): "Editor's Introduction: The New Economic History and the the Industrial Revolution," in *The British Industrial Revolution*, ed. by J. Mokyr, 2 ed. 6
- (2002): *The Gifts of Athena*, Princeton and Oxford: Princeton University Press. 1, 2, 37
- (2005): "The Intellectual Origin of Modern Economic Growth," *Journal of Economic History*, 65, 285–351. 1, 2, 6, 37
- (2009): *The Enlightened Economy*, New Haven and London: Yale University Press. 1, 2
- (2016): *A Culture of Growth: Origins of the Modern Economy*, Princeton University Press. 37
- MOSER, P. (2005): "How Do Patent Laws Influence Innovation? Evidence from Nineteenth-Century World's Fairs." *American Economic Review*, 95, 1214. 5, 16, 53
- (2012): "Innovation without Patents: Evidence from World's Fairs," *Journal of Law and Economics*, 55, pp. 43–74. 2, 5, 16, 53
- NEWMAN, M. E. J. (2010): *Networks: An Introduction*, Oxford University Press. 51
- NICHOLAS, T. (2011): "Cheaper Patents," *Research Policy*, 40, 325–339. 5
- NORTH, D. C. AND B. R. WEINGAST (1989): "Constitutions and Commitment: The Evolution of Institutional Governing Public Choice in Seventeenth-Century England," *The Journal of Economic History*, 49, 803–832. 1
- NUVOLARI, A. AND V. TARTARI (2011a): "Bennet Woodcroft and the Value of English Patents, 1617–1841," *Explorations in Economic History*, 48, 97–115. 16, 57
- (2011b): "Bennet Woodcroft and the Value of English Patents, 1617-1841," *Explorations in Economic History*, 48, 97–115. 5
- NUVOLARI, A., V. TARTARI, AND M. TRANCERO (2021): "Patterns of Innovation during the Industrial Revolution: A Reappraisal Using a Composite Indicator of Patent Quality," *Explorations in Economic History*, 82. 5, 16, 57, 61
- NUVOLARI, A., G. TORTORICI, AND M. VASTA (2023): "British-French Technology Transfer from the Revolution to Louis Philippe (1791–1844): Evidence from Patent Data," *The Journal of Economic History*, 83, 833–873. 4, 5
- O'BRIEN, P. AND C. KEYDER (2011): *Economic Growth in Britain and France 1780-1914*, New York: Routledge. 27

- PERETTO, P. F. (2015): “From Smith to Schumpeter: A Theory of Take-off and Convergence to Sustained Growth,” *European Economic Review*, 78, 1–26. 4
- POMERANZ, K. (2000): *The Great Divergence*, Princeton: Princeton University Press. 1
- RIDOLFI, L. AND A. NUVOLARI (2021): “L’histoire Immobile? A Reappraisal of French Economic Growth Using the Demand-Side Approach, 1280–1850,” *European Review of Economic History*, 25, 405–428. 4, 7
- SCHANKERMAN, M. AND A. PAKES (1986): “Estimates of the Value of Patent Rights in European Countries During the Post-1950 Period,” *The Economic Journal*, 96, 1052–1076. 15, 57
- SQUICCIARINI, M. P. (2020): “Devotion and Development: Religiosity, Education, and Economic Progress in Nineteenth-Century France,” *American Economic Review*, 110, 3454–3491. 4
- SQUICCIARINI, M. P. AND N. VOIGTLÄNDER (2015): “Human Capital and Industrialization: Evidence from the Age of Enlightenment,” *The Quarterly Journal of Economics*, 130, 1825–1883. 4
- SULLIVAN, R. J. (1989): “England’s “Age of Invention”: The Acceleration of Patents and Patentable Invention during the Industrial Revolution,” *Explorations in Economic History*, 26, 424–452. 5
- (1990): “The Revolution of Ideas: Widespread Patenting and Invention during the English Industrial Revolution,” *The Journal of Economic History*, 50, 349–362. 5
- THOMAS, M. (1984): “An Input–Output Approach to the British Economy, 1890–1914,” Ph.D. thesis, Oxford. 15, 45, 70
- TRAJTENBERG, M. (1990): “A Penny for Your Quotes: Patent Citations and the Value of Innovations,” *The RAND Journal of Economics*, 21, 172–187. 15, 57
- VOIGTLÄNDER, N. AND H.-J. VOTH (2006): “Why England? Demographic Factors, Structural Change and Physical Capital Accumulation during the Industrial Revolution,” *Journal of Economic Growth*, 11, 319–361. 1
- VOTH, H.-J., B. CAPRETTINI, AND A. TREW (2022): “Fighting for Growth: Labor Scarcity and Technological Progress During the British Industrial Revolution,” Mimeo. 1
- WOODCROFT, B. (1854): *Titles of Patents of Invention, Chronologically Arranged*, London: George Edward Eyre and William Spottiswoode. 43

A Concordant technology categories

This appendix details how we construct 123 concordant technology categories into which both British and French patents can be reliably classified. Our “informed approach” exploits information on how the respective national historical patent offices classified a given invention by considering patents that were filed in both countries. We then establish a systematic classification into concordant technology categories by cleaning, combining, and regrouping this starting information.

A.1 Starting sample: Bi-national patents

To create a starting sample, we manually identified 1,148 patents filed in both countries. These bi-national patents give us a list of how British *subcategories* correspond to French *subcategories* (keywords).

Bi-national patents before 1844. The search started with patents in Britain and France stating a foreign (French or British) address. We determined a match based on the inventor’s name, the patent title, and the temporal proximity of the patent date. As a result, 167 French patents filed before 1844 could be linked to a British patent, and 127 British patents filed before 1844 could be linked to a French patent.

Most likely, there were more bi-national patents, but the main difficulty in identifying them is that (i) patents sometimes report the name of the patent agent in one country and the name of the inventor in the other, making name-based search and identification infeasible; and that (ii) patent titles were sometimes altered or abbreviated during patenting abroad, making title based search and identification difficult. Regarding temporal proximity, we found that a diffusion lag of one year (and sometimes more) was not unusual, and accounted for it in the matching process.

British patents in France after 1844. After the French patent law reform of 1844, the French patent data explicitly state foreign patents’ country of origin and original filing date. The information exists because the reform principally recognized the priority of foreign patents while stipulating that the maximum patent duration in France would start with the foreign filing date.

Out of 916 “British patents” in France during 1844–1852, we can assign 855 (93%) to the original British patent as listed in Woodcroft (1854, Vol.2). The remainder is lost due to alterations of the patent title (e.g. shortening during translation) or the usage of patent agents in Britain (in which case the British patent lists a different inventor than the French “British” patents).⁴⁹

⁴⁹Curiously, these 855 French patents correspond to 808 British patents: some British patents were split into two when patented abroad.

A.2 Establishing concordance

Given the starting sample's list of how British *subcategories* correspond to French *subcategories* (keywords), our goal is to achieve exact correspondence and aggregation to technology *categories*.

To make the British classification correspond to the French, we aggregate, combine, and condense existing (sub)categories into new categories as necessary. For example, the concordant category “Flour Milling,” which existed already in the French classification system, is aggregated in the following British categories > subcategories: “Agricultural Produce > Apparatus Used In Filling Flour-Sacks; Fastenings For Flour-Sacks”, “Agricultural Produce > Grinding and Crushing Corn and Other Grain and Seeds” and “Grinding, Cutting, and Crushing > Corn and Other Grain, -Mills For Grinding”. This procedure reduces the initial 147 British technology categories to 133 refined categories.

To make the French classification correspond to the British, we assign French technology subcategories (keywords) to the new set of 133 (potentially concordant) categories. In a few cases, we cannot resolve ambiguity and do not assign keywords to any category. For example, one bi-national patent was classified into the French keyword “Yarn and fabric singeing” and into two British categories, “Spinning and Preparing For Spinning” and “Cloth Fulling, Dressing, Cutting, and Finishing.” Yet, both seem equally relevant—and too distinct to warrant aggregation. In sum, we can reliably assign French keywords to 123 concordant technology categories.⁵⁰

⁵⁰To nine (refined) British categories we cannot assign any French keyword: “Alarms, Snares, and Vermin Traps”, “Assurance: Preventing Forgery and Fraud”, “Bearings, Wheels, Axles, and Driving-Bands”, “Boring, Drilling, Punching”, “Chains and Chain-Cables”, “Cutting, Sawing, and Shaping”, “Friction,-Diminishing and Destroying”, “Papier-Mache and Japanned Wares”, and “Springs and Buffers”. We lost another category “Safes and Other Depositories” because there were no French patents in this category before 1844.

B Mapping from technologies to output sectors

This appendix describes the method we use for mapping technology categories to output sectors. The starting point for this exercise is the set of output sectors included in the input-output matrix constructed by Thomas (1984). This dataset includes estimates of the weight of each sector in consumption (the β parameters in our model). We need to map these output sectors to technology categories in order to obtain the α parameters in our model.

It is worth noting that, even in modern settings, generating mappings from technology categories to industries is challenging. One method relies on manual assignment of patents to both technology categories and industries undertaken by patent offices (Evenson, Kortum, and Putnam, 1988; Kortum and Putnam, 1997). Another approach uses algorithms to compare keywords in industry descriptions to the language in patents (Lybbert and Zolas, 2014). A third approach uses the fact that firms are classified by industry and file patents which are classified into technology categories (Dorner and Harhoff, 2018). Our approach is similar to the third of these alternatives.

Our mapping exploits the occupation information available in the patent data. Our assumption is that on average individuals working in a particular industry are more likely to be inventing technologies associated with that industry. The British patent data include roughly 7,000 occupation titles for inventors of patents filed between 1700 and 1849. A substantial fraction of these occupations unambiguously match to industries present in the IO table. We manually review these occupational titles and map them to the output sectors. We then use the fact that patents are classified into technology categories to generate a mapping between technology categories and output sectors.

To provide a sense of the types of occupations corresponding to different industries, Table B.1 lists by IO industry the three most common “topics” contained in occupations that help us to establish unambiguous matches. Generally, we do not match generic occupations that refer to professions or class/status (e.g. merchant, manager, worker, officer) unless they are qualified by a topic that refers unambiguously to one industry. For example, we do not match “engineers” to the industry “Engineering” because the unqualified occupation title refers to a profession rather than an industry. However, we match coal mining (colliery) engineer to the coal mining industry because in this case, the qualifying topic is unambiguous.⁵¹

Once we have a mapping from occupations to industries, the mapping from technology categories to industries is straightforward given that occupations are associated with

⁵¹Some professions are ambiguous even if qualified by a topic, for example “coal merchant” or “cloth merchant” because we do not know if this occupation worked in industry or in the excluded distribution services. One exception to the rule are composite occupations like “woollen manufacturer and merchant” because there the “manufacturer” clearly indicated involvement in production.

Table B.1: Information used for matching input–output industries to occupations

Input–output industry	Most common occupation theme		
	ranked 1st	2nd	3rd
Agriculture, Forestry, etc	farmer	agriculturalist	planter
Coal Mining	coal	colliery	viewer
Other mining	quarry	quarryman	engineer
Coke ovens	coke	burner	breeze
Iron and Steel	iron	steel	founder
Non-ferrouse metals	brass	founder	tin
Engineering	machine	agricultural	engine
Metal Goods, NES	tool	lock	wire
Shipbuilding	ship	builder	shipwright
Railway Rolling stock	railway	builder	carriage
Cotton and silk	cotton	spinner	silk
Woolen and worsted	wool	spinner	worsted
Hosiery and lace	lace	hosier	hosiery
Other textiles	carpet	elastic	cloth
Jute, hemp, and linen	flax	spinner	rope
Textile finishing	dyer	finisher	printer
Clothing	hat	tailor	clothier
Boot and shoe	boot	shoe	gutta-percha
Leather and fur	leather	harness	currier
Food processing	miller	baker	sugar
Drink	brewer	water	distiller
Tobacco	cigar	tobacco	snuff
Chemicals	chemist	oil	chemical
Paper	paper	card	stainer
Printing and publishing	printer	stationer	publisher
Rubber	india-rubber	rubber	gutta-percha
Timber trades	sawyer	mill	saw
Furniture	cabinet	dressing	case
Other wood	block	bobbin	wood
Building materials	brick	tile	stone
Building, etc.	builder	architect	painter
Misc. Manufactures	instrument	glass	watch
Gas, electricity, water	gas	meter	apparatus

The topics are obtained from breaking splitting the occupation string in parts, e.g. “iron founder” into “iron” and “founder”. The table excludes generic themes such as manufacturing, manufacturer, maker, worker, master, manager, agent, proprietor. Note that we do not match the occupations to industries based on individual themes but based on the information contained in the full occupation string.

patents which are classified into technology categories.⁵² The resulting probabilistic mapping from technology categories to industries appear quite reasonable. To illustrate this, Table B.2 lists for each industry the two most important technology categories with their respective α_{ik} value. The resulting mapping make sense: For example, for the industry *Agriculture, Forestry, etc*, the top technology category is *Agriculture* with $\alpha = 0.56$ and the second technology category is *Agricultural Produce* with $\alpha = 0.1$. The table also highlights how knowledge from some technology categories is used in several industries: For example, the *Spinning and Preparation* technology category has the highest α in both *Cotton and Silk* and *Woolen and Worsted* industries.

⁵²Two minor technology categories are missing because we were unable to map their associated occupations to any IO industry. These are “Diving, engines for diving” and “Maps and Globes”.

Table B.2: Most important technology category (largest α) by industry

Industry	1st technology category	α_{ik}	2nd technology category	α_{ik}
Agriculture, Forestry, etc	Agriculture	0.56	Agricultural Produce	0.1
Coal Mining	Pumps and Drainage	0.16	Salt and Saltpetre	0.14
Coke ovens	Stoves, Ovens, Kilns	0.5	Heating and Evaporating	0.5
Iron and Steel	Metallurgy	0.31	Agriculture	0.08
Non-ferrous metals	Lighting and Matches	0.1	Metallurgy	0.07
Engineering	Spinning and Preparation	0.39	Fabric Manufacture	0.22
Metal Goods, NES	Locks and Other Fastenings	0.1	Cutlery	0.08
Shipbuilding	Shipbuilding and Operation	0.52	Pumps and Drainage	0.13
Railway Rolling stock	Railways and Locomotives	0.5	Road Vehicles	0.25
Cotton and silk	Spinning and Preparation	0.58	Fabric Manufacture	0.26
Woolen and worsted	Spinning and Preparation	0.53	Fabric Manufacture	0.33
Hosiery and lace	Fabric Manufacture	0.79	Spinning and Preparation	0.08
Other textiles	Fabric Manufacture	0.66	Rubber and Gutta-Percha	0.08
Jute, hemp, and linen	Spinning and Preparation	0.54	Rope Manufacture	0.26
Textile finishing	Fabric Manufacture	0.27	Printing	0.24
Clothing	Fabric Manufacture	0.37	Wearing-Apparel	0.3
Boot and shoe	Footwear	0.58	Buttons and Fastenings	0.13
Leather and fur	Leather and Tanning	0.46	Saddlery and Harness	0.27
Food processing	Agricultural Produce	0.15	Sugar Manufacture	0.15
Drink	Brewing and Distilling	0.48	Steam Engines and Boilers	0.06
Tobacco	Fuel	0.4	Tobacco and Snuff	0.4
Chemicals	Chemical Products	0.22	Metallurgy	0.08
Paper	Paper and Pasteboard	0.56	Spinning and Preparation	0.13
Printing and publishing	Printing	0.49	Power and Propulsion	0.04
Rubber	Rubber and Gutta-Percha	0.5	Fabric Manufacture	0.5
Timber trades	Wood-Working	0.67	Lighting and Matches	0.33
Furniture	Furniture and Cabinet-ware	0.38	Wood-Working	0.09
Other wood	Shipbuilding and Operation	0.35	Arms and Ammunition	0.23
Building materials	Building Materials	0.4	Stone and Plaster Work	0.1
Building, etc.	Building Processes	0.18	Building Materials	0.07
Misc. Manufactures	Musical Instruments	0.19	Timekeeping Instruments	0.13
Gas, electricity, water	Measurement instruments	0.46	Gas Manufacture and Consumption	0.14

The shares α state how much each industry relies on each technology, with $\sum_i \alpha_{ik} = 1$. Source: See text.

C Validating the network measure using modern data

Because our approach to measuring the innovation matrix is novel, it is useful to provide some additional evidence showing that our approach provides an accurate measure of the underlying innovation network. To validate our approach, we turn to modern patent data, where we can observe both citations and individual identifiers for inventors that allow us to link their patents.

Our comparison focuses on the U.S. patent data provided in the 2015 version of PatStat. The PatStat database provides individual identifiers, International Patent Classification (IPC) technology categories for each granted patent, and bilateral patent citations. Using these inputs, we can construct and compare innovation matrices based on either citations or on the inventor-based approach that used in our main analysis. To keep the size of the networks manageable, we focus on the “three digit” IPC level (e.g., A41: Wearing Apparel) and classify each patent based on the first (primary) IPC code provided by the U.S. Patent and Trademark Office (PTO). The result is a 123 x 123 matrix, a similar level of detail to the technology classifications used in our main analysis.

Our inventor-based innovation network is constructed using the approach shown in Eq. 11. Our citation-based network is generated using the approach used in Liu and Ma (2023) as well as other modern studies:

$$\omega_{ij} = Cites_{ij} / \sum_l Cites_{il}$$

where $Cites_{ij}$ is the number of patents in category i citing patents in category j .

We focus on citations between U.S. patents for this measure. Also, because we are interested in knowledge flows that contribute to the development of new technologies, we limit our analysis to only those citations provided by the patent applicant in the original submission. This excludes other citations, such as those added by the patent examiner in the search phase or those added during opposition, which identify related technologies but may have been unknown to the inventor at the time of invention. After these cuts, we are left with a total of just over 30 million bilateral citations between U.S. patents.

D Additional results on the innovation network

D.1 Network plots

Figure D.3 presents a fully labeled version of Figure 2. Besides the larger clusters labeled in the main figure, this appendix figures features many more smaller clusters.

Figure D.1: Fully labeled joint innovation network



This is the fully labeled version of Figure 2. Plot generated using multidimensional scaling. The edges of the joint network are computed as $(\omega_{ij}^{UK} + \omega_{ij}^{FR})/2$.

D.2 Network comparison

Table D.3 regression results comparing node centrality in the British network to that in the French network. We consider four centrality measures: Eigenvector, degree, closeness, and

betweenness. Across all measures, British and French node centrality appear strikingly similar. This holds with small variations for both level and logarithmic specifications.

Table D.3: Comparing node centrality in the French and British networks

	Dep var: British node centrality							
	Eigenvector		Degree		Closeness		Betweenness	
	(1) Level	(2) Log	(3) Level	(4) Log	(5) Level	(6) Log	(7) Level	(8) Log
French node centrality	0.804*** (0.096)	0.729*** (0.079)	0.625*** (0.059)	0.529*** (0.082)	0.301 (0.215)	0.870** (0.429)	0.453*** (0.094)	0.601*** (0.154)
<i>N</i> (Obs = nodes)	119	113	119	113	119	113	119	35
<i>R</i> ²	0.625	0.400	0.524	0.243	0.019	0.304	0.176	0.325

The table documents that the observed centrality of network nodes (technology categories) is highly similar across countries. For all four centrality measures, we focus on the *sending* centrality of nodes (e.g., left eigenvector or out-degree centrality) in line with the theory and our argument. The centrality measures are standard in the network literature (e.g. Newman, 2010) and implemented in Stata by Miura (2012). OLS regressions, robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

D.3 Stability of innovation network over time

In this appendix we look at whether the innovation network is evolving over time, and if so, how fast it is changing. To do so, we have to grapple with the fact that we only observe the underlying matrix with error. As a result, if we take two snapshots of the matrix from two different sub-periods of our study period, they are certain to differ. The question is, how much of this difference is due to error in our ability to observe the underlying matrix, and how much of the difference reflects evolution of the matrix over time.

To study this question, it is useful to contrast two stylized models. In one, there is a fixed underlying matrix which we observe with error:

$$\omega_{ijt} = \bar{\omega}_{ij} + \epsilon_{ijt} \quad (15)$$

where $\bar{\omega}_{ij}$ is the ij 'th edge of the fixed underlying matrix.

Alternatively, we may think that the matrix is evolving over time, so that

$$\omega_{ijt} = \rho\omega_{ijt-1} + e_{ijt} \quad (16)$$

where e_{ijt} may be random error or a function of the knowledge stock.

We are interested in evaluating the extent to which the matrix is consistent with the second of these two views. To do so, we start with Eq. 16 and subtract out ω_{ijt-1} to obtain:

$$\omega_{ijt} - \omega_{ijt-1} = (\rho - 1)\omega_{ijt-1} + e_{ijt} \quad (17)$$

Using Eq. 17, we can obtain an estimate of ρ . Note that this is essentially a panel data analog of the setup for the Dickey–Fuller test. If we define $\delta = \rho - 1$, then we testing whether the model is a random walk amounts to testing the null hypothesis that $\delta = 0$. Alternatively, the model in Eq. 15 would suggest that we should observe $\delta = -1$.

Table D.4: Dickey–Fuller type regressions of two period-specific innovation networks

	Dep var: Edge change $\omega_{ij,t} - \omega_{ij,t-1}$				
	1834–1843		1840–1843		1824–1833
	(1)	(2)	(3)	(4)	(5)
	All edges	Diagonals	Off-diagonals	Same # pats p. period	Earlier period
Edge $\omega_{ij,t-1}$: 1824–1833	-0.837*** (0.030)	-0.705*** (0.100)	-0.970*** (0.007)	-0.840*** (0.031)	
Edge $\omega_{ij,t-1}$: 1790–1823					-0.803*** (0.035)
N (Obs = edges)	10500	97	10403	8649	9595
R^2	0.56	0.44	0.68	0.43	0.32

Dickey–Fuller type regressions of network changes on earlier networks. Using joint network edges throughout. OLS regressions. Standard errors clustered two-way by origin and destination technology category in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D.4 presents results based on Eq. 17 using innovation networks constructed using patents from subsets of our study period. In Columns 1–3, we compare the innovation network based on patents from 1834–43 to a network based on patents from ten years before, 1824–33. Column 1 looks across the full matrix, while Columns 2 and 3 separate the diagonal and off-diagonal elements. In all three cases, we can reject both the random walk model ($H_0: \delta = 0$) and the white noise model ($H_0: \delta = -1$), though the fact that we obtain estimates that are fairly close to zero suggests that the model is closer to the one in Eq. 15 than Eq. 16. In Column 4, we restrict the period used to construct the dependent variable to include the same number of patents as the period used to construct the explanatory variable, which is less than ten years because the number of patents was growing rapidly over time. In Column 5, we show similar results using earlier periods. Overall, these results strongly reject the random walk model, and though they do suggest that $\rho > 0$, we find evidence that ρ is not large. Overall, these findings suggest that the matrix we observe is fairly close to one characterized by a fixed underlying structure.

E Innovation network based on exhibition data

This appendix details our analysis of data on technologies exhibited at the Crystal Palace Exhibition in London 1851, the world’s first world’s fair.

As explained in the main text section 4.6, we use data collected by Moser (2005, 2012) covering 6,003 exhibits by British exhibitors and 1,675 exhibits by French exhibitors. Each entry includes information such as the exhibitor name and address and a brief description of the exhibit. Each exhibit was also classified into one of 29 categories.⁵³

Following the same procedure applied to the patent data, we manually reviewed all exhibits in order to link exhibits by the same exhibitors. Additionally, we manually linked individuals who exhibited at the Crystal Palace to inventors who appear in our patent data. Linking was based on the names and addresses of exhibitors. Manual review was necessary because there is substantial variation in the format of exhibitor names and addresses in the data. Of the British exhibits, 691 were linked to at least one other exhibition by the same exhibitor. For French exhibitors, 47 exhibits can be linked.

Using these linked exhibits, we calculate the similarity of different technology categories as follows. Denoting exhibitors (inventors) k and exhibit categories i , the strength of connections from category j to i is given by:

$$\omega_{ij}^{Exhibit} = \frac{E_{kij}}{\sum_k (E_{ki} + E_{kj})}$$

where E_{kij} is the number of exhibit pairs by the same exhibitor that can be constructed where one exhibit is classified into category i and the other is classified into category j , and E_i and E_j are the total number of exhibits in categories i and j , respectively.

Note that this approach is very similar to that applied to the patent data, except that the resulting exhibition network is undirected. The reason why we can’t distinguish E_{ij} from E_{ji} is that both exhibits are at the same time; we do not have information in which sequence the exhibitor created them. As a result, $\Omega^{Exhibit}$ is symmetric.⁵⁴

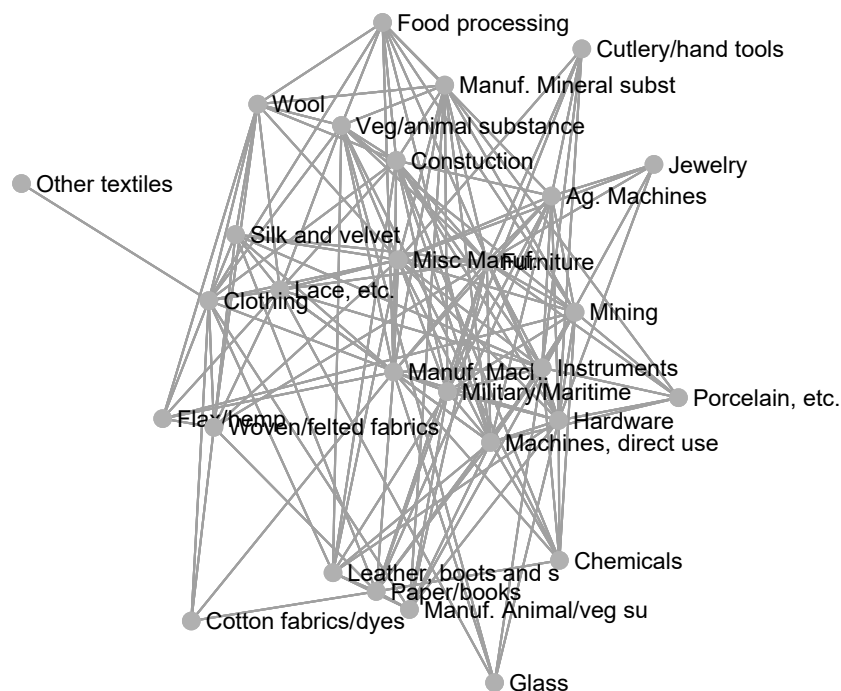
Figure E.2 presents the innovation network obtained by applying this method to the exhibition data.⁵⁵ The network features a dense central core, located just below and to the right of the center of the plot. The core includes technology types such as *Machines for direct use* (including steam engines and locomotives), *Manufacturing machines*, *Military/Maritime* (including shipbuilding), *Instruments*, *Mining*, and *Hardware*. Similarities between this

⁵³The original dataset includes 30 categories, but we omit category 30, which is for fine arts exhibits. For some French exhibits, it also provides secondary or tertiary categorizations for French exhibits. But as most exhibits have only one category, we focus on each exhibit’s first category.

⁵⁴The interpretation of the resulting exhibition-based innovation network is not as straightforward as the one obtained from a patent-based innovation network, since exhibiting in two categories might more likely reflect factors other than innovation spillovers.

⁵⁵Note that this technology space is built using links from both British and French exhibitors, though most of the available links come from British exhibitors.

Figure E.2: The technology network obtained from exhibition data



Plot is generated using multidimensional scaling.

innovation network and the patent-based one are already apparent.

To compare the exhibition-based innovation network to the patent-based innovation network more directly, we need to map patent technology categories into exhibition categories. One potential approach is to assign patent categories to exhibit categories based on their descriptions. However, this works well only for some categories but not for others. Instead, we adopt an approach that relies on inventors who both filed patents and exhibited at the Crystal Palace. Specifically, we use a list of British-based inventor-exhibitors from Hanlon (2023). The list was constructed by manually searching the British patent data for each exhibitor and also for inventors of exhibits mentioned in the exhibit descriptions, which might be different for example because ownership of a technology had been transferred. This manual review identified 634 exhibits associated with an individual that also patented, with links for 566 individuals in total (some individuals account for more than one exhibit).⁵⁶

Using the linked inventor-exhibitors, we assess whether the exhibition-based innova-

⁵⁶Note that it is impossible to associate exhibits with individual patents given the available information. In addition, exhibits may include more than one patented device.

Table E.5: Comparing patent-based and exhibit-based innovation network

	Dep var: Exhibition eigencentrality		
	Level	Log	Rank
	(1)	(2)	(3)
Patent eigencentrality	0.091*** (0.028)	0.041*** (0.014)	0.105** (0.043)
N (Obs = inventor–exhibitor)	560	560	560
R^2	0.015	0.015	0.011
Std. β	0.122	0.124	0.105

The table documents that British inventor–exhibitors were more likely to exhibit in central categories at the London World’s Fair 1851 if they had previously patented in more central technology categories.

Observations are individuals who both patented and exhibited in the Crystal Palace exhibition. For patents, eigencentrality is calculated as left eigenvector in line with theory and argument. For exhibits, eigencentrality is calculated as undirected eigenvector since the exhibition network is undirected. For inventor–exhibitor with multiple patents or exhibits, eigencentrality is averaged. OLS regressions, robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

tion network is similar to the patent-based innovation network. The basic assumption is that inventor–exhibitors patented technologies similar to those they exhibited. Under this assumption, we expect that if the exhibit-based innovation network is similar to the patent-based one, then inventors with more central patents within the patent network should also have more central exhibits within the exhibit network.

It is also informative to study how French and British exhibitors differed in terms of the types of technologies that they exhibited at the Crystal Palace. As a first look, Table E.6 describes the ratio of British to French exhibitors in the different exhibition data technology categories. For example, the ratio of British to French inventors was particularly high in areas such as *Machines for direct use* (including steam engines and locomotives) or *Military and Maritime Equipment* (including shipbuilding). In these categories, British inventors were also dominant in the patent data. Conversely, French exhibitors were relatively more prominent in technologies related to *Wearing Apparel*, *Food Processing*, and *Paper and Books*. In these technologies, French inventors also filed relatively more patents.

Table E.6: Distribution of British and French exhibits

Exhibition technology category	British	French	B/F ratio
Construction	196	8	24.50
Military and maritime equipment	349	30	11.63
Machines for direct use, engines, carriages, railway	388	42	9.24
Tapestry, lace, embroidery	283	32	8.84
Agricultural machinery	255	29	8.79
Mining and Minerals	357	43	8.30
General hardware, incl. locks and grates	610	78	7.82
Manu. of mineral subst. incl. tiles, cement, bricks	119	18	6.61
Furniture	327	59	5.54
Jewelry	113	21	5.38
Manu. of an. and veget. subst. incl. rubber, brushes	127	27	4.70
Misc. manufactures and small wares	258	61	4.23
Glass	83	23	3.61
Woven and felted fabrics	96	27	3.56
Leather, boots and shoes	271	78	3.47
Woollen and worsted	322	94	3.43
Veg. and animal substances used in manufacturing	136	41	3.32
Cotton fabrics and dyes	62	19	3.26
Scientific instruments	552	181	3.05
Clothing	223	76	2.93
Flax and hemp	72	25	2.88
Machines and tools for manufacturing	219	93	2.35
Cutlery and hand tools	39	21	1.86
Paper and bookbinding	154	88	1.75
Chemicals	129	77	1.68
Porcelain, china, earthenware	60	37	1.62
Food processing	123	77	1.60
Silk and velvet	79	108	0.73
Mixed fabrics, incl. Shawls	0	54	0.00

The table describes the number exhibits by British and French exhibitors in each technology in the exhibitions data as well as the ratio of British to French exhibits. The table is sorted by the ratio, with the “most British” technology categories appearing at the top and the “most French” categories appearing at the bottom.

F Adjusting for patent quality

This appendix describes how we adjust for patent quality in the main text, and documents the robustness to alternative ways of adjusting for patent quality. Quality adjustments are important because patent counts frequently miss variation in the quality of the innovation represented by each patent. This is a common concern in studies relying on patent data, and thus researchers routinely adjust patent counts using quality measures such as patent renewals (Schankerman and Pakes, 1986) or patent citations (Trajtenberg, 1990; Harhoff et al., 1999). In our historical setting, no such measure is systematically available for both countries. However, we have alternative measures for Britain and France that, while different from each other, can be used to roughly account for innovation quality.

F.1 Historical measures of patent quality

For British patents, we use measures provided by two standard sources, Nuvolari and Tartari (2011a) and Nuvolari et al. (2021), that are similar to patent citations. The basis of their quality measures is a list of historical citations to patents from the engineering–technical and legal literatures as of 1852. Digitized by Nuvolari and Tartari (2011a), this list was originally compiled by the head of the British Patent Office, Bennet Woodcroft; they thus refer to the patent quality index built from this source as the Woodcroft Reference Index (WRI). In this list, each patent is cited at least once, up to a maximum of 23 citations.⁵⁷

Nuvolari et al. (2021) update and augment the index from Nuvolari and Tartari (2011a) by collecting, from a set of modern reference works on historical technology, modern citations to the patent and mentions of the inventor. This augmented patent quality index is referred to as the Bibliographic Composite Index (BCI).⁵⁸ Nuvolari et al. (2021) argue that the BCI provides a more accurate measure of patent quality than the WRI index, so we use it as our primary indicator of British patent quality, though we also test robustness using just the WRI. The left-hand panel of Figure F.3 presents a histogram for the BCI patent quality measure.

For French patents, we have an indication of patent quality that comes from the choice of patent length and the payment of fees to keep patents in force. The assumption here is that inventors only chose long patent lengths and paid additional fees for higher-quality inventions. To be specific, at the time of filing, French inventors could choose between three different patent lengths: 5, 10, or 15 years. Longer patents were more expensive: A 5 year patent cost 300 Francs, a 10 year patent 800 Francs, and a 15 year patent 1500 Francs.

⁵⁷We use the citation counts for this index and do not remove fixed effects as in Nuvolari and Tartari (2011a).

⁵⁸We focus on citation counts, summing historical and modern citations, since citation counts naturally have a positive scale. This differs from Nuvolari et al. (2021), who propose an index of all three measures that is centered on 0 and includes negative values. We thus also use the first principal component (rescaled into the positive domain) of historical and modern citations and inventor mentions for robustness checks.

In addition, inventors only had to pay half of the initial fee up front. The remainder was due within six months. Thus, a second measure of patent quality comes from whether the inventor paid the second half of the fee to keep the patent in force beyond one year. A third indication of patent quality comes from the payment of fees to make additions to the patent specification. Presumably, this was only worthwhile for patented invention that proved to be more successful. Together, these different fees partition patent into eight bins: 5 or 10 or 15 year terms, each of which could be renewed or not renewed, and with or without additions. The right-hand panel of Figure F.3 presents a histogram for the renewal-fee based patent quality measure for France. Note that there are really just 8 fee levels here, though it appears that there are more because some fee categories are split across multiple bins. However, some of these bins are fairly sparsely populated, so for our analysis we will combine them into five bins of increasing cost.

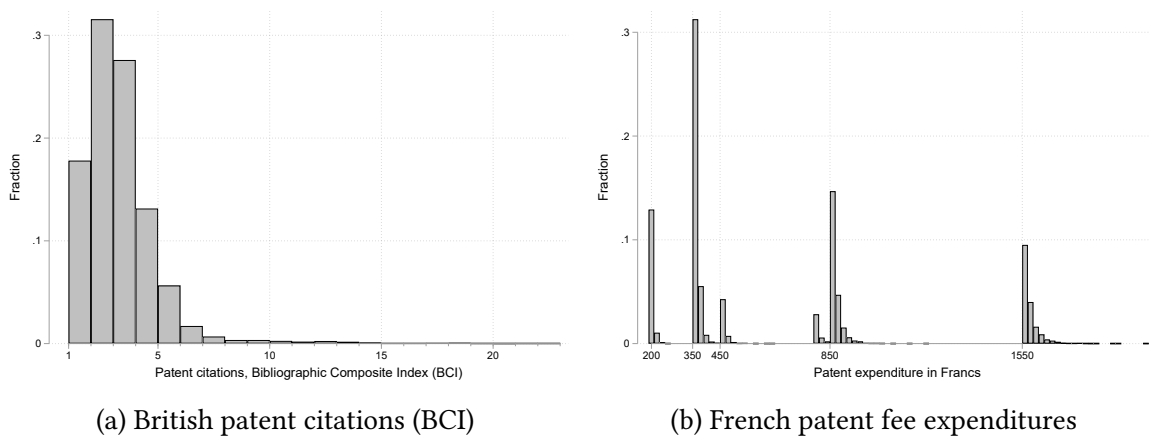


Figure F.3: Raw quality measures in Britain and in France

Table F.7 depicts the five quality bins that arise naturally from the various fees paid for patents in France. To generate a correspondence between the French and British patent quality measures, we have mapped the different levels of British patent citations to these bins, as described in the table, in order to generate similar distributions across the various patent quality levels. We use the five quality levels as cardinal measures of relative bin quality in the baseline, but results are very similar if we instead use the average monetary expenditures in the French bins as scale of relative bin quality. Using the bins shown in the table, and applying the same cardinal values to each bin, generates a comparable patent quality measure that can be applied to patents of both countries. The key assumption here is that the underlying distributions of patent quality are similar.

F.2 Alternative quality adjustments

An alternative way of mapping the distributions into the same scale is μ - σ -equating. Specifically, we can compute the means (μ) and standard deviations (σ) of the British and

Table F.7: Patent quality bins by British citation and French expenditure characteristics

Level	British quality	Cdf	French quality	Cdf
1	one cite	18%	5 years not renewed	14%
2	two cites	49%	5 years renewed	52%
3	three cites	77%	10 years not renewed, 15 years not renewed, 10 years renewed no addition	75%
4	four cites	90%	10 years renewed with addition, 15 years re- newed no addition	92%
5	five or more cites	100%	15 years renewed with addition	100%

French quality measures and then express the French quality index QI_{FR} of patent p in terms of the UK quality distribution:

$$QI_{p,FR}^{UK} = \mu_{UK} + \sigma_{UK}(QI_{p,FR} - \mu_{FR})/\sigma_{FR}$$

We apply this to the universe of patents from 1791 to 1844, the period when patents are available in both countries in order to obtain an alternative comparable measure of patent quality. As with our preferred measure, the key assumption underlying this alternative measure is that the mean and standard deviation of the underlying distributions of patent quality in the two countries are similar.

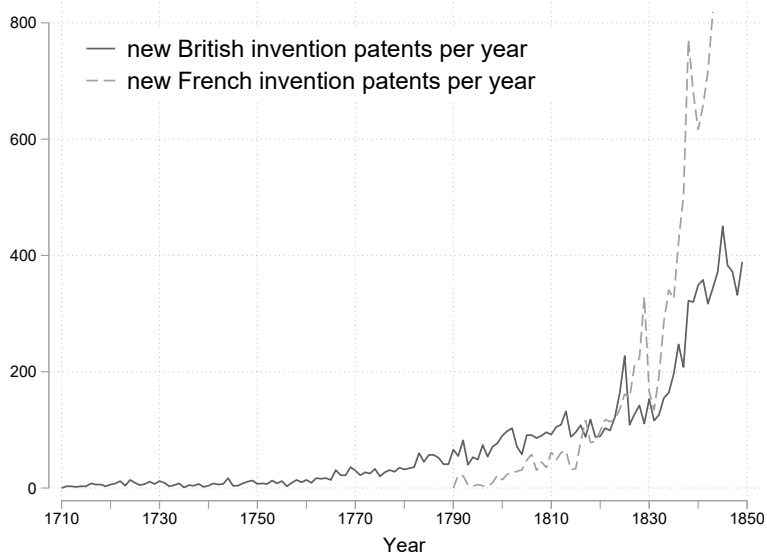
F.3 Descriptive patterns after adjusting for patent quality

Figure F.4 compares the time-series of raw patent counts in Britain and France (top panel) to the quality adjusted time-series (bottom panel). We can see that the main effect here is to bring the growth of British quality-adjusted patents closer to France in the bottom panel compared to what we see in the raw patent counts in the top panel. This must reflect that later British patents tended to be of higher quality than early patents.

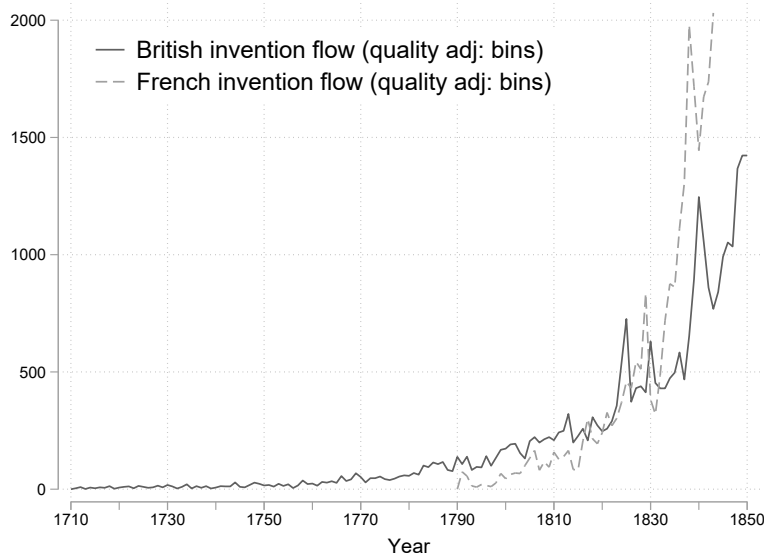
F.4 Robustness to alternative quality adjustments

Table F.8 presents results using the same specification applied in Table 2 in the main text, but with a variety of alternative ways of adjusting for patent quality. Column 1 repeats the baseline for reference, which used the British BCI citation measure and bins both British and French patents into five quality levels. Column 2 applies a different cardinal scale by weighting the five quality levels with the French average patent expenditures.⁵⁹ Column 3 applies the “ μ - σ -equating” approach described in subsection F.2. Column 4 uses the same quality measure as in Column 3, but uses British patents from 1791 instead

⁵⁹We normalize the scale by setting the lowest level to one.



(a) Patent counts



(b) Adjusted via quality bins

Figure F.4: Invention patents over time in Britain and France

Table F.8: Patent quality: Further robustness for main estimation

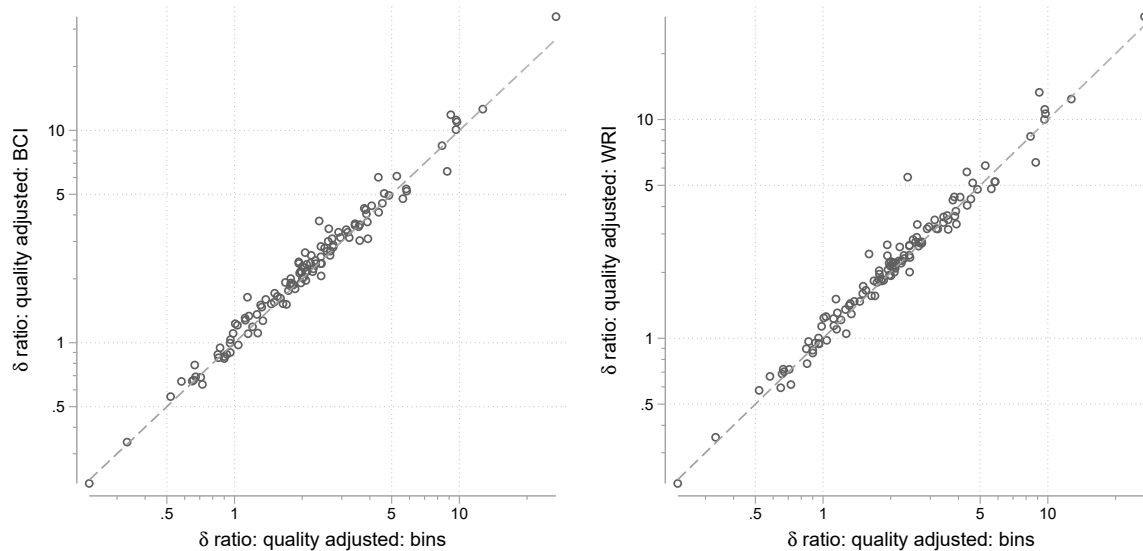
	Dep var: $\Delta \ln q_{c,i,t}$						
	(1) BCI bin (baseline)	(2) BCI bin, expend.	(3) BCI $\mu-\sigma$	(4) BCI $\mu-\sigma$, q_0^{UK} 1790	(5) WRI, binned	(6) WRI $\mu-\sigma$	(7) PCA $\mu-\sigma$
$\ln \chi_{t-1}$	0.102*** (0.011)	0.105*** (0.012)	0.112*** (0.015)	0.136*** (0.016)	0.104*** (0.011)	0.112*** (0.015)	0.115*** (0.015)
$\ln \chi'_{t-1}$	0.004 (0.002)	0.005* (0.003)	0.004 (0.003)	-0.008* (0.004)	0.003 (0.002)	0.003 (0.003)	0.003 (0.002)
$\ln q_{t-1}$	-0.106*** (0.011)	-0.109*** (0.012)	-0.116*** (0.015)	-0.129*** (0.015)	-0.107*** (0.011)	-0.115*** (0.014)	-0.118*** (0.015)
N	9325	9325	9325	9214	9325	9325	9325
Num. FE (ξ_i)	240	240	240	240	240	240	240
F -stat	26.8	23.9	24.8	21.7	26.8	25.2	26.1
implied λ	0.112	0.116	0.123	0.137	0.113	0.122	0.126
implied ι	0.963	0.955	0.964	1.000	0.968	0.972	0.972

The table documents the robustness of our main estimates with respect to different ways of quality adjusting patents. Column 1 repeats the baseline results for reference, which adjust for quality by binning British BCI citations and French patent expenditures into five quality levels. Column 2 applies a different cardinal scale by weighting the five quality levels with the French average patent expenditures. Column 3 uses the British BCI citations but maps French patent expenditures onto them via $\mu-\sigma$ -equating. Column 4 uses the same quality measure as in Column 3, but uses British patents from 1791 instead of 1761. Column 5 uses the British WRI citations for binning. Column 6 uses the British WRI citations and maps French patent expenditures onto them via $\mu-\sigma$ -equating. Column 7 uses the British principal component of all three Nuvolari et al. (2021) indicators and maps French patent expenditures onto them via $\mu-\sigma$ -equating.

of 1761. Column 5 returns to our preferred “binned” approach to relating the two patent quality measures, but uses the WRI quality index in place of the BCI index to measure the quality of British patents. Column 6 uses the British WRI citations and maps French patent expenditures onto them via $\mu-\sigma$ -equating. Column 7 uses the British principal component of all three Nuvolari et al. (2021) indicators and maps French patent expenditures onto them via $\mu-\sigma$ -equating. All specifications give quite similar results.


In Figure F.5, we compare the ratio of knowledge access efficiencies in the two countries (δ_{UK}/δ_{FR}) obtained from our baseline method for quality adjustment (x-axis) to those obtained when with alternative methods of adjusting (or not adjusting) for patent quality (y-axis). Panel F.5a compares our baseline against the BCI citation measure, onto which we map French patent expenditures via $\mu-\sigma$ -equating. Panel F.5b compares our baseline against the WRI citation measure (also using $\mu-\sigma$ -equating). Panel F.5c compares the baseline against the case of using raw patent counts i.e. not adjusting for quality at all. All

alternative methods yield very similar knowledge access ratios compared to the baseline.



(a) BCI $\mu-\sigma$ vs. baseline

(b) WRI $\mu-\sigma$ vs. baseline



(c) Patent counts (no adj.) vs. baseline

Figure F.5: Robustness of δ ratios to adjusting for patent quality

G Analysis using Google N-grams data

This appendix provides additional details on the construction of the Google n-grams data used in Section 8.1 as well as some additional results.

G.1 Data construction

Our goal is to construct a dataset that provides a country-by-technology measure of knowledge access in Britain and France. We start with the patent titles, which are already mapped to concordant technology categories. We use these titles for patents associated with each technology category to obtain a set of keywords that are strong predictors of patents falling into each category. To begin, we parse the titles into vectors of unigrams or bigrams.⁶⁰ We then run univariate patent-level regressions looking at how well each unigram or bigram predicts each technology category.⁶¹ For each gram, we retain the technology category that received the highest t-statistic, and additionally applied a cutoff t-statistic of $t \geq 5$ for unigrams and a cutoff of 3 for bigrams. This identifies the candidate set of potential keywords.

Next, we feed the candidate keywords from each language into a Large Language Model that translates them into the other language.⁶² We then take these translated terms back to the patent data and look at how predictive they are of the technology category. Only when both a keyword and its translation are strongly predictive of the same technology category, using the same criteria as above, we consider that a useful keyword pair.

To give a sense of the results of this exercise, Table G.9 provides the top-10 keyword pairs for two technology categories, Spinning, and Agriculture. We can see that, in general, this exercise delivers reasonable results. It is worth noting that there are some surprises in the list. “Drill,” for instance, shows up as an agricultural technology in the list below. A modern reader might instead have thought of it as a tool, forgetting that seed drills were one of the most important agricultural technologies of the period. Our approach tells us that, in fact, the term was actually much more likely to appear in patents related to agriculture than in, for example, machinery and tools. While there are clearly imperfections in the keyword lists that our method produces, we believe that in general it is effectively identifying keywords that are strongly associated with particular technology types.

Once we have our keyword pairs, we feed them into Google N-grams. For French

⁶⁰We exclude frequent stop words, such as determiners, prepositions, and linking verbs.

⁶¹We univariate regressions here, rather than including all grams in a single regression and applying LASSO-style methods, because we are interested in the predictive potential of each keyword, not the marginal additional predictive power conditional on other keywords.

⁶²We used OpenAI’s GPT 5.1, v.2025-11-13, reasoning = none, temperature = 0.2. The model was assigned the role of a bilingual (English/French) terminology expert and was instructed to provide one or multiple equivalent translations as output, while preserving grammatical structure. Input terms were provided in batches, one for each technology category. The category name was given as context.

Table G.9: Examples for equivalent keyword pairs

English term	French term
<i>Technology category: Spinning</i>	
bobbins/spools	bobines
carding	cardage
carding engine	machine à carder
carding engines	machines à carder
carding wool	cardage de la laine
cardings/cards	cardes
combed	peigné/peignée/peignées/peignés
combing wool	peignage de la laine
combing/heckling	peignage
cotton	coton
<i>Technology category: Agriculture</i>	
cattle	bétail
food	aliment/nourriture
oats	avoine
preparing food	préparation des aliments
agricultural	agricole
cultivation/culture/tillage	culture
distributing	distribuant
drill	semoir
fertilizing	fertilisation
harrow	herse

The table lists examples of equivalent keywords pairs that our prediction cum translation algorithm generates from patent titles.




Figure G.6: Scatter plot, corresponding to Column 1 of Table 5

terms, we use the French corpus, and for English words we use the British English corpus, restricting our search to material produced between 1750 and 1850. From N-grams, we recover the number of documents associated with each keyword.

G.2 Additional results

Table 5 in the main text presents our baseline analysis of the Google N-grams data. In this subsection, we provide some additional results and robustness checks supporting those in the main text.

In Figure G.6, we present a scatterplot corresponding to the specification shown in Column 1 of Table 5 in the main text. The main take-away from this plot is that the results are not simply driven by one or two outlier technology categories.

Table G.10 shows results estimated at the keyword-pair level, rather than collapsing keyword pairs to the level of technology categories. We weight these regressions so that keyword pairs from each technology category have equal weight, i.e., so that technology categories with many keyword pairs do not get more weight than those with relatively few keyword pairs. Results are qualitatively similar to those in the main text, but smaller in magnitude and less precisely estimated. This is related to the fact that, because the dependent variable ($\ln \delta$ ratio) only varies at the technology category level, the within-category variation in the independent variable (\ln document ratio) acts similarly to classical measurement error.

Table G.10: Results estimated at the keyword-pair level

	Dep var: $\text{Ln } \delta^B / \delta^F$ ratio					
	1750–1800		1770–1830		1800–1850	
	(1)	(2)	(3)	(4)	(5)	(6)
Ln document ratio	0.048*		0.045*		0.044	
	(0.025)		(0.025)		(0.028)	
Ln English documents		0.052*		0.043		0.040
		(0.027)		(0.026)		(0.030)
Ln French documents		-0.048*		-0.046*		-0.047
		(0.026)		(0.025)		(0.028)
<i>N</i> (Obs = gram)	1009	1009	1061	1061	1063	1063
Clusters (category)	113	113	113	113	113	113
R^2	0.008	0.008	0.006	0.006	0.004	0.004

OLS regressions with sampling weights adjusting for the number of grams per technology category. Standard errors clustered at technology category level in parentheses.

G.3 Using the N-grams data to evaluate the importance of accounting for the innovation network

We can also use the data from Google N-grams to assess the value of using the structure of our model to back out theoretically grounded knowledge access cost parameters. Our knowledge access cost parameters are derived from estimates of research effort obtained from the patent data using Eq. 13. This estimating accounts for how innovation rates are affected by spillovers through the innovation network. Suppose that, instead, we simply treated the ratio of patents in each technology category in each country as a measure of the knowledge access costs. Would this provide an equally good measure of knowledge access? To answer this question, we can study how well these two alternatives predict the availability of documents in the Google N-grams corpus.

Columns 1 of Table G.11 compares the ratio of estimated knowledge access cost parameters in Britain vs. France to the ratio of documents in the two countries associated with each technology type. As shown in the main text, this association is strong. Column 2 instead uses the ratio of patents in a technology category in the two countries. Here, we also see a positive association, but slightly weaker than when using our knowledge access cost parameter ratio. In Column 3, we run a horse-race between these two alternatives. We find that our estimated knowledge access cost parameter ratio, while losing significance, remains predictive of the document ratio, while the ratio of patents has no predictive power. The same is true for other time periods, in Columns 4 and 5. Thus, it appears that by using the structure of the model to back out knowledge access costs, we obtain values

that are closer to the actual availability of knowledge than if we had simply relied on patent ratios. In other words, using the structure of the model and accounting for the innovation network appears to be important for obtaining valid knowledge access costs estimates.

Table G.11: Testing for the importance of accounting for the innovation network

	Dep var: Ln document ratio				
	Period 1750–1800			1770–1830	1800–1850
	(1)	(2)	(3)	(4)	(5)
Ln δ^B/δ^F ratio	0.132** (0.061)		0.110 (0.109)	0.091 (0.085)	0.088 (0.077)
Ln patent share ratio		0.089 (0.060)	0.022 (0.099)	0.010 (0.078)	0.004 (0.071)
<i>N</i> (Obs = category)	113	113	113	113	113
R^2	0.040	0.031	0.041	0.033	0.031

We calculate the (ln) patent share ratio using the baseline quality adjusted patents. OLS regressions, robust standard errors in parentheses.

H Estimating the technology sharing parameter (η)

This appendix describes how we estimate the technology sharing parameter η .

H.1 Deriving the estimating equation

Given the assumption on technology sharing in intermediate goods production, model equation 2 becomes

$$\begin{aligned}\ln y_{kt} &= \sum_i \alpha_{ik} \left[\int_0^1 \eta \ln(q_{it}(v) l_{kit}(v)) dv + \int_0^1 (1 - \eta) \ln(q'_{it}(v) l_{kit}(v)) dv \right] \\ &= \sum_i \alpha_{ik} \left[\eta \ln(q_{it} l_{kit}) + (1 - \eta) \ln(q'_{it} l_{kit}) \right]\end{aligned}$$

Writing out the assumption that a share domestic varieties are partly under monopoly, while foreign varieties are produced with domestic labor under perfect competition, we obtain,

$$\begin{aligned}\ln y_{kt} &= \sum_i \alpha_{ik} \left[\eta \left(\int_0^{1-e^{-\phi}} \ln(q_{it} l_{kit}^{\text{mon}}) ds + \int_{1-e^{-\phi}}^1 \ln(q_{it} l_{kit}^{\text{pc}}) ds \right) + (1 - \eta) \ln(q'_{it} l_{kit}^{\text{pc}}) \right] \\ &= \sum_i \alpha_{ik} \left[\eta \left(\int_0^1 \ln q_{it}(v) + \int_0^{1-e^{-\phi}} \ln l_{kit}^{\text{mon}}(v) dv + \int_{1-e^{-\phi}}^1 \ln l_{kit}^{\text{pc}}(v) dv \right) \right. \\ &\quad \left. + (1 - \eta) \ln(q'_{it} l_{kit}^{\text{pc}}) \right] \\ &= \sum_i \alpha_{ik} \left[\eta \left(\ln q_{it} + (1 - e^{-\phi}) \ln\left(\frac{1}{1 + \lambda} \alpha_{ik} \beta_k \bar{L}_t\right) + e^{-\phi} \ln(\alpha_{ik} \beta_k \bar{L}_t) \right) \right. \\ &\quad \left. + (1 - \eta) (\ln q'_{it} + \ln(\alpha_{ik} \beta_k \bar{L}_t)) \right] \\ &= \sum_i \alpha_{ik} \left[\eta \left(\ln q_{it} + (1 - e^{-\phi}) \ln \frac{1}{1 + \lambda} + \ln \alpha_{ik} \beta_k \bar{L}_t \right) \right. \\ &\quad \left. + (1 - \eta) (\ln q'_{it} + \ln(\alpha_{ik} \beta_k \bar{L}_t)) \right] \\ &= \sum_i \alpha_{ik} \left(\eta \ln q_{it} + (1 - \eta) \ln q'_{it} - \eta \ln(1 + \lambda) \left(1 - \frac{q_{it}}{r_{it} \chi_{it}}\right) + \ln \alpha_{ik} \beta_k \bar{L}_t \right) \\ &= \eta Q_{kt} + (1 - \eta) Q'_{kt} - \eta \Phi_{kt} + \ln \bar{L}_t + A_k\end{aligned}$$

where in the last line, we defined $Q_{kt} \equiv \sum_i \alpha_{ik} \ln q_{it}$, $Q'_{kt} \equiv \sum_i \alpha_{ik} \ln q'_{it}$, $\Phi_{kt} \equiv \sum_i \alpha_{ik} \ln(1 + \lambda) \left(1 - \frac{q_{it}}{r_{it} \chi_{it}}\right)$, and $A_k \equiv \sum_i \alpha_{ik} \ln \alpha_{ik} \beta_k$, and used that $\sum_i \alpha_{ik} = 1$.

Next, we take the long difference between two periods, $\tau = T$ and $\tau = t_0$:

$$\ln y_{kT} - \ln y_{kt_0} = \eta Q_{kT} + (1 - \eta) Q'_{kT} - \eta Q_{kt_0} - (1 - \eta) Q'_{kt_0} - \eta \Phi_{kT} + \eta \Phi_{kt_0} + \ln \bar{L}_T - \ln \bar{L}_{t_0}$$

This equation simplifies as follows: First, note that $\Delta Q_{k,T-t_0} = \sum_i \alpha_{ik} \ln q_{iT} - \sum_i \alpha_{ik} \ln q_{it_0} = \sum_i \alpha_{ik} \ln \frac{q_{iT}}{q_{it_0}}$. Second, $\ln \bar{L}_T - \ln \bar{L}_{t_0}$ does not vary across k and goes into the constant. Third, to compute the Φ_k terms, which capture the market entry channel effect on output of the arrival rate of ideas, we would need time varying measures of \hat{r} . But this we do not have—we have estimated an average value over the period 1800–1843. With the average \hat{r} , we could calculate an average Φ_k term, but this would drop out when taking differences. Thus, we drop the Φ_k terms as an empirical approximation made for tractability. In sum, the equation becomes

$$\Delta \ln y_{T-t_0,k} = \eta \Delta Q_{T-t_0,k} + (1 - \eta) \Delta Q'_{T-t_0,k} + \text{const} . \quad (18)$$

H.2 Data

We digitize panel data on historical output at the industry–by–decade level for France from the authoritative source, Markovitch (1965). The production data comes in two forms: One is a quantity index, the other is the value of production.⁶³ The data are reported as decadal averages. The first two decades are 1781–90 and 1803–12; from 1815–24 up until the mid 20th century, the data are reported without gaps, with decades centering on the round years 1820, 1830, and so on.⁶⁴

We hand link the 23 French industries to the British input-output industries for which we already constructed industry weights by technology category, α_{ik} , ending up with data for 16 industries: *Gas, Coal, Building Materials, Metal production, Metal products, Chemicals, Rubber, Tobacco and Matches, Textile manufacturing, Leather, Paper and Cardboard, Press and Publishing, Construction and Public Works, Food, Clothing and Fabric Finishing, and Wood and Furniture*.⁶⁵ As non-zero output is first reported for *Gas* and *Rubber* in the 1840 (1835–44) decade, our balanced estimation panel includes 14 industries.

⁶³The quantity index is given in Table 1 (Markovitch, 1965), setting $I_{1938} = 100$. The value of production in period prices is given in Table 3, in million Francs (Markovitch, 1965, p.). We cannot back out the actual production levels since neither the quantities of 1930 nor the period prices are reported.

⁶⁴The early data gaps are due to the French Revolution and the Wars of the Sixth and Seventh Coalition.

⁶⁵Three French industries are not included in the British input-output data by Thomas (1984): *Electricity, Petroleum and fuel oil*, and the *Glass industry*. Three more French industries are lost due to the need for aggregation: *Extraction of metallic ores* and *extraction of various ores* to the British *Other mining* (i.e. not coal mining) industry, *Extraction of construction materials* and *Ceramics and manufacturing of construction materials* to the British *Building materials* industry, and *Chemical Industries* and *Lipids* to the British *Chemicals* industry. Finally, we cannot map patents onto the *Other mining* industry because of a lack of occupation–patent links in Britain during this period.

H.3 Estimation

The empirical counterpart of Equation 18 is

$$\Delta \ln y_k = \gamma_0 + \gamma_1 \Delta Q_k + \gamma_2 \Delta Q'_k + \varepsilon_k, \quad (19)$$

where the unit of observation k is a French industry; the outcome $\Delta \ln y$ is the growth in industrial output between t_0 and T ; and the main explanatory variables are ΔQ and Q' , the changes in the weighted sums of domestic (French) and foreign (British) knowledge stocks between t_0 and T . The constant γ_0 controls for potential changes in population. Our main goal is to estimate $\eta = \gamma_1 = 1 - \gamma_2$. We will use this structural constraint to help identifying η . As study period, we choose $T = 1840$ and $t_0 = 1810$, which is the longest difference for which we have both patents and output available in both countries.

Table H.12 reports estimates of the technology sharing parameter using various specifications and data. In the odd columns, we estimate simple OLS regressions as a reference that uses only variation in changes of the domestic knowledge stock. Column 1 suggests that $\eta = 0.685$, using the baseline method to compute q . In the even columns, we implement the structural equation by controlling for the changes of the foreign knowledge stock and constraining coefficients to sum to one across the domestic and foreign terms. Column 2 suggests that $\eta = 0.622$. We observe a similar pattern across the baseline and robustness specifications: The implied estimates of η are generally larger for the simple OLS, but broadly similar to the structural constrained regressions.

The result is essentially unchanged if we use another method of quality adjusting patent counts, as shown in columns 3 and 4 where we instead use the BCI ($\mu - \sigma$) quality adjustment. But results are influenced by how we define the initial knowledge stocks q_0 . In columns 1 through 4, we count British patents q^{UK} from 1761—approximately when the Industrial Revolution begins—and French patents from 1791—when the French patent law was enacted. This may be giving an unfair advantage to the British Q' , since additions to the French knowledge stock were probably not zero in the period 1760–1790. We probe robustness by equalizing the initial knowledge stocks q_0 across countries. In columns 5 and 6, we assume that the French q^{FR} equaled that of Britain q^{UK} during the period 1760–1790, which produces a substantially larger estimate for $\eta \approx 0.8$. In columns 7 and 8, we assume that the knowledge stock was zero before 1790 across all technologies in both countries, which gives a smaller estimate for $\eta \approx 0.56$.

Table H.12: The importance of technology sharing for production (η estimation)

	Dep var: Δy (1840–1810) (industrial output growth)							
	Baseline		BCI ($\mu-\sigma$)		q_0 1760		q_0 1790	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	cns reg	OLS	cns reg	OLS	cns reg	OLS	cns reg
ΔQ (1840–1810)	0.685 (0.210)	0.622 (0.162)	0.675 (0.201)	0.616 (0.137)	0.888 (0.276)	0.801 (0.211)	0.685 (0.210)	0.561 (0.165)
$\Delta Q'$ (1840–1810)		0.378 (0.162)		0.384 (0.137)		0.199 (0.211)		0.439 (0.165)
N (industries)	14	14	14	14	14	14	14	14
R^2	0.38		0.36		0.45		0.38	
$Root\ MSE$	0.32	0.31	0.32	0.31	0.30	0.29	0.32	0.30
implied η	0.685	0.622	0.675	0.616	0.888	0.801	0.685	0.561

OLS regressions (odd columns) and constrained LS regressions such that $\gamma_{\Delta Q} + \gamma_{Q'} = 1$ (even columns). The outcome is the growth of French industrial production between the 1840 decade (1835–44) and 1810 decade (1803–12). We measure Q , Q' as of 1843 and 1810, respectively. *Baseline* (columns 1, 2): We calculate knowledge stocks q as in the baseline, which adjust quality via five quality bins and starts counting British patents q^{UK} from 1760 and French patents from 1790. *BCI* (columns 3, 4): We adjust quality via the BCI method, which uses patent citations for q^{UK} and maps French patent expenditure onto them via $\mu-\sigma$ -equating. *q_0 1760* (columns 5, 6): We again use the baseline quality bins, but now let both countries' q_0 start in 1760 by assigning the British q^{UK} between 1760 and 1790 to France. *q_0 1790* (columns 7, 8): Using the baseline quality bins, we let both countries' q_0 start in 1790 by excluding British patents before 1790, i.e. setting q^{UK} to 0 before 1790. Note that column 7 replicates column 1, because the French knowledge stocks remain unaffected. All regressions include a constant (not reported). Robust standard errors in parentheses.

I Model simulation appendix

This appendix present additional results for our baseline simulation as well as a range of robustness exercises. We begin by providing some additional figures showing results from the baseline simulation presented in the main text.

I.1 Additional results for the baseline simulation

Figure I.7 plots output in our baseline simulations under either the close economy case (top panel) the spillovers case (middle panel) or with spillovers and technology sharing (bottom panel). It is interesting to see that output was initially lower in the UK, but growing more rapidly. This reflects both the fact that population was lower in the UK and that a greater share of labor was allocated to research as a results of the lower knowledge access costs.

Figure I.8 plots the level of researchers in the two economies. The level of research is higher in the UK as a result of lower knowledge access costs. In both economies, the level of research is increasing on the transition path and then levels off as the economy approaches balanced growth, after about 3-4 decades. Note that allowing knowledge spillovers across countries reduces the number of researchers in Britain and increases the number of researchers in France. This is because, in the spillovers case, British inventors are dependent on the lower French knowledge stock, while French inventors benefit from the higher British knowledge stock. Note also that technology sharing does not affect the level of researchers. In our model, technology sharing affects output but not researcher allocation. Figure I.9 plots the average arrival rate of technologies, which shows a pattern similar to researchers.

Figure I.10 plots the growth rate in the two economies under different scenarios. It is interesting to see that the simulation implies growth rates that are higher than those actually experienced by the economies we study. This should not be surprising given that our model does not incorporate any of the frictions that are likely to slow down the adoption of new technologies in real economies. This highlights the fact that our quantitative results should be interpreted as implied growth rate differences before the impact of any other frictions affecting the implementation of new technologies, such as those related to credit availability or concerns about the security of property rights, are accounted for.

I.2 Simulation robustness

Here, we study how differences in initial conditions affect our simulation results.

Note that the model's functional form implies that, with a common uniform initial knowledge stock q_0 , the level of q_0 does not influence results. To see this, note that a uniform q_0 factors out of the ratio q_{it}/χ_{it} in the first period. As a result, the effective

research effort r_{it} and the innovation arrival rate ϕ_{it} are completely independent of q_0 .

We thus consider the impact of allowing variation in initial technology levels across technology categories. To do so in a realistic way, we start with the q_0 values observed in Britain in 1790. We use values from 1790 because this is near the beginning of our study period and our goal here is to obtain realistic technology levels around the time of the onset of the Industrial Revolution. French q_0 s cannot be observed at such an early date, which is why we only use British values. We then add an arbitrary constant to ensure that no technology category begins the simulation with a technology level of zero, and use the resulting initial technology level as the starting conditions for both countries. To give a sense of the extent of variation in initial technology levels in this exercise, Figure I.11 provides a histogram of initial q_0 values when we add an arbitrary constant equal to two.

Figure I.12 presents results with these initial technology levels that vary across categories. In the top panel we use the British distribution in 1790 plus 2, while in the bottom panel we instead add 100. We can see that allowing some variation in initial technology level does impact the results, but the basic patterns remain very similar to those presented in our baseline simulation. Thus, allowing realistic variation in initial technology levels across categories also does not appear to substantially affect our simulation results.

I.3 Counterfactual simulations

This section presents additional results for a set of counterfactual exercises that can help us isolate different aspects of Britain's overall growth advantage. In our first exercise, we look at the impact of the innovation network. To do so, we consider a counterfactual where the network is set to the identity matrix, so that all innovation spillovers occur within the same technology category.

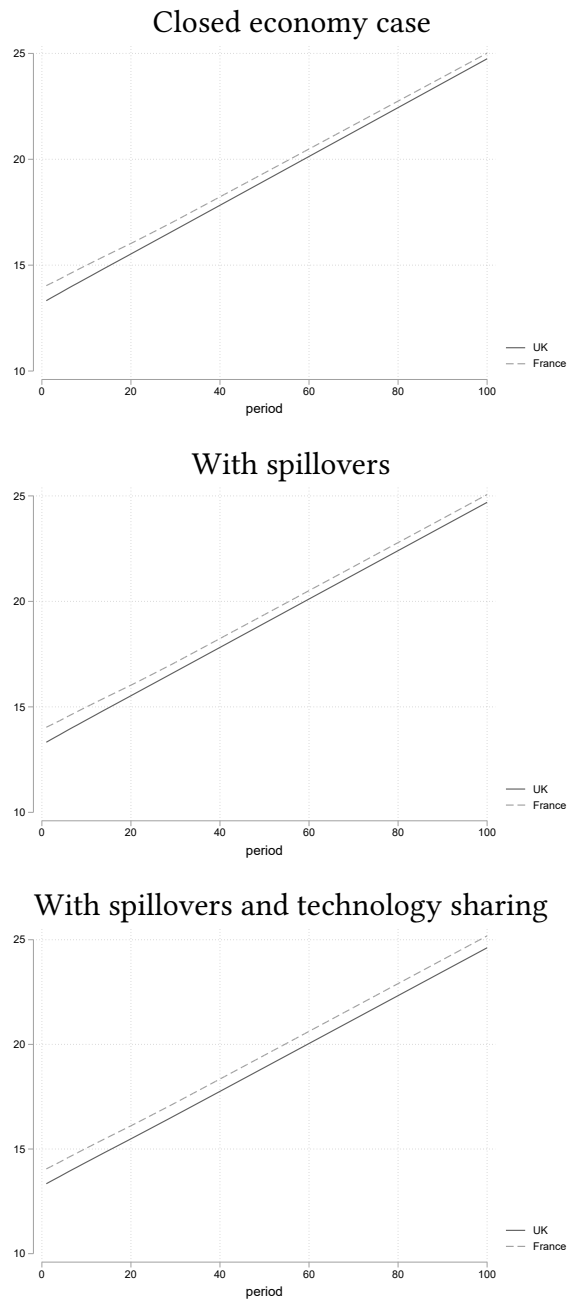
Figure I.13 plots the growth rate differences with the innovation matrix set to the identity matrix under the closed economy, cross-country knowledge spillovers, and technology sharing scenarios. The closed economy line, which corresponds to the dash-dotted line in Figure 6, shows a stable latent British growth advantage, which is more than 0.5 percentage points higher than the short-run British advantage of about 1 percentage point in the baseline. The knowledge spillovers line shows a similar latent British growth advantage for two decades, before it drops sizeably when France starts to catch-up rapidly over the next two decades. Yet, the British advantage remains higher than the baseline closed economy scenario for almost a century. The technology sharing line shrinks the British advantage, but it remains sizable in this scenario even after 50 years. These findings highlight two things. First, accounting for the innovation network has a meaningful impact on our quantitative results, in particular on the dynamics of the British growth advantage. Second, Britain's growth advantage over France was not driven by the way its knowledge access cost vector operated through the innovation matrix. If anything, the British distribution of knowledge access costs across technologies appears to have been disadvantageous relative

to the structure of the innovation network.

Note, however, that even without the innovation network, differences in the distribution of knowledge access costs across technologies can still matter for growth because of the way that different technologies fit into the production function in different output sectors (the α parameters) and because of the weight of those output sectors in consumption (the β parameters). To examine the impact of this channel, we conduct a second counterfactual where we eliminate any differences in the *distribution* of knowledge access costs across technology types in both countries, while still allowing differences in the overall *level* of knowledge access costs between the two countries. Specifically, we set knowledge access costs in each country equal to the weighted average knowledge access cost across all technology types, where the weights are based on the consumption weight of each technology category: $\bar{\delta}^c = \sum_i \delta_i^c (\sum_k \alpha_{ik} \beta_k)$.


Figure I.14 presents the results of this counterfactual exercise. Comparing these results to the previous counterfactual in Figure I.13, we can see that Britain's growth advantage has been reduced by over 80%. This shows that Britain's distribution of knowledge access costs was advantageous with respect to the importance of different technologies in overall consumption.

Figure I.7: Log output in our baseline simulation




The figure plots log output in the UK and France from our baseline simulation for the closed economy case (top panel), the case with cross-country knowledge spillovers (middle panel) and the case with both knowledge spillovers and technology sharing (bottom panel).

Figure I.8: Researchers in our baseline simulation




The figure plots the level of research in the two economies in our baseline simulation under alternative scenarios.

Figure I.9: Innovation arrival rate in our baseline simulation




The figure plots the average innovation arrival rate across all technology categories in the two economies in our baseline simulation under alternative scenarios.

Figure I.10: Growth rates in our baseline simulation




The figure plots the growth rate in the two economies in our baseline simulation under alternative scenarios.

Figure I.11: Histogram of initial technology levels using British q in 1790 plus 2




The figure plots a histogram of the initial technology levels in a simulation that begins with q^{UK} observed in 1790 plus an arbitrary constant of two.

Figure I.12: Simulations allowing variation in initial q_0 across technology categories




(a) British q in 1790 + 2




The figure plots the difference in growth rates in two economies, one with the British knowledge access costs and population and the other with the French knowledge access costs and population. In the top panel, the initial technology level in all industries is set to q^{UK} in 1790 plus 2, while in the bottom panel it is set q^{UK} in 1790 plus 100. The closed economy case does not allow knowledge spillovers or technology sharing. The knowledge spillover line allows knowledge spillovers using our estimate of ι . The technology sharing line additionally allows technology sharing using $\eta = 0.7$.

Figure I.13: Growth rate differences with the innovation network set to the identity matrix



The figure plots the difference in growth rates in two economies, one with the British knowledge access costs and population and the other with the French knowledge access costs and population. The parameters are the same as in our baseline simulation in the main text, except that we have set $\omega_{ij} = 0 \forall i \neq j$ and $\omega_{ij} = 1 \forall i = j$.

Figure I.14: Growth rate differences without variation in δ s across technology categories



The figure plots the difference in growth rates in two economies, one with the British knowledge access costs and population and the other with the French knowledge access costs and population. The parameters are the same as in our baseline simulation in the main text, except that (i) we have set $\omega_{ij} = 0 \forall i \neq j$ and $\omega_{ij} = 1 \forall i = j$ and (ii) all of the δ_i^c for each country have been set to the weighted average across all technology categories within that country, $\bar{\delta}^c = \sum_i \delta_i^c (\sum_k \alpha_{ik} \beta_k)$.