Innovation Networks in the Industrial Revolution *

Lukas Rosenberger
LMU Munich, CESifo

W. Walker Hanlon
Northwestern, NBER, CEPR

Carl Hallmann
Jane Street

October 2026

Abstract

We offer a new approach to understanding technological progress and economic growth in historical settings that combines modern growth theory, detailed micro data, and new methods for measuring key parameters using historical data. We extend the disaggregated growth model from Liu and Ma (2023) to incorporate research effort allocations that are determined in decentralized markets subject to country- and technology-specific innovation frictions, cross-country knowledge spillovers, and technology sharing. We apply this framework to study economic growth in Britain and France during the Industrial Revolution. We show how key parameters of the model can be estimated using historical data, even off the balanced growth path. These parameters allow us to simulate counterfactuals, which show that Britain enjoyed substantial growth benefits from having lower innovation frictions in technologies that served more demand and that generated more innovation spillovers for other technologies. Finally, we use our estimated parameters to evaluate Joel Mokyr's influential hypothesis that knowledge access costs played a key role in shaping innovation patterns during the Industrial Revolution.

*We thank Isabelle Hauf-Pisoni for outstanding research assistance. We thank Enrico Berkes, Davide Cantoni, David de la Croix, Guillaume Daudin, Quoc-Anh Do, Martin Fiszbein, Carola Frydman, Leander Heldring, Rick Hornbeck, Naomi Lamoreaux, Joel Mokyr, Bang Nguyen, Alessandro Nuvolari, Sebastian Ottinger, Monika Schnitzer, Claudia Steinwender, Uwe Sunde, Alex Trew, Fabian Waldinger, and seminar participants at Northwestern, LMU Munich, NBER Summer Institute, Chicago Fed, Mannheim, IMF, LEAP, UW Milwaukee, Peking University, Southern Denmark, Louvain, Bayreuth, HKU Quantitative History Webinar, "Productivity Revolutions" workshop Manchester, EEA Barcelona, EHES Vienna, and EHS Newcastle 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: Hanlon, whanlon@northwestern.edu; Rosenberger, lukas.rosenberger@econ.lmu.de.

1 Introduction

This paper offers a new approach for studying technological progress and economic growth in the past, combining modern growth theory, detailed micro data, and new ways of measuring key parameters in historical settings. Over the last decades, substantial progress has been made both in the development of disaggregated models of economic growth that lend themselves to quantitative application and in the digitization of detailed historical statistics to which these models can be applied. Despite this progress, these methods have not been widely applied to understand economic growth in the past, perhaps because important challenges remain in tailoring modern theories to the more limited data available in historical settings.

We begin by developing an endogenous growth model featuring multiple countries, multiple technology types and multiple production sectors. Our model starts from the framework of Liu and Ma (2023), in which technological progress in a particular technology depends on both research effort and technology spillovers across technology types through an "innovation network." We extend this framework in several ways. First, in our model, the allocation of research effort is determined in a decentralized market subject to country-technology-specific *innovation frictions*. These innovation frictions, the key primitives in our model, are wedges that represent barriers—such as the knowledge access costs emphasized by Joel Mokyr (Mokyr, 2002, 2005, 2009)—that make it harder for potential researchers in a country to work in a particular technological domain. Second, our model features multiple countries, which are linked by both knowledge spillovers, where ideas in one country facilitate innovation in another, and technology sharing, where technologies developed in one country are applied to production in another. Third, our model allows for time-varying differences in the institutional and business environment in each country, which affect the extent to which inventors can profit from new discoveries. Fourth, we decouple technologies from industries. This allows for the possibility that some technologies are more "general", i.e., used across a broader set of production sectors than others. It also facilitates a cleaner mapping from the model to the data, since such variation is a salient feature of the empirical setting that we consider.

We apply this framework to study technological progress and economic growth during one of the most important periods of economic history, the Industrial Revolution. Specifically, we are interested in one of the most enduring questions about this episode: why was Britain able to achieve more rapid technological progress and economic growth than other European countries during the early stages of the Industrial Revolution? This topic has attracted enormous attention, leading to 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, 2026), advantageous natural resources (Pomeranz, 2000; Fernihough and O'Rourke, 2021), and lower knowledge access costs (Mokyr, 2002, 2005, 2009).[1] Yet existing studies rarely provide quantitative estimates of the growth contribution of the various hypothesized advantages, making it difficult to assess how much different factors mattered.

To apply our framework to this historical setting, we begin by bringing together detailed data on patenting in Britain from 1700 to 1849 as well as in France, the natural comparison country (Crafts, 1977), from 1791 to 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] However, key pieces of information, such as systematic patent citations or R&D expenditures, were not collected during the Industrial Revolution.

Thus, bringing our theory to the data in this historical setting requires overcoming challenges not faced by studies focused on modern settings. First, in order to construct the innovation network, we need to measure spillovers across technology types. Previous studies in modern settings use patent citations, but systematic citation data are not available before 1945 (e.g., Nicholas, 2010).[4] 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. We validate this novel approach in two ways: using modern data, we show that our inventor-based method comes very close to replicating the innovation network obtained from patent citations, and using historical data on the full text of British patents from our study period, we show that the obtained network is similar to one based on patent text similarity. Developing and validating this new approach, which allows the study of innovation networks further back in time than is currently possible, is one contribution of our study.

Second, we need to estimate our model without information on research effort, which is typically measured using R&D expenditures. We show how the model can be used to back out research effort, as well as full vectors of country-by-technology innovation frictions, using only the available patent data. Importantly, the validity of our estimates

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

[2] Both 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.

[3] Of course, not every useful invention was patented, as Moser (2012) has shown. To account for this, we also replicate some of our key parameter measures using data from non-patented inventions exhibited at the 1851 Crystal Palace exhibition. These data generate patterns that are consistent with the patent data.

[4] Systematic patent citations exist only for the post-war period: for example, in the United States from February 1947 (Nicholas, 2010), in Germany from the late 1950s, in France under the law of 2 January 1968 (Loi n° 68-1), and in Britain under the Patents Act 1977.

does not require an assumption that we are on the balanced growth path. This matters for our application, since the economies we study were unlikely to be on a balanced growth path during the Industrial Revolution.

Using these inputs, together with the structure of the theory, we show how one can jointly estimate the parameters of the model: important macro parameters, the vectors of country-by-technology-specific innovation frictions, and time-varying parameters representing the institutional and business environment in each country. Consistent with available historical evidence, our estimated innovation frictions show that Britain maintained an advantage in mechanical technologies, such as cylinders, pistons, pulleys and joints, that played a central role in the Industrial Revolution. In contrast, our estimates indicate that France had a relative advantage in a number of consumer goods, such as bedding and bathing machines, cooking and baking, and, interestingly, aerial conveyances (i.e., air balloons). Our estimates of the institutional and business environment show that the British environment was improving relative to that in France between 1803 and 1815, corresponding to the disruptions in France that accompanied the Napoleonic Wars, but that the environment in France improved rapidly soon after the conflict ended. Our model also delivers estimates of the importance of both cross-country knowledge spillovers and (more tentatively) cross-country technology sharing during the Industrial Revolution. We find that both types of cross-country spillovers were substantial and contributed meaningfully to closing the growth gap between Britain and France, consistent with available historical evidence (Harris, 2017).

Our parameter estimates allow us to examine, quantitatively, whether Britain grew more rapidly during the early stages of the Industrial Revolution because of how British inventive effort was distributed across different technology types. Economic historians have long recognized that Britain maintained an advantage in certain technologies characteristic of the Industrial Revolution, such as steam engines and machine parts, as we find.[5] Continental inventors, however, were also innovative, but often in other technology types, such as chemicals and consumer products (Harris, 2017). Whether these differences matter for growth depends on whether some technologies were "better" for growth than others. Our framework formalizes two channels that might make some technologies contribute more to growth: knowledge spillovers through the innovation network, and the share of demand served by each technology.

To assess whether Britain benefited from having a relative advantage in the "right" technologies, we use our model and parameter estimates to conduct a series of counterfactual simulations. In particular, we simulate two economies, one endowed with Britain's innovation friction parameters and one with France's, that start from a common initial technology level in 1800. Our baseline simulation shows that the British economy's output per worker grew substantially faster in the first two decades of the Industrial Revolution.

[5]See, e.g., Landes (1969) and Allen (2009a).

By the middle of the nineteenth century, however, the simulated French economy closed this gap and caught up relative to the initial levels, as the environment improved following the end of the Napoleonic Wars. We then "turn off" two different channels of the model, the innovation network and differences in the demand shares, to assess how much Britain's relative advantage in certain technologies, and disadvantage in others, contributed to growth. We find that roughly half of Britain's growth advantage over the first two decades of the Industrial Revolution disappears when we turn off the heterogeneity in the growth contribution of different technologies. This tells us that the *distribution* of Britain's innovation frictions, relative to France, gave rise to a meaningful growth advantage. The remaining differences between Britain and France were due either to differences in the average level of innovation frictions between the two countries, or differences in their institutional environments, which cannot be separated in our empirical framework.

One reason why estimating the underlying innovation friction parameters is valuable is that they provide a theoretically grounded starting point for evaluating existing theories explaining Britain's advantage during the Industrial Revolution. To demonstrate how our estimates can be used to better understand the sources of growth during the Industrial Revolution, we consider one of the most influential theories of Britain's advantage: knowledge access costs. Mokyr argues that,

> "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."[6]

Despite the influence of this argument, quantitatively evaluating the importance of knowledge access for growth during the Industrial Revolution has proven to be challenging. To make progress on this problem, we compare our estimated country-technology innovation frictions to measures of knowledge access.

Our analysis begins with the hypothesis that one important determinant of knowledge access was the availability of printed material related to each technology type in a particular language. Put differently, we want to look for a systematic relationship between available printed material in Britain vs. France related to each technology type and our estimated innovation friction parameters. We start by using the patent titles to derive a set of equivalent keywords in English and French that are strongly predictive of patents being classified in both countries in 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 innovation friction parameters to the availability of printed material in each language. We find that technologies where we estimate lower

[6] Mokyr (2005), pp. 295–296. Italics in original.

innovation frictions in the UK relative to France have substantially more printed material in English relative to French published between 1750 and 1800, and vice versa. This shows that knowledge access costs, as represented by the availability of printed material in each country, are correlated with the innovation frictions recovered from our structural estimation, consistent with Mokyr's hypothesis.

This study contributes to existing work in both economic history and growth economics. Economic historians have produced 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; Caprettini, 2026) as well as a broader set of work studying historical knowledge access institutions (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 have emphasized the importance of knowledge access during the Industrial Revolution, most prominently Joel Mokyr, we are the first to produce a theoretically grounded empirical evaluation of this idea.

A second related strand of work applies (typically 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 micro data.

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). As discussed above, we start with the framework of Liu and Ma (2023), but extend their model in several useful dimensions.

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 (2011), 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 innovation frictions at the 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 parameters of the model. In Section 8, we conduct counterfactual simulations to evaluate some of the sources of Britain's growth advantage during the early nineteenth century. Finally, in Section 9, we use our estimated parameters, together with data from Google Ngrams, to study Mokyr's hypothesis that knowledge access costs played a key role in explaining innovation and productivity differences in Britain vs. France during the Industrial Revolution.

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. Though there is now evidence that sustained economic growth in Britain began earlier (Bouscasse, Nakamura, and Steinsson, 2025), it is clear that growth accelerated decisively around the late eighteenth or early nineteenth century, first in Britain and then across Western Europe. There is broad agreement that technological progress played a central role in this process.[7]

Britain's economic growth during this period has often been contrasted with that of France, a natural comparison country because it was also large, relatively advanced, and nearby.[8] Recent estimates of industrial growth suggest that industrial production in Britain grew at annual rates of 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).

Of course, there were many factors that could have influenced growth rates in Britain

[7] See, e.g., Landes (1969), Mokyr (1990), Mokyr (1999), Allen (2009a).

[8] Leading 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.

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, which ended in 1815.[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 Napoléon, 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.

In particular, we are interested in the observation that Britain and France excelled in different technological domains. Historical evidence suggests that Britain was a leader in many of the characteristic technologies of the first Industrial Revolution: steam engines, spinning machines, metalworking, railroads, etc. One reflection of this lead was the effort that Continental countries (as well as Americans) put into obtaining these British technologies (Harris, 2017). Another reflection was the leading position that Britain occupied in many of the industries reliant on these technologies. However, French inventors, as well as others on the Continent, were also innovative. "There is," wrote Harris (2017, p. 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.).[10] Our framework will help us document these differences, evaluate their implications for economic growth, and enable the study of factors that caused them, such as knowledge access barriers.

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 main ways. Most importantly, 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 innovation

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

[10] See, e.g., Harris (2017), pp. 540–543, p. 560, who discusses important French innovations in these technologies.

frictions, δ_{ci}, the key primitives of our theory.[11] We also allow countries to differ in their time-varying institutional environments, represented by a set of parameters γ_{ct} that reflect the probability that innovators will successfully obtain property rights over their invention. Our theory features two types of cross-country spillovers: 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.[12] Finally, we decouple technology categories from goods sectors. Rather than mapping technologies one-to-one 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. Besides adding realism and facilitating a cleaner mapping from model to data, this structure allows some technologies to be more "general," in the sense of facing greater industry demand.

We formulate the model in discrete time to facilitate the mapping onto the data. We begin by describing a closed-economy model. We then generalize to incorporate multiple countries and allow knowledge spillovers and technology sharing between them.

Preferences and Production The model features a representative consumer in country c with utility at time t that is a function of discounted log consumption $Y_{c,t}$ in period t and all future periods:

$$V_{ct} = \sum_{s=0}^{\infty} e^{-\rho s} \ln Y_{c,t+s}$$

The aggregate consumption good in period t, Y_{ct}, is a Cobb-Douglas combination over consumption from sectors $k = 1...K$:

$$Y_{ct} = \prod_{k=1}^{K} y_{ckt}^{\beta_k} \tag{1}$$

where the β_k parameters give the consumption shares for each sector and $\sum_k \beta_k = 1$. The price of the aggregate consumption good in a country is denoted P_{ct} and can be normalized to one for one country.

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,

$$\ln y_{ckt} = \sum_{i}^{I} \alpha_{ik} \int_0^1 \ln u_{cikt}(v) dv \tag{2}$$

[11] In 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 rationalize differences in allocations across countries except as a product of differential consumption preferences.

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

where $u_{cikt}(v)$ is the amount of intermediate input of variety v produced using technology type i used in production sector k 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 $q_{cit}(v)$ and labor:

$$u_{cikt}(v) = q_{cit}(v) l_{cikt}(v) \tag{3}$$

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

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

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

We require that $q_{cit} > 0$ for all i, t.

Innovation In each period, innovations arrive in each variety according to a Poisson process with rate ϕ_{cit} per period, so that the number of innovations in variety v during a period is Poisson-distributed with mean ϕ_{cit}. 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 $(1 + \lambda)^{N_v} q_{cit}(v)$.

The expected number of innovations per variety, ϕ_{cit}, is increasing in the level of research effort in technology type i, $r_{cit} \geq 0$, and in knowledge spillovers through the innovation network, X_{cit}, and decreasing in the current technology level q_{cit}[13]:

$$\phi_{cit} = \ln\left(1 + \frac{r_{cit} X_{cit}}{q_{cit}^{1+\kappa}}\right). \tag{4}$$

In the closed-economy version of the model, we define $\ln X_{cit} = \sum_{j=1}^{I} \omega_{ij} \ln q_{cjt}$ where the parameters ω_{ij} are the elements of the innovation network $\Omega_{I\times I}$ that govern the level of innovation spillovers from j to i. We assume that $\sum_{j=1}^{I} \omega_{ij} = 1$ to allow for the possibility of sustained economic growth. Later, we will modify our expression for these innovation spillovers to differentiate between within-country and cross-country spillovers.

The κ parameter in Eq. 4, which we assume satisfies $\kappa \geq 0$, allows for either an endogenous growth model ($\kappa = 0$) or a semi-endogenous growth model ($\kappa > 0$) as suggested by Jones (1995).

There are several attractive features of the structure in Eq. 4, which differs from that used by Liu and Ma (2023). First, it is well defined even for technology types with zero research effort. Second, it accommodates endogenous growth when $\kappa = 0$ and semi-

[13]This 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.

endogenous growth when $\kappa > 0$. Third, as we will see, this framework allows us to cleanly incorporate endogenous research effort under more realistic assumptions than in previous work.[14]

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_{cit}, so:

$$\ln q_{ci,t+1} = \ln q_{cit} + \phi_{cit} \ln(1+\lambda) \tag{5}$$

Resource constraints and innovation frictions In the model, a given population of workers L_{ct}, each endowed with one unit of labor, can choose between research and production work. Let R_{cit} be the amount of labor doing research in technology i in country c. We then have the following resource constraint:

$$\sum_{k=1}^{K}\sum_{i=1}^{I} l_{cikt} + \sum_{i=1}^{I} R_{cit} = L_{ct} \qquad l_{cikt} = \int_0^1 l_{cikt}(v)dv \tag{6}$$

If an individual chooses to become a researcher, then they must expend a fraction $(1-\delta_{ci})$ of their time in order to develop the skills and knowledge necessary to undertake research in technology type i. The term $(1-\delta_{ci})$ measures the *innovation friction*. After paying the time cost to acquire the needed skills and knowledge, researchers are left with $r_{cit} = R_{cit}\delta_{ci}$ of their time remaining as effective research effort.

Decentralized researcher and worker allocation The allocation of researchers across sectors is determined through the market. When a new quality improvement arrives, the researcher that produced it receives a one-period monopoly on producing that intermediate variety with that quality level.[15] 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_{ct}. Note also that the Cobb-Douglas structure of the final goods aggregator (Eq. 1) ensures that the revenue expended on each final good is $\beta_k P_{ct} Y_{ct}$. Using these, and solving the optimization problem for the firms that produce good k, we can derive the amount of labor employed by perfectly competitive intermediate firms that use technology i to produce inputs for

[14] In particular, the structure used in Liu and Ma (2023) is not defined when research effort approaches zero and requires strong assumptions—a single monopolist research firm in each technology category—to generate clean analytical expressions in the decentralized version.

[15] That 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 last innovator holds the monopoly.

sector k:

$$l_{cikt}^{P.C.}(v) = \frac{\alpha_{ik}\beta_k P_{ct} Y_{ct}}{w_{ct}} = \alpha_{ik}\beta_k L_{ct}$$

using the fact that product market clearing implies $w_{ct}L_{ct} = P_{ct}Y_{ct}$.[16]

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,

$$l_{cikt}^{Mon}(v) = \frac{\alpha_{ik}\beta_k P_{ct} Y_{ct}}{(1+\lambda)w_{ct}}$$

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

$$\pi_{cikt} = \frac{\lambda}{1+\lambda}\alpha_{ik}\beta_k P_{ct} Y_{ct}$$

We introduce a set of parameters $\gamma_{ct} \in [0,1]$ to capture the institutional environment. Specifically, γ_{ct} is the fraction of successful innovations where the inventor is able to obtain and enforce property rights and extract monopoly rents for the use of that technology. In the remaining $1-\gamma_{ct}$ of innovations, the technology becomes freely available to anyone.

Free entry into research implies that individuals will choose to become researchers until the expected return to being a researcher is equal to the wage. A researcher's expected return, ERR_{cit}, depends on the fraction of varieties that receive at least one innovation, $1-e^{-\phi_{cit}}$, the fraction of these that obtain monopoly protection (γ_{ct}), divided by the number of researchers R_{cit}, times the monopoly profit:[17]

$$ERR_{cit} = \frac{(1-e^{-\phi_{cit}})\gamma_{ct}\sum_k \pi_{cikt}}{R_{cit}}$$

Setting ERR_{cit} equal to the wage rate, using $r_{cit} = R_{cit}\delta_{ci}$, and using that total income must equal the total wage bill ($P_{ct}Y_{ct} = w_{ct}L_{ct}$), we obtain the following expression for effective research effort:

$$r_{cit} = \delta_{ci}C_i\gamma_{ct}L_{ct} - \frac{q_{cit}^{1+\kappa}}{X_{cit}}. \tag{7}$$

where $C_i = \frac{\lambda}{1+\lambda}\sum_k \alpha_{ik}\beta_k$.

Plugging Eq. 7 into Eq. 4 and then into Eq. 5, and writing $\Delta \ln q_{ci,t} \equiv \ln q_{ci,t} - \ln q_{ci,t-1}$, we obtain the following expression for the change in technology levels in a particular

[16]Note that free entry into research implies that the expected return to a researcher is driven down to the wage of production workers, so that total income in the economy is $w_{ct}L_{ct}$.

[17]Note that if multiple innovations arrive in a variety in the same period, only one receives monopoly profits.

country-technology:

$$\Delta \ln q_{ci,t} = \ln(1+\lambda)\left[\ln \delta_{ci} C_i + \ln \gamma_{c,t-1} L_{c,t-1} + \ln X_{ci,t-1} - (1+\kappa)\ln q_{ci,t-1}\right] \tag{8}$$

Eq. 8 will be the foundation for our empirical analysis. It tells us that we can decompose the factors driving technological progress into (i) country-by-technology innovation frictions δ_{ci}, (ii) a set of fixed technology-level demand factors C_i, (iii) time-varying aggregate country conditions represented by $\gamma_{ct} L_{ct}$, and (iv) time-varying country- and technology-specific forces driven by existing technology levels, either operating directly (q_{cit}) or through the innovation network (X_{cit}).

3.1 Extensions: Knowledge spillovers and technology sharing

We now extend the model to allow for two different types of cross-country spillovers. First, we incorporate cross-country *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.

In our baseline closed-economy model, innovation spillovers are given by $\ln X_{cit} = \sum_{j=1}^{I} \omega_{ij} \ln q_{cjt}$. We extend this to two countries by setting $\ln \tilde{X}_{cit} = \iota \ln X_{cit} + (1-\iota)\ln X'_{cit}$, where we use $'$ to denote the foreign country and where $\iota \in [0.5, 1]$ is the parameter governing the importance of domestic knowledge spillovers for domestic innovation. Thus, the two-country version of Eq. 8 is,

$$\Delta \ln q_{ci,t} = \ln(1+\lambda)\left[\ln \delta_{ci} C_i + \ln \gamma_{c,t-1} L_{c,t-1} + \iota \ln X_{ci,t-1} + (1-\iota)\ln X'_{ci,t-1} - (1+\kappa)\ln q_{ci,t-1}\right] \tag{9}$$

where $\ln X'_{cit} = \sum_{j=1}^{I} \omega_{ij} \ln q'_{cjt}$ represents spillovers to country c from the foreign country.

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 technology:

$$\ln y_{ckt} = \sum_i \alpha_{ik} \int_0^1 \left[\eta \ln(q_{cit}(\nu) l_{cikt}(\nu)) + (1-\eta)\ln(q'_{cit}(\nu) l_{cikt}(\nu))\right] d\nu \tag{10}$$

where $q_{cit}(\nu)$ and $q'_{cit}(\nu)$ denote, respectively, domestic and foreign technology from the vantage point of country c, and $\eta \in [0.5, 1]$ determines the importance of each technology in the production process. Later, we will estimate η from production data and examine in the simulations how much technology sharing influences relative growth rates and output levels in our setting.

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—which is our primary object of interest, given that we are studying the take-off of modern economic growth—resulting in permanent level differences in output.

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.[18] 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.[19] We add to these data technology classifications produced by the BPO. Patents are classified into one or more of 602 technology subcategories that aggregate hierarchically into 147 technology categories.[20]

The French patent data used in our analysis were digitized by the French national patent office (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

[18]For 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.

[19]Most 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 are unknown.

[20]See Hanlon (2023) for additional details about the data.

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 new inventions.[21] 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.[22]

4.2 Linking inventors' 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.[23]

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 nations' 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, we construct a list of how British *sub*categories correspond to French *sub*categories (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.

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

[22] Another difference between the French and British patent systems is that inventors in France could choose 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.

[23] The 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.

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 some technologies map unambiguously to one industry, while others 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.[24]

The key challenge is how to map these industries to patenting technology categories.[25] 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's data also provide 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. Our model includes country–year fixed effects, which we will estimate relative to an omitted reference year, and this will allow for aggregate cross-country differences in patent quality. To further account for differences in patent quality within country, across technology types, we will adjust patent counts using measures similar to those in the literature, such as patent citations (Trajtenberg, 1990; Harhoff, Narin, Scherer, and Vopel, 1999) and the payment of renewal fees (Schankerman and Pakes, 1986). Specifically, we have two different measures of patent quality for the two countries we study. For British patents, we have quality measures provided by Nuvolari and Tartari (2011) 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

[24]The 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 candles*, *Oils and paint*, *Explosives*), and exclude the *Motor and Cycle* industry because it did not yet exist during our period. Horrell, Humphries, and Weale (1994) provide an input-output matrix for the British economy in 1841, which is much less detailed.

[25]No 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.

patents in force. In Appendix G, 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 on 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.[26] 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.[27] 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 spillovers between different technology categories. Specifically, the strength of spillovers from technology category j to category i is typically measured as $\omega_{ij}^{cite} = Cites_{ij} / \sum_l Cites_{il}$,

[26] This exhibition was followed by others, such as the Expositions 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.

[27] The 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, but as most exhibits have only one category we focus on each exhibit's first category.

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 are 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 lead to a subsequent discovery in technology category i. If we found that inventors in category i had often previously patented in j but not in 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

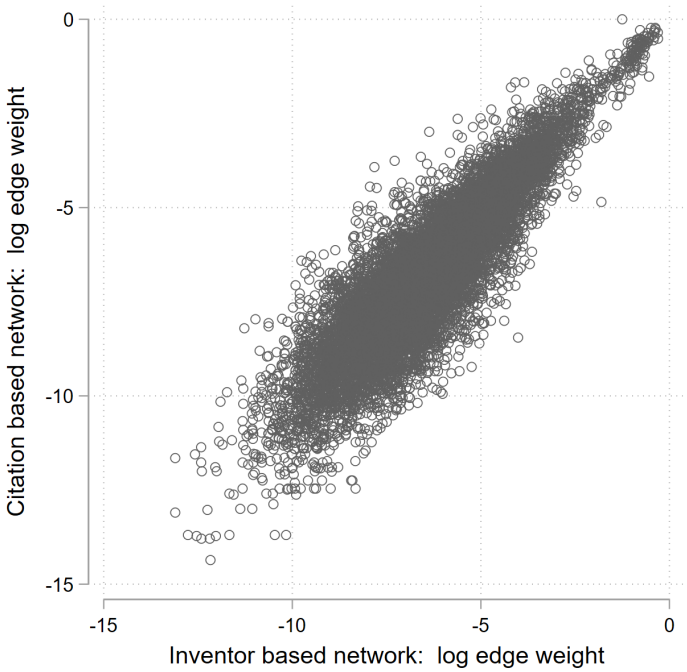
$$\omega_{ij} = \frac{\sum_m P_{mij}}{\sum_m P_{mi}} \tag{11}$$

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 P_{mi} 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).[28]

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 joint matrix can be constructed in several ways; each will require judgment about what relative weight should be attached to patents from each country. As patents in the two countries are the

[28] One 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 toss a coin for cases without a 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 using U.S. patents from 1970–2014 based on data from PatStat. For data and method, see text.

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.

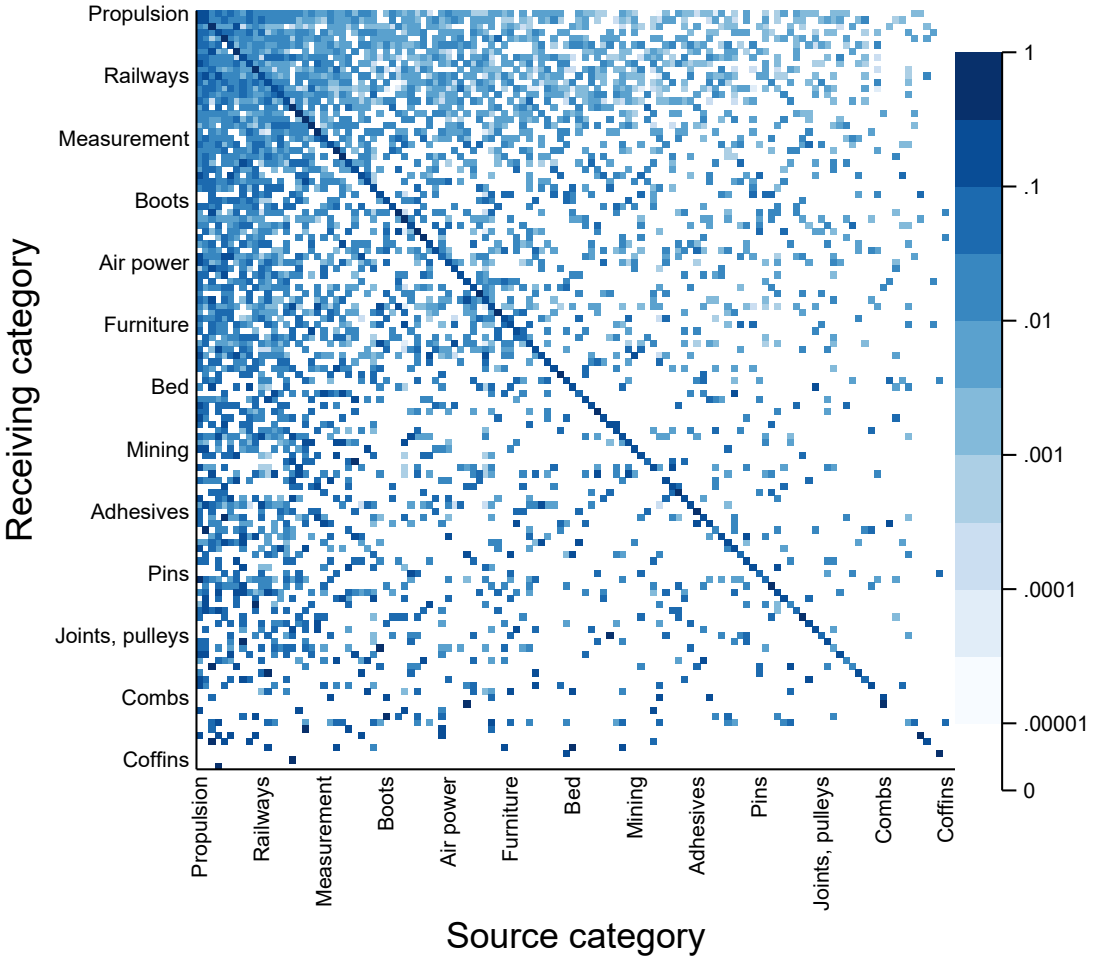
As a second validation of our method, focusing specifically on the historical context we study, we take advantage of data on the full text of all British patents from 1700–1850 provided by Berkes, Chen, and Tranchero (2025). Using the method developed by Kelly, Papanikolaou, Seru, and Taddy (2021), we calculate the bilateral similarity between every patent pair. We then use these to construct a directed network of spillovers across technology categories over time. Regression results, in Appendix Table D.1, show that this text-based spillover network is highly correlated with the network obtained using our method. This provides a second, historical, validation of our method.

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 order to generate the first estimates of innovation networks during the Industrial Revolution. We offer two approaches to visualizing the results. Figure 2 presents a heat map showing the strength of spillovers across different technology categories based on the joint matrix constructed using both British and French patents. The main take-away from this figure is that technologies differed quite substantially in the extent to which they generated spillovers. Some, such as propulsion, railways, or measurement, appear to produce substantial spillovers, while categories such as coffins produced very little.

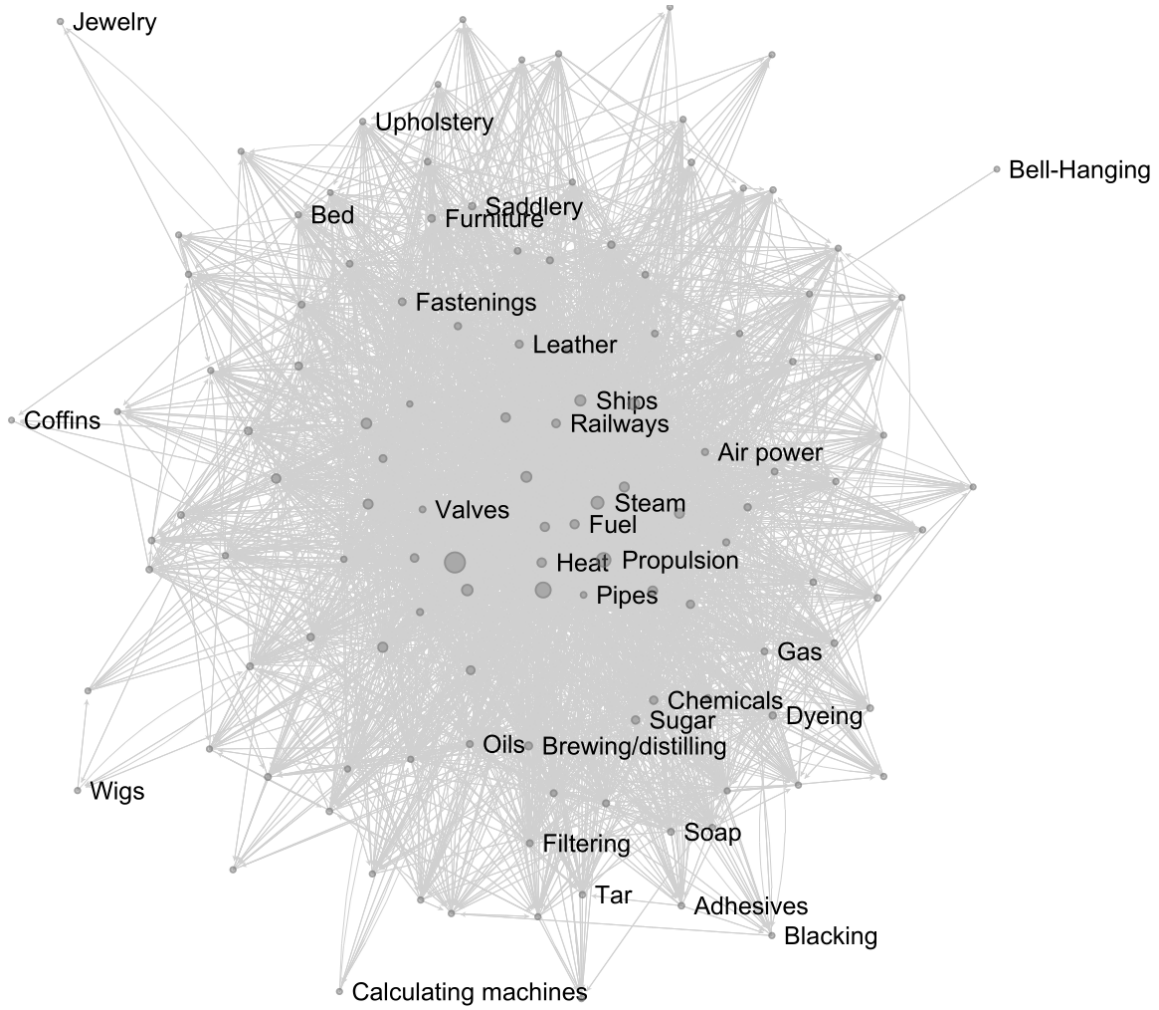
An alternative visualization is shown in Figure 3. In this figure, each technology category is a node, the size of the node reflects the number of patents filed in that category, and the location of the node is determined by the strength of connections between that node and every other node in the network as determined by a multidimensional scaling algorithm. There are several interesting patterns in this figure. 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, Dyeing, 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.

Figure 2: Heat map of the innovation network during the Industrial Revolution



This is a heat map showing the strength of spillovers through the innovation network based on the joint matrix constructed from British and French patent data using our inventor-based method. Rows are sorted by left (sending) eigencentrality. Note that for readability category labels are included for only a subset of technology categories.

Figure 3: The joint innovation network visualized



Plot generated using multidimensional scaling. The edges of the joint network are computed as $(\omega_{ij}^{UK} + \omega_{ij}^{FR})/2$. For readability, labels are included for only a subset of technology categories. See Figure E.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***	0.148***	0.246***	0.233***	0.527***	0.857***	0.896***	1.091***
	(0.014)	(0.014)	(0.036)	(0.036)	(0.060)	(0.050)	(0.046)	(0.064)
Receiving node (i) FE		✓		✓		✓		✓
Sending node (j) FE		✓		✓		✓		✓
N (Obs = edges)	14884	14884	14884	14884	650	614	650	614
Estim. FE coef.		243		243		158		158
R^2	0.05	0.28	0.11	0.14	0.44	0.72	0.39	0.60
Within R^2		0.01		0.10		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 123×123 = 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

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

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

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 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. Across all specifications, but specifically when we focus on the intensive margin, we find a strong similarity of the network edges across countries, which is suggestive of a common underlying network structure.

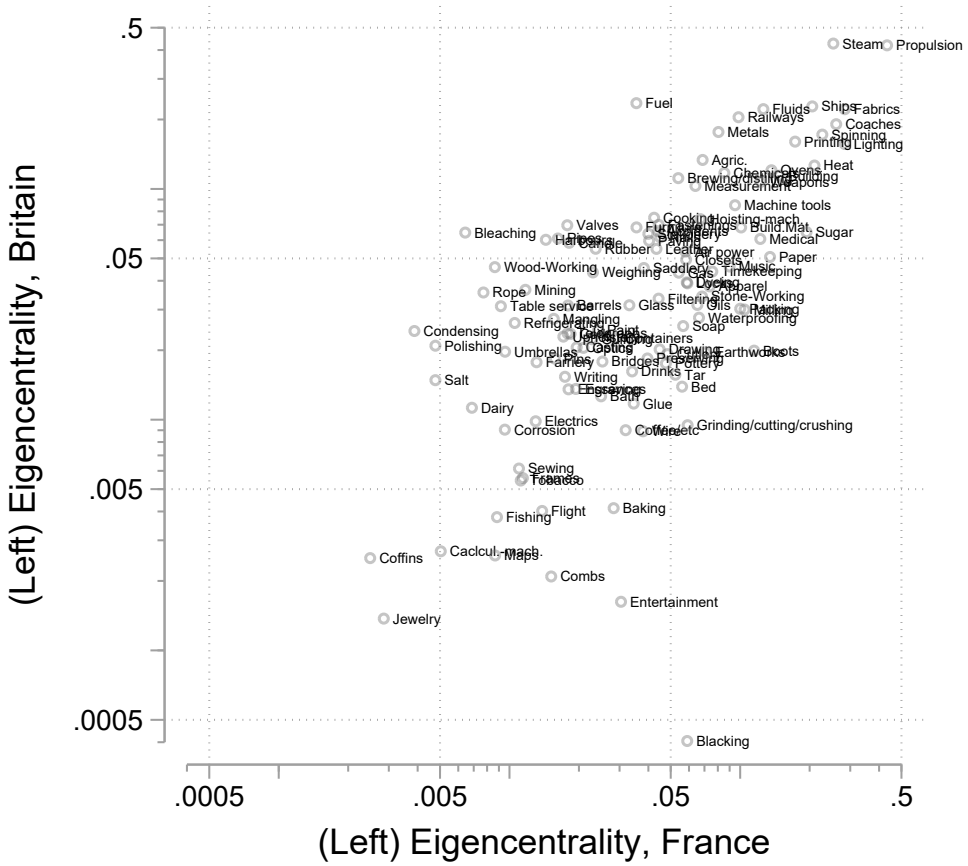
Node centrality An alternative approach to assessing matrix similarity is to focus on the centrality of the network nodes, which provides a useful way to summarize the shape of the network. In the network literature, there exist different metrics for node centrality; which is most appropriate depends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

Figure 4 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 E.1, 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 who exhibited 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 F), except that the exhibition data do 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 F, appears similar to the network obtained from the patent data. However, we can examine that similarity more formally by using the links that

Figure 4: 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.

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 E.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, matrix stability appears to be a reasonable assumption for the purposes of 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 key model parameters

The structure of the model, together with the patent data, can be used to recover estimates of the key model parameters from only observed research output (quality-adjusted patents). The starting point for this exercise is Eq. 9. We can rewrite this as an estimating equation:

$$\Delta \ln q_{cit} = \beta_X \ln X_{ci,t-1} + \beta_{X'} \ln X'_{ci,t-1} + \beta_q \ln q_{ci,t-1} + \xi_{ci} + \theta_{ct} + \varepsilon_{cit} \tag{12}$$

where ξ_{ci} are a full set of country-by-technology fixed effects and θ_{ct} are country-by-year fixed effects, with $\theta_{c,t0} = 0$ in the reference year.

We map this equation to our data by setting $q_{cit} - q_{ci,t-1} = Patents_{cit}$ where $Patents_{cit}$ is the number of quality-adjusted patents in country c, technology i, and period t. This implies that $q_{ci,t-1} = \sum_{\tau=0}^{t-1} Patents_{ci\tau}$. For our baseline estimation of Eq. 12, we focus on the period 1800–1843, when we have a substantial number of patents in both countries (though we also provide robustness checks using alternative periods) and choose 1800 as the reference year.[29] The three estimated coefficients, β_X, $\beta_{X'}$, and β_q, capture the three unknown macro parameters λ, κ, and ι:

$$\lambda = e^{\beta_X + \beta_{X'}} - 1; \quad \iota = \frac{\beta_X}{\beta_X + \beta_{X'}}; \quad \kappa = \frac{-\beta_q}{\beta_X + \beta_{X'}} - 1\,.$$

Table 2 presents our estimates of the main coefficients from Eq. 12 for a range of alternative specifications. Column 1 is our baseline specification, where we pool both countries and estimate results over the full 1800–1843 period using quality-adjusted patents. At the bottom of the table, we report the values of the key parameters, λ, ι, and κ, implied by our coefficient estimates. It is useful to note that our estimate of $\lambda = 0.092$ is not too far from the value found in previous work: Acemoglu, Akcigit, Alp, Bloom, and Kerr (2018) estimate a $\lambda = 0.132$ 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 of around 0.75, which suggests that cross-country

[29] One 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), but we examine the robustness of our results to other choices.

knowledge spillovers were meaningful during this period, though domestic knowledge spillovers were still more important overall. The estimate $\kappa = 0.246$ indicates that our data are best fit by a semi-endogenous growth model where rising technology levels impose a mild drag on the productivity of research effort.

Table 2: Main estimation table

	Dep var: $\Delta \ln q_{c,i,t}$				
	(1)	(2)	(3)	(4)	(5)
	Baseline	Patent counts	q_0^{UK} 1790	Late 1816–43	Endogenous growth
$\ln \chi_{t-1}$	0.066***	0.064***	0.069***	0.072**	0.084***
	(0.023)	(0.024)	(0.023)	(0.034)	(0.015)
$\ln \chi'_{t-1}$	0.022**	0.021**	0.021*	0.034**	0.029***
	(0.009)	(0.009)	(0.013)	(0.017)	(0.010)
$\ln q_{t-1}$	−0.110***	−0.110***	−0.123***	−0.115***	−0.113***
	(0.012)	(0.010)	(0.012)	(0.015)	(0.012)
N	9152	9145	9043	6251	9152
R^2	0.33	0.36	0.34	0.38	
F-stat	24.88	27.22	24.59	27.86	24.98
implied λ	0.092	0.090	0.095	0.112	0.119
implied ι	0.754	0.752	0.765	0.679	0.742
implied κ	0.246	0.283	0.357	0.076	0

The table presents OLS estimates of Eq. 12 in columns 1 through 4, and constrained LS estimates in column 5 such that $\beta_X + \beta_{X'} = -\beta_q$. The top panel lists estimated coefficients, with robust standard errors in parentheses, while the bottom panel gives the implied values for λ, ι, and κ. All specifications include a full set of country-by-technology category fixed effects, plus a set of country-by-year fixed effects with 1800 as omitted reference year. The first column gives the baseline specification, which focuses on outcomes in the period 1800–43, 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 X. The second column uses patent counts in place of quality-adjusted patents. The third column uses only British patents filed after 1790 to construct knowledge levels, to better match the time period available for France. The fourth column looks only at a later period, starting after the end of the Napoleonic Wars. In the fifth column, we show results for an endogenous growth model where κ is constrained to be equal to zero.

Columns 2–4 present results obtained using different input data. It is notable that the estimates are very stable across these alternatives. This tells us that factors such as adjusting for patent quality (Column 1 vs. Column 2), using different sets of patents to construct initial technology levels (Column 3), or looking only at later outcomes (Column 4), do not meaningfully impact the results. The biggest difference is in Column 4, which focuses on the period after the end of the Napoleonic Wars, but we note that those estimates are not statistically distinguishable from our baseline results. Column 5 presents results

obtained from an endogenous growth model where κ is constrained to be zero.

Next, we use the estimated fixed effects $\hat{\xi}_{ci}$ and $\hat{\theta}_{ct}$ to recover the parameter vectors of the innovation frictions δ_{ci} and the institutional environment γ_{ct} as follows:

$$\hat{\delta}_{ci} = e^{\hat{\xi}_{ci}/\ln(1+\lambda)} \cdot \frac{1}{C_i \, L_{c,t0} \, \gamma_{c,t0}} \tag{13}$$

$$\hat{\gamma}_{ct} = e^{\hat{\theta}_{ct}/\ln(1+\lambda)} \cdot \frac{L_{c,t0} \, \gamma_{c,t0}}{L_{c,t}} \tag{14}$$

Given $\hat{\lambda}$, the constant $C_i = \frac{\lambda}{1+\lambda} \sum_k \alpha_{ik}\beta_k$ is data. We also have data for population L_{ct} from Bolt and van Zanden (2025) and O'Brien and Keyder (2011).[30] The parameter vectors are identified up to a normalizing constant $\gamma_{c,t0}$ of the reference year t_0.

Obtaining the innovation friction parameters δ_i^c is an important contribution of our paper. To give a sense of them, Table 3 presents a listing of the technologies where Britain had the largest advantage (left panel) and those where France had the largest advantage (right panel). We focus on the cross-country ratio of innovation frictions $\delta_i^{UK}/\delta_i^{FR}$, which is identified up to a constant that we normalize to 1 for the purpose of exposition. In the left panel, we can see that Britain had a large advantage in metalworking and mechanical technologies (e.g., cylinders, pistons and rollers; joints and pulleys; grinding and polishing) as well as technologies strongly associated with factory production (e.g., preventing accidents). In contrast, we can see that France's advantages appear concentrated in food production and consumer products, as well as, notably, aerial conveyances (i.e., air balloons), an area where France was undoubtedly the early technological leader. These estimates are consistent with historical sources, which emphasize how "Britain and France had followed different paths to industrialization, and that the British model was not the only paradigm possible" (Crouzet, 2003, p. 10). While Britain excelled in mechanical technologies, "The French…successfully developed the production of labour-intensive, quality, fashion and luxury goods" (Crouzet, 2003, p. 11).

Our regressions also deliver a full set of country-by-time fixed effects. These can tell us something about *relative* changes in the innovation environment in the two countries across the study period. Plotted in Figure 5, these estimates show that starting around 1820 France experienced a substantial improvement relative to Britain. This corresponds to the end of "the economic disasters of the revolutionary and Napoleonic periods" (Crouzet, 2003, p. 10), which likely improved the outlook for inventors in France. We can also see the short-lived impact of the July Revolution of 1830, which set back French innovation in 1830–31. In contrast, our estimates suggest a generally stable innovation environment in

[30]The Maddison Project Database (Bolt and van Zanden, 2025) reports yearly population starting in 1820. We amend this with decadal population from O'Brien and Keyder (2011) for 1802–1813 and 1815–1824 to backcast population, assuming constant exponential growth rates. If taken at the decadal midpoint, O'Brien and Keyder's decadal population figures line up almost exactly with the Maddison population figures.

Table 3: Technologies with the largest innovation friction ratios across countries

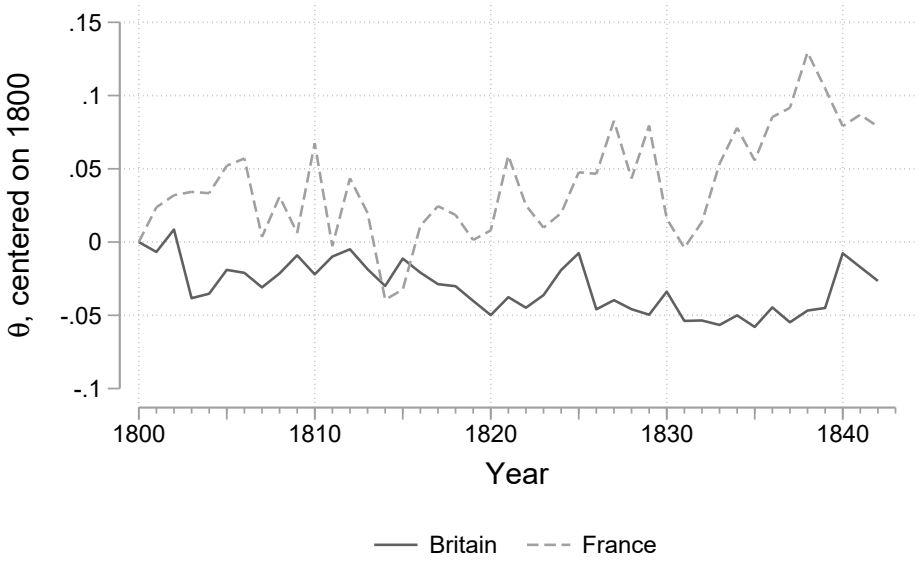
Top British category	δ ratio	Top French category	δ ratio
Cylinders, Pistons, Rollers	171	Aerial Conveyances	0.069
Joints and Pulleys	96	Tar, Resin, Oil Distillation	0.095
Grinding and Polishing	75.5	Cooking and Baking	0.114
Mouldings and Ornaments	11.7	Coffee, Tea and Cocoa	0.121
Preventing Accidents	8.4	Games, Exercises, Amusements	0.138
Pipes and Tubes	7.8	Blacking	0.145
Bleaching, Washing, Scouring	6.7	Fishing	0.176
Rope Manufacture	5.7	Bed and bedding	0.184
Farriery and Veterinary medicine	5.4	Baths and Bathing-Machines	0.185
Brushes	3.9	Electricity and Magnetism	0.205

Ratios of estimated British to French innovation frictions, δ^{UK}/δ^{FR}, calculated from Eq. 13 using data and the baseline estimates of Table 2, Column 1. The ratio is normalized by the geometric mean across technology categories so that country-specific constants cancel.

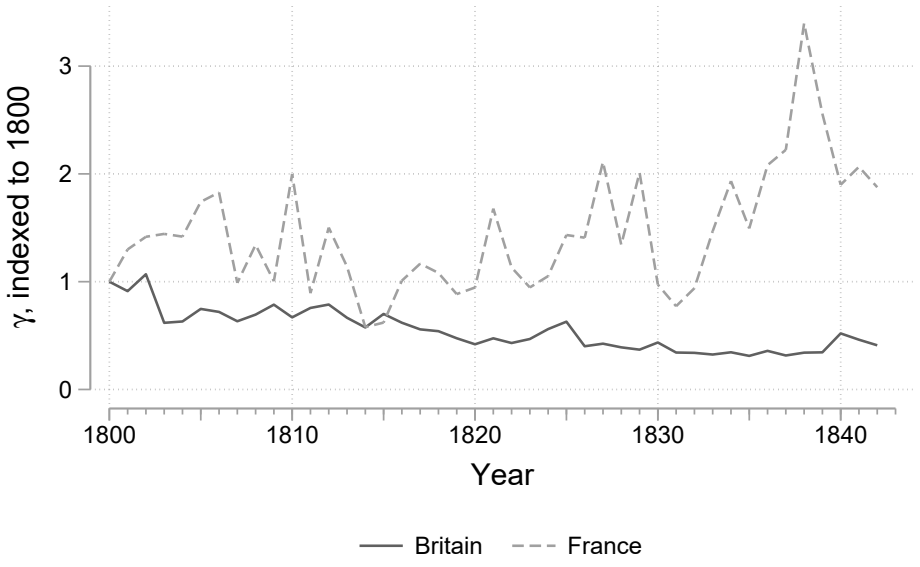
the UK with perhaps a mild downward drift.[31]

[31]This might reflect a combination of rising foreign competition and the fact that enforcement of patent rights against foreign infringers was imperfect.

Figure 5: Relative time paths of country–year shocks



(a) θ: country–year fixed effects, centered on 1800



(b) γ: adjusted by population growth, indexed to 1800

Panel (a) depicts the country-year fixed effects from the baseline regression in Table 2, Column 1, which are centered in both countries on 1800. Panel (b) depicts the implied values of γ, recovered via Eq. 14 and indexed in both countries to 1800.

8 Growth implications on the transition path

We now use the model to examine the growth implications of the differences in the innovation frictions that we have estimated. An important feature of our estimation procedure is that it delivers valid parameter estimates either on or off the balanced growth path. This means that we can use our estimates to consider patterns of growth during a period of transition, such as the one that must have characterized the Industrial Revolution.

We simulate two economies, representing Britain and France, that start from the same initial technology levels but differ in their innovation frictions, δ_{ci}, and in their institutional environment, γ_{ct}. Specifically, we endow each economy with the estimated $\hat{\xi}_{ci}$ and $\hat{\theta}_{ct}$ values from our preferred specification, in Column 1 of Table 2, which incorporate these features as well as the influence of country population differences. The macro-parameters in this simulation, λ, κ, and ι, also come from the same estimation. The only input we are missing is η, which affects how the technology level translates into output when we allow for technology sharing across the economies. 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 nineteenth century, as described in more detail in Appendix J.[32] Our results indicate a high level of technology sharing, with an average estimate of η of 0.71 (0.5 would indicate perfect sharing, 1 complete reliance on domestic technology). However, it is important to recognize that this estimate should be interpreted with caution. This is one dimension where further improvements in data availability have the potential to substantially improve the results that methods such as the one developed in this paper can generate.

We have to choose an initial technology level for our simulation. In our baseline results, we start with a uniform technology level of $q = 10$ for all technologies. The simulation is unaffected by the exact level of technology, but the assumption of uniformity is meaningful. In Appendix I, we present robustness exercises where we allow realistic variation in initial technology levels. These produce results that are quite similar to our baseline findings.

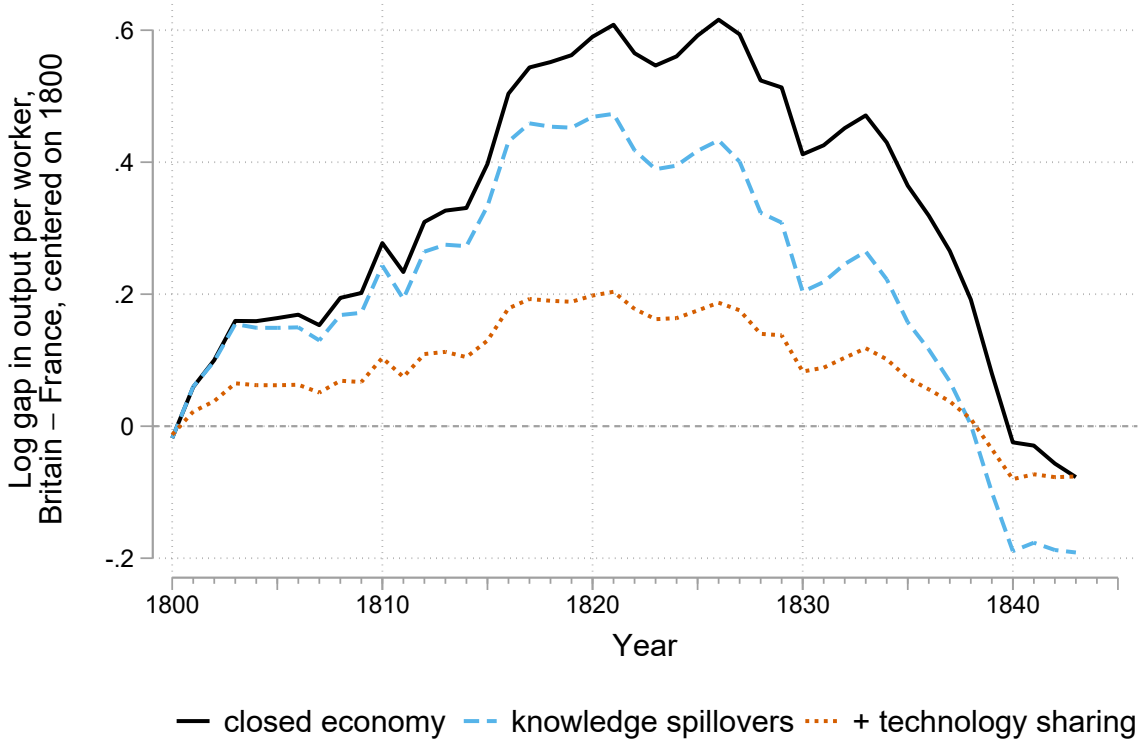
Figure 6 shows how the two simulated economies behave across the study period, relative to the output per worker differences in 1800 (the omitted year in our estimation).[33] Consider first the heavy black line, which corresponds to our closed economy model. We can see that the simulated British economy grows more rapidly across the first two decades. After 1820, France begins to catch up, and rapidly so after 1830. Below, we will show that Britain's initial advantage is due, in part, to how Britain's innovation frictions are

[32] We focus on France because no similarly detailed industry output data are available for the UK in the early nineteenth century. We 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.

[33] For log output per worker, see Appendix Figure I.1 which reports it for the two economies separately.

distributed across different technology categories. The reversal after 1820 is due largely to improvements in the innovation environment in France (i.e., the γ_{ct} terms in the model), as suggested by the patterns shown in Figure 5. The timing of this reversal makes sense in historical perspective, since it corresponds to a period of recovery after the massive disruptions of the Napoleonic Wars had ended. Note that the simulation does *not* predict that the French level of output per worker would be above the British one by the early 1840s: the simulation is *relative* to any pre-existing output per worker gap in 1800.

Figure 6: Simulated difference in log output per worker starting from a common initial technology level



The figure plots the difference of log output per worker in two simulated economies, one with the estimated parameters of Britain and the other with the estimated parameters of France, starting from a common initial technology level. The simulation uses the estimated parameters corresponding to the specification in Column 1 of Table 2.

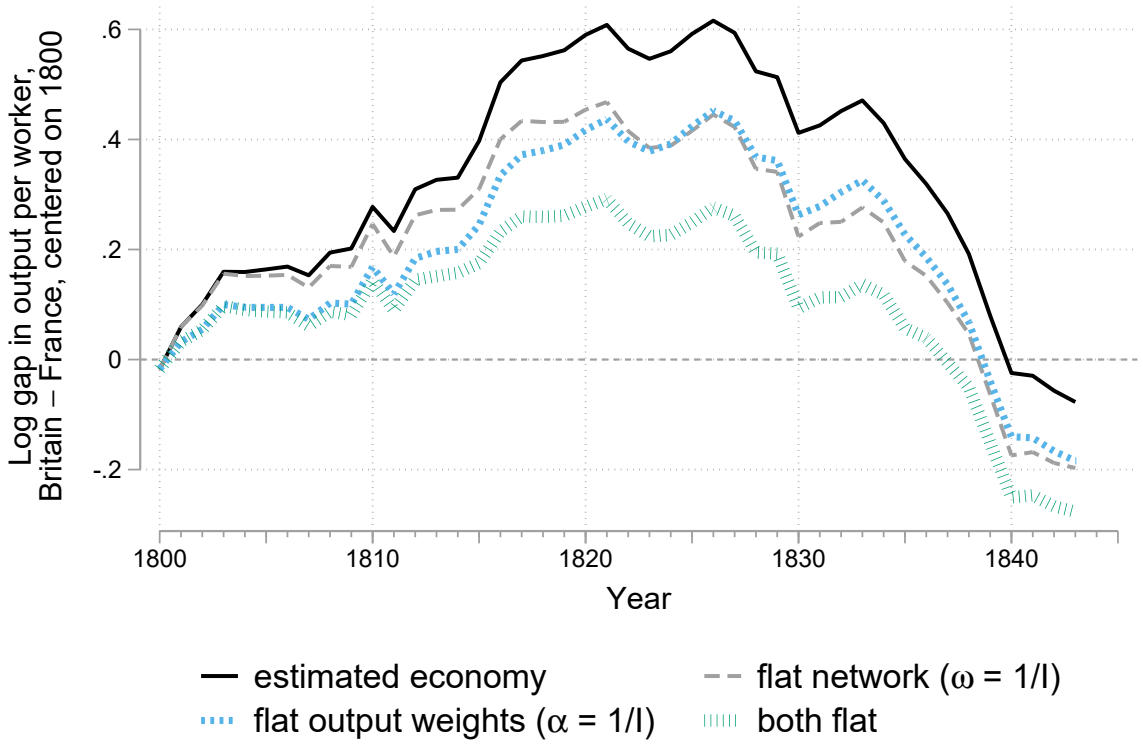
The middle line in the graph corresponds to the extension with knowledge spillovers. We can see that this reduces Britain's relative advantage, as French innovators benefit from the knowledge spillovers from Britain. But this effect only emerges over time: Britain had to build a knowledge lead first, and France to fall behind, before catch-up growth became meaningful.

The bottom dotted line then adds in the possibility of technology sharing. This pushes the output per worker gap towards parity by allowing each economy to use technology developed in the other for production. Remember that this is based on a value for η

that is estimated from French data but assumed to be the same in Britain. Insofar as the assumption that Britain was as reliant on French technology as France was reliant on British technology is implausible, the bottom dotted line will represent a lower bound on growth differences between the two countries. Still, the model predicts an output advantage in Britain relative to differences in 1800 across most of the study period.

Next, we conduct counterfactual simulations to study the source of Britain's advantage. Figure 7 starts by reproducing the heavy black line from the closed economy scenario, which is also the top line on this graph. We consider counterfactuals where we "turn off" each of the two main channels in the model through which the distribution of innovation frictions across technologies can generate a growth advantage.

Figure 7: Effect of equalizing output weights or network spillovers on simulated output per worker gap



The figure plots the difference in log output per worker in two simulated economies, one with the estimated parameters of Britain and the other with the estimated parameters of France, starting from a common initial technology level. The simulation uses the estimated parameters corresponding to the specification in Column 1 of Table 2. The estimated economy corresponds to the closed economy scenario from the previous Figure 6. The three counterfactuals either equalize centrality in the innovation network ($\omega_{ij} = 1/I \; \forall i, j$, gray dashed line) or equalize generality in production ($\alpha_{ik} = 1/I \; \forall i$, blue short dashed line) or equalize both (green hashed line).

In the short-dashed blue line, we turn off differences in the production demand served by technologies by setting the α_{ik} parameters to be equal for all i. Equalizing demand

across technologies substantially reduces Britain's output per worker advantage. This tells us that part of Britain's initial growth advantage was due to having lower innovation frictions in technologies that served a larger share of demand. In the dashed gray line, we instead turn off the role of the innovation network by equalizing spillovers across the full matrix Ω.[34] Relative to the baseline case, this also reduces Britain's growth advantage by a meaningful amount, which tells us that Britain also had an advantageous distribution of innovation frictions relative to the innovation network. Finally, the hashed green line turns off both advantages, which further reduces Britain's growth advantage.

In sum, these results tell us that a substantial part of Britain's initial growth advantage was due to having a beneficial distribution of innovation frictions relative to both demand and innovation network spillovers.[35] The remainder was due to differences in the average level of innovation frictions across all technology categories and differences in the overall institutional environment.

Another factor that affects relative output in the simulation is the growth drag represented by the κ parameter. In Appendix Figure I.2, we show that turning this effect off by setting $\kappa = 0$ substantially increases Britain's growth advantage. This makes sense because this growth drag will make innovation more costly in the more advanced economy. The result underscores that allowing for semi-endogenous growth has important quantitative implications for the different growth trajectories of the economies we study.

9 Application: Innovation frictions and knowledge access costs

In this section, we demonstrate how our estimates of innovation frictions can be leveraged to assess existing theories of Britain's growth advantage during the Industrial Revolution. We focus specifically on Mokyr's influential argument that knowledge access costs played an important role in Britain's success during the Industrial Revolution. An ideal evaluation of Mokyr's theory would compare our estimated innovation frictions to a measure of the availability of (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 that 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

[34] An alternative approach to turning off the innovation network is to set it to the identity matrix, which corresponds to a model without any innovation network. We show results using this approach in Appendix Figure I.3.

[35] The influence of these factors shown in Figure 7 is relatively conservative. Counterfactuals that allow for realistic differences in initial technology levels across technology types show larger effects (Appendix Figure I.5).

(Dittmar, 2011).

Specifically, we examine the availability of technology-specific published knowledge in (British) English vs. French from the mid-18th century up to 1800 using the Google Books corpus (Michel et al., 2011).[36] The assumption here is that the availability of written knowledge relevant for each technology type in a particular language is one important element of country-specific knowledge access.[37]

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.[38] We then translate these candidate keywords and verify that the translated terms are strong predictors for the same technology category in the other country.[39] Further details of this procedure and some examples of the produced equivalent keyword pairs are available in Appendix H.

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.[40] We focus on documents published between 1750 and 1800 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 innovation friction 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 4 is conducted at the level of technology categories. An analysis at the level of keyword pairs can be found in Appendix Table H.2.

Panel A of Table 4 compares the log ratio of documents associated with each technology category in the Google Books corpus to the log ratio of the estimated innovation frictions. The columns correspond to the different specifications of the estimating equation 12

[36]Of course, both English and French are spoken in countries other than Britain 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.

[37]This exercise is inspired in part by Juhász et al. (2025).

[38]For 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.

[39]One 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 *machine à vapeur*. It is worth noting that a single keyword in one language may be part of more than one matched pair of keywords.

[40]Ngrams 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 4: Comparing estimated innovation frictions to Google Books documents

	Dep var: Ln δ^B/δ^F ratio				
	Baseline (1)	Patent counts (2)	q^{UK} 1790 (3)	Late 1816–43 (4)	Endogenous growth (5)
Panel A: Document ratio					
Ln document ratio 1750–1799	0.484*	0.550**	0.591**	0.443*	0.386*
	(0.251)	(0.240)	(0.268)	(0.225)	(0.200)
N (Obs = category)	109	109	109	109	109
R^2	0.047	0.059	0.058	0.054	0.048
Panel B: English and French documents					
Ln English documents 1750–1799	0.369	0.458*	0.445*	0.361	0.291
	(0.247)	(0.249)	(0.267)	(0.222)	(0.196)
Ln French documents 1750–1799	−0.447*	−0.520**	−0.544**	−0.417*	−0.356*
	(0.243)	(0.236)	(0.260)	(0.220)	(0.193)
N (Obs = category)	109	109	109	109	109
R^2	0.056	0.065	0.069	0.060	0.057

Each column refers to the corresponding column in Table 2, implementing a different specification of the estimating Eq. 12 which gives rise to different estimated innovation frictions. The document ratio is the ratio of English to French documents, where English refers to the British English corpus and French to the French corpus. We aggregate document frequencies from Google Books N-gram Viewer over the 1750–1799 period. OLS regressions, robust standard errors in parentheses.

displayed in Table 2, which give rise to different estimates of the innovation frictions. We focus on publications from the period 1750–1800 to exclude publications about the patents from which we estimated the innovation frictions. We find a positive and statistically significant association between the ratio of estimated innovation frictions and the measure of knowledge contained in published sources in the two languages. In Panel B, we separate out documents in English and in French. We find that relatively lower innovation frictions in the UK compared to France are positively related to the availability of documents in English, and negatively related to the availability of documents in French.

In Appendix H, we provide additional robustness checks on these results. We show that similar but noisier 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 H.1 provides a scatterplot, corresponding to the specification in Column 1 of Table 4, which shows that our results are not driven by just one or two outliers.

Overall, these results suggest that knowledge access costs, as represented by the availability of printed material related to particular technologies, can help explain the pattern of innovation frictions that emerges from our estimation. This provides support for Mokyr's argument that knowledge access costs played an important role in explaining innovation patterns during the Industrial Revolution.

10 Conclusions

This study pioneers a new approach to understanding growth during the Industrial Revolution that combines modern growth theory with detailed micro-data. The advantage of this approach is that it allows us to provide a rigorous, theoretically grounded, quantitative analysis of the sources of Britain's growth advantage during this crucial period of economic history. Our results show that Britain benefited from innovation advantages in key technologies, which contributed meaningfully to higher British growth during the early stages of the Industrial Revolution. In addition, our analysis of the Google Ngrams data suggests that access to printed knowledge may explain part of the distribution of innovation frictions that we observe, as hypothesized by Joel Mokyr (Mokyr, 2002, 2005, 2009).

To produce these results, we have extended an endogenous growth model with an innovation network from prior work to incorporate decentralized choices by potential researchers who face rich variation across countries and technologies in innovation frictions. We have also introduced and validated a new empirical method that allows the measurement of innovation networks further back in history than is possible using the standard citation-based approach. This opens up possibilities for studying how the innovation network evolves and shapes technological progress over longer periods of time. Finally, we have shown how one can use the structure of the model to back out measures of

research effort and innovation frictions 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.

We view this paper as a starting point. It benefits from recent improvements in economists' ability to bring growth models to the data, together with the rapidly expanding set of micro-data available for studying the Industrial Revolution as well as other historical contexts. However, we expect that further improvements in both theory and data construction are likely to make our approach even more attractive. This paper also suggests directions for future improvement. For example, one of the least satisfying aspects of the current analysis is the estimation of technology sharing, which relies on a relatively small dataset for just one country. Expanding the availability of production data, whether from historical censuses or other sources, has the potential to improve these estimates, and thus our understanding of the role of cross-country technology sharing in the past.

References

Acemoglu, D., U. Akcigit, H. Alp, N. Bloom, and W. Kerr (2018): "Innovation, Reallocation, and Growth," *American Economic Review*, 108, 3450–3491. 25

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. 2, 3, 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. 6

Berkes, E., M. L. Chen, and M. Tranchero (2025): "300 Years of British Patents," Hugging Face dataset, accessed 6 August 2026. 19, 50

——— (2026): "300 Years of British Patents," *Research Policy*, 55, 105347. 50

Bolt, J. and J. L. van Zanden (2025): "Maddison-Style Estimates of the Evolution of the World Economy: A New 2023 Update," *Journal of Economic Surveys*, 39, 631–671. 27

Bottomley, S. (2014a): *The British Patent System during the Industrial Revolution 1700–1852: From Privilege to Property*, Cambridge University Press. 6

——— (2014b): "Patenting in England, Scotland and Ireland during the Industrial Revolution, 1700-1852," *Explorations in Economic History*, 54, 48–63. 6

——— (2019): "The Returns to Invention during the British Industrial Revolution," *The Economic History Review*, 72, 510–530. 6

Bouscasse, P., E. Nakamura, and J. Steinsson (2025): "When Did Growth Begin? New Estimates of Productivity Growth in England from 1250 to 1870," *Quarterly Journal of Economics*, 140, 835–888. 6

Broadberry, S., B. M. S. Campbell, A. Klein, M. Overton, and B. van Leeuwen (2015): *British Economic Growth, 1270–1870*, Cambridge: Cambridge University Press. 6

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. 6

Burton, M. D. and T. Nicholas (2017): "Prizes, Patents and the Search for Longitude," *Explorations in Economic History*, 64, 21–36. 6

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. 5

Caprettini, B. (2026): "You Only Weave Twice: Espionage and Industrialization in 18th Century France," Mimeo. 5

Cervellati, M. and U. Sunde (2005): "Human Capital Formation, Life Expectancy, and the Process of Development," *American Economic Review*, 95, 1653–1672. 5

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. 5

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. F. (1977): "Industrial Revolution in England and France: Some Thoughts on the Question,"Why Was England First?"," *The Economic History Review*, 30, 429–441. 2, 6

——— (1995): "Macroinventions, Economic Growth, and Industrial Revolution'in Britain and France," *The Economic History Review*, 48, 591–598. 6

Crouzet, F. (1996): "La Première Révolution Industrielle," in *Histoire de La France Industrielle*, ed. by M. Lévy-Leboyer, Larousse. 6

——— (2003): "The Historiography of French Economic Growth in the Nineteenth Century," *Economic History Review*, LVI, 215–242. 27

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. 5

Dittmar, J. E. (2011): "Information Technology and Economic Change: The Impact of the Printing Press," *Quarterly Journal of Economics*, 126, 1133–1172. 5, 34

Dittmar, J. E. and R. R. Meisenzahl (2024): "The Research University, Invention, and Industry: Evidence from German History," Working paper. 5

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. 2

Galor, O. (2011): *Unified Growth Theory*, Princeton University Press. 5

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. 5

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," *Economic Journal*, 135, 1749–1781. 6, 13, 14, 56

Hansen, G. D. and E. C. Prescott (2002): "Malthus to Solow," *American Economic Review*, 92, 1205–1217. 5

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, 59

Harris, J. (2017): *Industrial Espionage and Technology Transfer*, London: Routledge. 3, 6, 7

Horrell, 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," Working paper. 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

Jones, C. (1995): "R & D-Based Models of Economic Growth," *Journal of Political Economy*, 103, 759–784. 9

Juhász, R. (2018): "Temporary Protection and Technology Adoption: Evidence from the Napoleonic Blockade," *American Economic Review*, 108, 3339–76. 5, 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. 5, 34

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. 5

Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy (2021): "Measuring technological innovation over the long run," *American Economic Review: Insights*, 3, 303–320. 19, 50

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. 6

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. 6

Khan, Z. (2005): *The Democratization of Invention: Patents and Copyrights in American Economic Development, 1790-1920*, Cambridge: Cambridge University Press. 6

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. 3, 6, 19

Liu, E. and S. Ma (2023): "Innovation Networks and R&D Allocation," NBER Working Paper 29607. 0, 1, 5, 7, 8, 9, 10, 17, 23, 24, 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. 6

Markovitch, T. J. (1965): "L'industrie Française de 1789 à 1964. Sources et Méthodes," *Cahiers de L'I.S.E.A.*, AF4. 76

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. 34

Miura, H. (2012): "Stata Graph Library for Network Analysis," *Stata Journal*, 12, 94–129. 53

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 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, 36

——— (2005): "The Intellectual Origin of Modern Economic Growth," *Journal of Economic History*, 65, 285–351. 1, 2, 4, 36

——— (2009): *The Enlightened Economy*, New Haven and London: Yale University Press. 1, 2, 36

Moser, P. (2005): "How Do Patent Laws Influence Innovation? Evidence from Nineteenth-Century World's Fairs." *American Economic Review*, 95, 1214–1236. 6, 16, 55

——— (2012): "Innovation without Patents: Evidence from World's Fairs," *Journal of Law and Economics*, 55, pp. 43–74. 2, 6, 16, 55

Newman, M. E. J. (2010): *Networks: An Introduction*, Oxford University Press. 53

Nicholas, T. (2010): "The Role of Independent Invention in U.S. Technological Development, 1880–1930," *The Journal of Economic History*, 70, 57–82. 2

——— (2011): "Cheaper Patents," *Research Policy*, 40, 325–339. 6

North, D. C. and B. R. Weingast (1989): "Constitutions and Commitment: The Evolution of Institutions Governing Public Choice in Seventeenth-Century England," *The Journal of Economic History*, 49, 803–832. 1

Nuvolari, A. and V. Tartari (2011): "Bennet Woodcroft and the Value of English Patents, 1617-1841," *Explorations in Economic History*, 48, 97–115. 6, 15, 59

Nuvolari, A., V. Tartari, and M. Tranchero (2021): "Patterns of Innovation during the Industrial Revolution: A Reappraisal Using a Composite Indicator of Patent Quality," *Explorations in Economic History*, 82. 6, 15, 16, 59, 63

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. 5, 6

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. 5

Pomeranz, K. (2000): *The Great Divergence*, Princeton: Princeton University Press. 2

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. 5

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, 59

Squicciarini, M. P. (2020): "Devotion and Development: Religiosity, Education, and Economic Progress in Nineteenth-Century France," *American Economic Review*, 110, 3454–3491. 5

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. 5

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. 6

——— (1990): "The Revolution of Ideas: Widespread Patenting and Invention during the English Industrial Revolution," *The Journal of Economic History*, 50, 349–362. 6

Thomas, M. (1984): "An Input–Output Approach to the British Economy, 1890–1914," Ph.D. thesis, Oxford. 15, 45, 76

Trajtenberg, M. (1990): "A Penny for Your Quotes: Patent Citations and the Value of Innovations," *The RAND Journal of Economics*, 21, 172–187. 15, 59

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. 2

Voth, H.-J., B. Caprettini, and A. Trew (2026): "Fighting for Growth: Labor Scarcity and Technological Progress During the British Industrial Revolution," *American Economic Review*, 116, 3837–3874. 2

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 *sub*categories correspond to French *sub*categories (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 difficulties in identifying them are 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 we 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 in principle 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 use of patent agents in Britain (in which case the British patent lists a different inventor than the French "British" patents).[41]

[41] Curiously, 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 *sub*categories correspond to French *sub*categories (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 from 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.[42]

[42]To 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.[43]

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

[43]Some professions are ambiguous even if qualified by a topic, for example "coal merchant" or "cloth merchant" because we do not know whether these individuals worked in industry or in the excluded distribution services. One exception to the rule is composite occupations like "woollen manufacturer and merchant" because there the "manufacturer" clearly indicates involvement in production.

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

	Most common occupation theme		
Input–output industry	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 splitting the occupation string into 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.[44] The resulting probabilistic mapping from technology categories to industries appears 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 makes sense: For example, for the industry *Agriculture, Forestry, etc*, the top technology category is *Agriculture* with $\alpha = 0.65$ and the second technology category is *Power and propulsion* with $\alpha = 0.07$. 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 the *Cotton and Silk* and *Woollen and Worsted* industries.

[44]Two 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	0.65	Power and Propulsion	0.07
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	0.09
Non-ferrouse 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	0.32
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	Sugar Manufacture	0.15	Flour milling	0.13
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	Fabric Manufacture	0.5	Rubber and Gutta-Percha	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	0.39	Stone and Plaster Work	0.1
Building, etc.	Construction	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 either on citations or on the inventor-based approach 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 (USPTO). The result is a 123×123 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 Validating the network measure using historical patent text similarity

This appendix describes a second approach for validating our method, one that focuses on the historical context that we study. Since patent citations are not available in that context, we instead focus on patent texts using methods developed by Kelly et al. (2021).

This analysis starts with the texts for all British patents from 1700–1850 compiled by Berkes, Chen, and Tranchero (2026) and available at Berkes et al. (2025). This dataset includes patent identifiers and patent texts, with the texts already helpfully broken down into word tokens. We begin by calculating two objects. The first is term frequency, defined as,

$$TF_{pw} = \frac{c_{pw}}{\sum_{w'} c_{pw'}}$$

where c_{pw} represents the count of term w in patent p. This statistic reflects the frequency of term w in the text of patent p. The second is the backward-inverse-document-frequency,

$$BIDF_{wt} = \log\left(\frac{\#\text{ patents filed before } t}{1 + \#\text{ patents filed before } t \text{ that contain } w}\right)$$

This object is used to overweight terms that most clearly distinguish the content of the patent from previous patents or, conversely, to deflate the importance of terms that appear regularly in prior patents. We then calculate,

$$TFBIDF_{pw} = TF_{pw} \times BIDF_{w,t(p)}$$

where $t(p)$ is the filing year of patent p. This captures the importance of term w in patent p, deflated by how common the term was in earlier patents, i.e., by the share of patents filed before $t(p)$ that contain w. We collect these values over all terms into the vector $TFBIDF_p$ and normalize it to unit length:

$$V_p = \frac{TFBIDF_p}{\|TFBIDF_p\|}$$

Next, we calculate the cosine similarity, which lies in [0,1] when all term weights are non-negative, between the normalized vectors for every pair of patents p and p' filed within 20 years of each other, where patent p' followed patent p:[45]

$$sim_{pp'} = V_p \cdot V_{p'}$$

Kelly et al. (2021) show that this statistic is correlated with citations between patents.

[45] In practice, we require that patent p' was filed at least one year but no more than 20 years after p. We do not include pairs filed in the same year.

We will use it as the primary building block for our text-similarity-based spillover measure.

Finally, we need to aggregate the patent-pair-level similarity measure $sim_{pp'}$ to the technology-category-pair level. To do this, we calculate,

$$SIM_{j\to i} = \frac{\sum_{p\in j}\sum_{p'\in i} sim_{pp'}}{TOT_j + TOT_i}$$

where TOT_j is the sum of $sim_{pp'}$ over all pairs with the earlier patent p in j, and TOT_i is the sum of $sim_{pp'}$ over all pairs with the later patent p' in i. Note that this is a directed similarity matrix, where the direction runs from earlier (sending) patents in j to later (receiving) patents in i. The inclusion of the two terms in the denominator is important to ensure that the similarity measure is not driven by the number of patents in particular technology categories.

On visual inspection, the $SIM_{j\to i}$ statistic yields generally reasonable results. For example, the most similar category to spinning according to this measure (other than spinning itself) is fabric production and composition, and the most similar category to steam engines is motive power and propulsion, followed by fuel.

Table D.1 presents a set of regression results comparing our measure of spillovers between technology categories, ω_{ij}, to the text-similarity-based measure, $SIM_{j\to i}$. We can see that our measure is strongly correlated with the text-similarity-based matrix. This provides a second piece of evidence validating our measure of the innovation network.

Table D.1: Validating our method against text-similarity based spillovers

Dependent variable: patent-text-based spillover measure				
	$SIM_{j\to i}$	$SIM_{j\to i}$	$SIM_{j\to i}$	$\ln SIM_{j\to i}$
	(1)	(2)	(3)	(4)
ω_{ij}	0.02337***	0.01133***	0.01738***	
	(0.003265)	(0.003328)	(0.002732)	
$\ln \omega_{ij}$				0.1214***
				(0.006339)
Sample	All	$\omega_{ij} > 0$	All	All
Technology category FEs			Yes	Yes
Observations	15,128	3,781	15,128	3,779
R^2	0.093	0.029	0.547	0.811

OLS regressions, robust standard errors in parentheses. $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

E Additional results on the innovation network

E.1 Network plots

Figure E.1 presents a fully labeled version of Figure 3.

Figure E.1: Fully labeled joint innovation network



This is the fully labeled version of Figure 3. Plot generated using multidimensional scaling. The edges of the joint network are computed as $(\omega_{ij}^{UK} + \omega_{ij}^{FR})/2$.

E.2 Network comparison

Table E.1 presents 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 E.1: 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	0.803***	0.713***	0.629***	0.525***	0.386	0.832*	0.469***	0.541***
	(0.098)	(0.081)	(0.064)	(0.082)	(0.264)	(0.421)	(0.100)	(0.160)
N (Obs = nodes)	118	113	118	113	118	113	118	35
R^2	0.616	0.382	0.515	0.237	0.025	0.291	0.185	0.311

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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

E.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:

$$\omega_{ijt} = \bar{\omega}_{ij} + \epsilon_{ijt} \tag{15}$$

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

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

$$\omega_{ijt} = \rho\omega_{ijt-1} + e_{ijt} \tag{16}$$

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:

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

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 $\sigma = \rho - 1$, then testing whether the model is a random walk amounts to testing the null hypothesis that $\sigma = 0$. Alternatively, the model in Eq. 15 would suggest that we should observe $\sigma = -1$.

Table E.2: 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	−0.829***	−0.608***	−0.971***	−0.829***	
	(0.032)	(0.094)	(0.008)	(0.034)	
Edge $\omega_{ij,t-1}$: 1790–1823					−0.797***
					(0.036)
N (Obs = edges)	9991	94	9897	8281	9207
R^2	0.56	0.38	0.68	0.43	0.31

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 E.2 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: $\sigma = 0$) and the white noise model (H_0: $\sigma = -1$), though the fact that we obtain estimates of σ that are fairly close to -1, i.e. ρ close to zero, suggests that the model is closer to the white noise model in Eq. 15 than the random walk in 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. Thus, the matrix we observe is fairly close to one characterized by a fixed underlying structure.

F Innovation network based on exhibition data

This appendix details our analysis of data on technologies exhibited at the Crystal Palace Exhibition in London in 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.[46]

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 exhibit by the same exhibitor. For French exhibitors, 47 exhibits were 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.

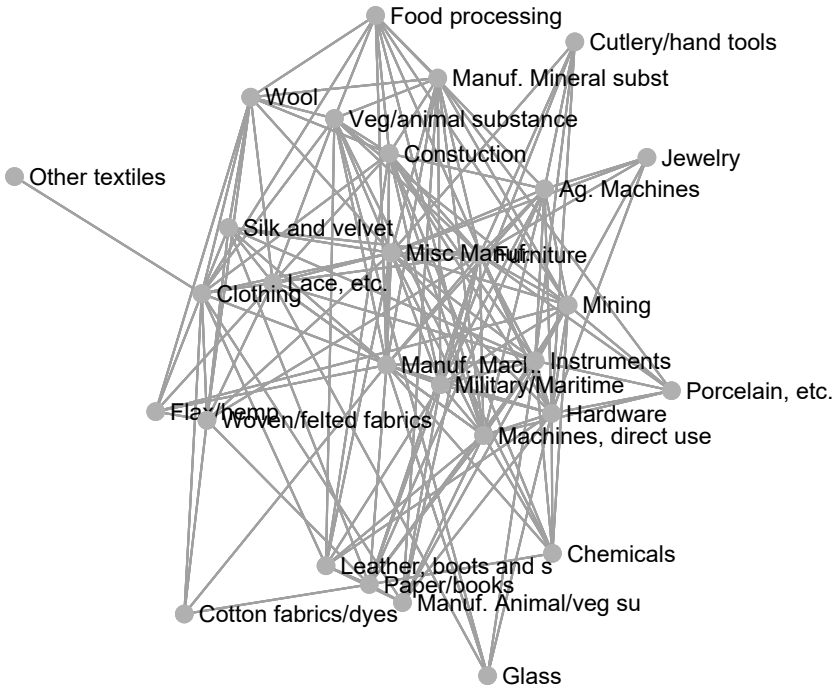
Figure F.1 presents the innovation network obtained by applying this method to the exhibition data.[47] The network features a dense central core, located just below and to the right of the center of the plot. The core includes technology types such as *Machines for direct use* (including steam engines and locomotives), *Manufacturing machines*, *Military/Maritime* (including shipbuilding), *Instruments*, *Mining*, and *Hardware*. Similarities between this 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

[46] The 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. But as most exhibits have only one category, we focus on each exhibit's first category.

[47] Note 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 F.1: The technology network obtained from exhibition data



Plot is generated using multidimensional scaling.

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 who also patented, with links for 566 individuals in total (some individuals account for more than one exhibit).[48]

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

[48]Note 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 F.1: Comparing patent-based and exhibit-based innovation networks

	Dep var: Exhibition eigencentrality		
	Level	Log	Rank
	(1)	(2)	(3)
Patent eigencentrality	0.091***	0.041***	0.105**
	(0.028)	(0.014)	(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–exhibitors with multiple patents or exhibits, eigencentrality is averaged. OLS regressions, robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

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 F.2 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 F.2: 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 of 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.

G 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.

G.1 Historical measures of patent quality

For British patents, we use measures provided by two standard sources, Nuvolari and Tartari (2011) 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 (2011), 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.[49]

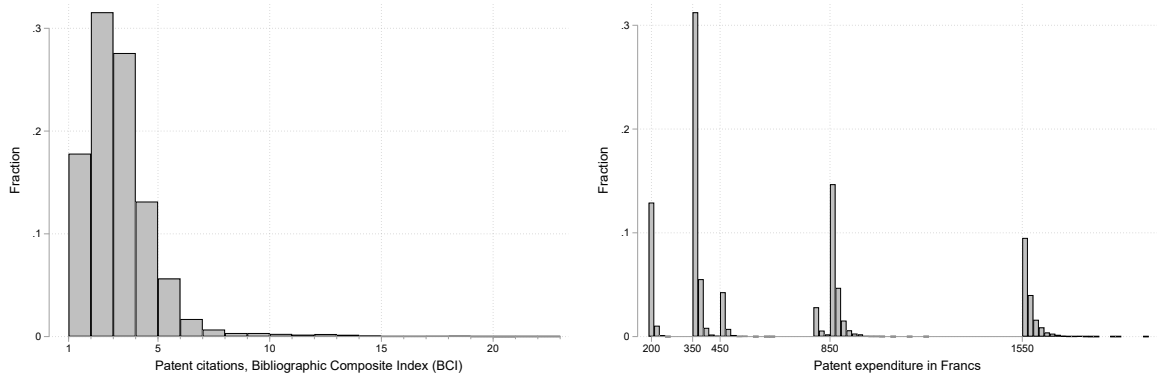
Nuvolari et al. (2021) update and augment the index from Nuvolari and Tartari (2011) by collecting, from a set of modern reference works on historical technology, modern citations to the patent and mentions of the inventor. This augmented patent quality index is referred to as the Bibliographic Composite Index (BCI).[50] Nuvolari et al. (2021) argue that the BCI provides a more accurate measure of patent quality than the WRI, 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 G.1 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 1,500

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

[50] We 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.

Francs. 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 the first half a 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 inventions that proved to be more successful. Together, these different fees partition patents into bins: 5 or 10 or 15 year terms, each of which could be not renewed or renewed, and the renewed ones could be without or with additions. The right-hand panel of Figure G.1 presents a histogram for the renewal-fee-based patent quality measure for France. Note that there are really just 9 fee levels here, though it appears that there are more because some fee categories are split across multiple plot bins. However, some of these bins are fairly sparsely populated, so for our analysis we will combine them into five bins of increasing cost.



(a) British patent citations (BCI)

(b) French patent fee expenditures

Figure G.1: Raw quality measures in Britain and in France

Table G.1 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 a 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.

G.2 Alternative quality adjustments

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

Table G.1: 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 (without and with addition)	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 renewed 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:

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

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.

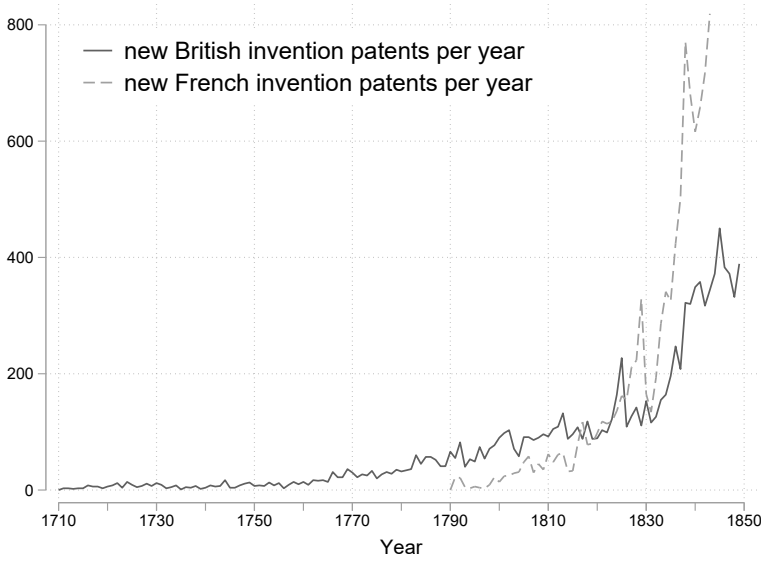
G.3 Descriptive patterns after adjusting for patent quality

Figure G.2 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.

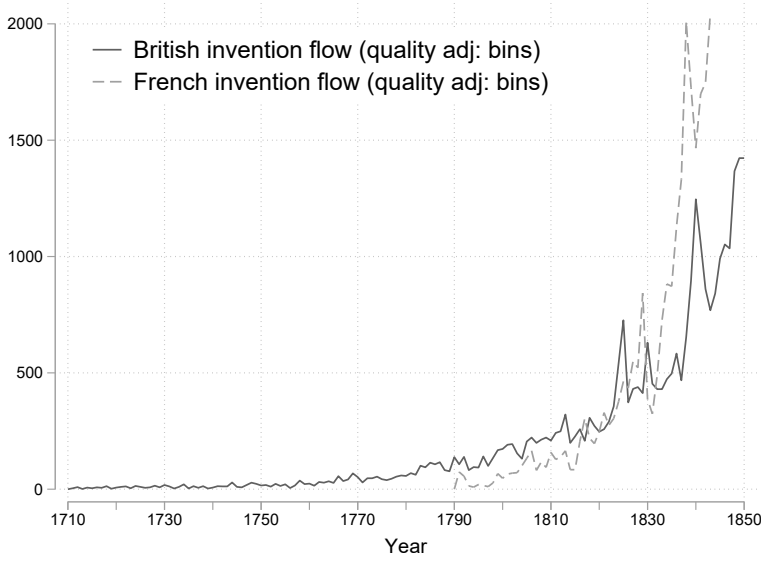
G.4 Robustness to alternative quality adjustments

Table G.2 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.[51] Column 3 applies the "$\mu-\sigma-$equating" approach described in subsection G.2. Column 4 uses the same quality measure as in Column 3, but uses British patents from 1791 instead

[51]We normalize the scale by setting the lowest level to one.



(a) Patent counts



(b) Adjusted via quality bins

Figure G.2: Invention patents over time in Britain and France

Table G.2: Patent quality: Further robustness for main estimation

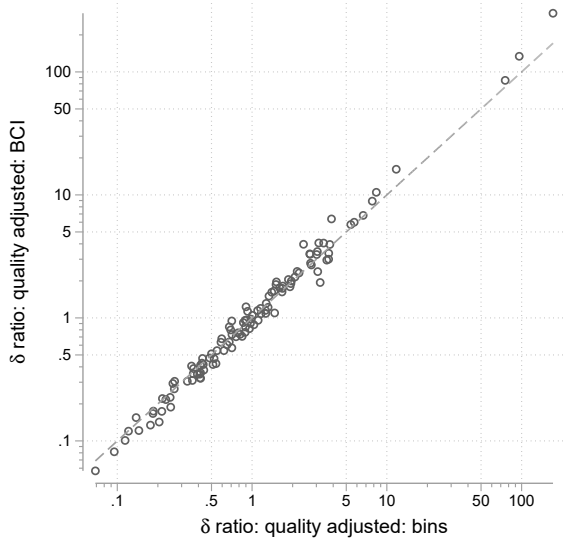
	Dep var: $\Delta \ln q_{c,i,t}$						
	(1) BCI bin (baseline)	(2) BCI bin, expend.	(3) BCI μ–σ	(4) BCI μ–σ, q_0^{UK}1790	(5) WRI, binned	(6) WRI μ–σ	(7) PCA μ–σ
$\ln \chi_{t-1}$	0.066***	0.058**	0.065**	0.066***	0.068***	0.068***	0.065**
	(0.023)	(0.024)	(0.026)	(0.024)	(0.023)	(0.026)	(0.028)
$\ln \chi'_{t-1}$	0.022**	0.021**	0.023**	0.021*	0.022**	0.023**	0.024**
	(0.009)	(0.010)	(0.009)	(0.013)	(0.009)	(0.010)	(0.009)
$\ln q_{t-1}$	−0.110***	−0.110***	−0.116***	−0.128***	−0.110***	−0.116***	−0.116***
	(0.012)	(0.013)	(0.016)	(0.014)	(0.012)	(0.015)	(0.016)
N	9152	9152	9152	9043	9152	9152	9152
R^2	0.33	0.31	0.31	0.32	0.33	0.32	0.32
F-stat	24.9	21.9	22.5	22.3	24.9	22.9	23.4
implied λ	0.092	0.083	0.093	0.091	0.094	0.096	0.093
implied ι	0.754	0.730	0.738	0.755	0.755	0.746	0.733
implied κ	0.246	0.388	0.311	0.466	0.228	0.266	0.298

This 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 μ–σ–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 μ–σ–equating. Column 7 uses the British principal component of all three Nuvolari et al. (2021) indicators and maps French patent expenditures onto them via μ–σ–equating.

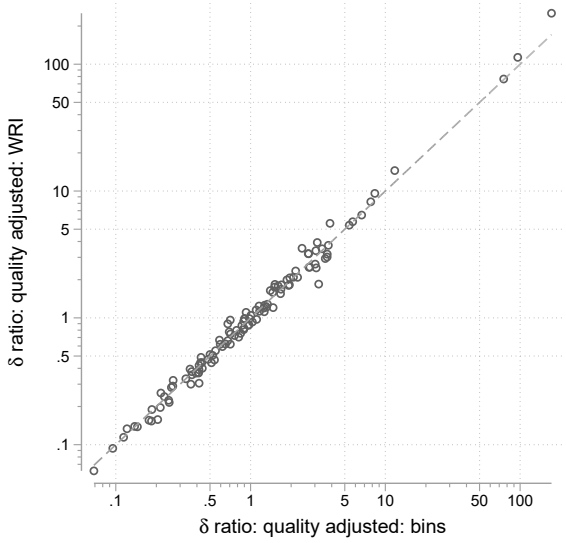
of 1761. Column 5 returns to our preferred "binned" approach to relating the two patent quality measures, but uses the WRI quality index in place of the BCI index to measure the quality of British patents. Column 6 uses the British WRI citations and maps French patent expenditures onto them via μ–σ–equating.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

In Figure G.3, 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 with alternative methods of adjusting (or not adjusting) for patent quality (y-axis). Panel G.3a compares our baseline against the BCI citation measure, onto which we map French patent expenditures via μ–σ–equating. Panel G.3b compares our baseline against the WRI citation measure (also using μ–σ–equating). Panel G.3c compares the

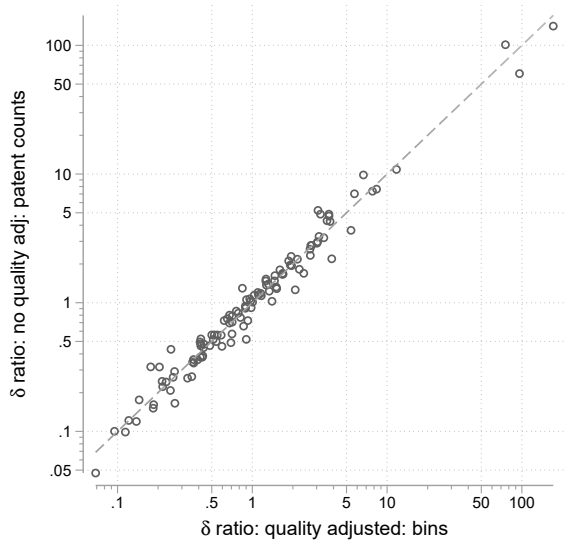
baseline against the case of using raw patent counts, i.e., not adjusting for quality at all. All alternative methods yield very similar knowledge access ratios compared to the baseline.



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



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



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

Figure G.3: Robustness of δ ratios to adjusting for patent quality

H Analysis using Google Ngrams data

This appendix provides additional details on the construction of the Google Ngrams data used in Section 9 as well as some additional results.

H.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.[52] We then run univariate patent-level regressions looking at how well each unigram or bigram predicts each technology category.[53] For each gram, we retain the technology category that received the highest t-statistic, and additionally apply 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.[54] We then take these translated terms back to the patent data and look at how predictive they are of the technology category. Only when a keyword and its translation are strongly predictive of the same technology category, using the same criteria as above, do we consider that a useful keyword pair.

To give a sense of the results of this exercise, Table H.1 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 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 Ngrams. For French

[52] We exclude frequent stop words, such as determiners, prepositions, and linking verbs.

[53] We use 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.

[54] We 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. The batch script with the full prompt is available in the replication files.

Table H.1: Examples for equivalent keyword pairs

English term	French term
Technology category: Spinning	
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
Technology category: Agriculture	
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 keyword pairs that our algorithm generates from patent titles.

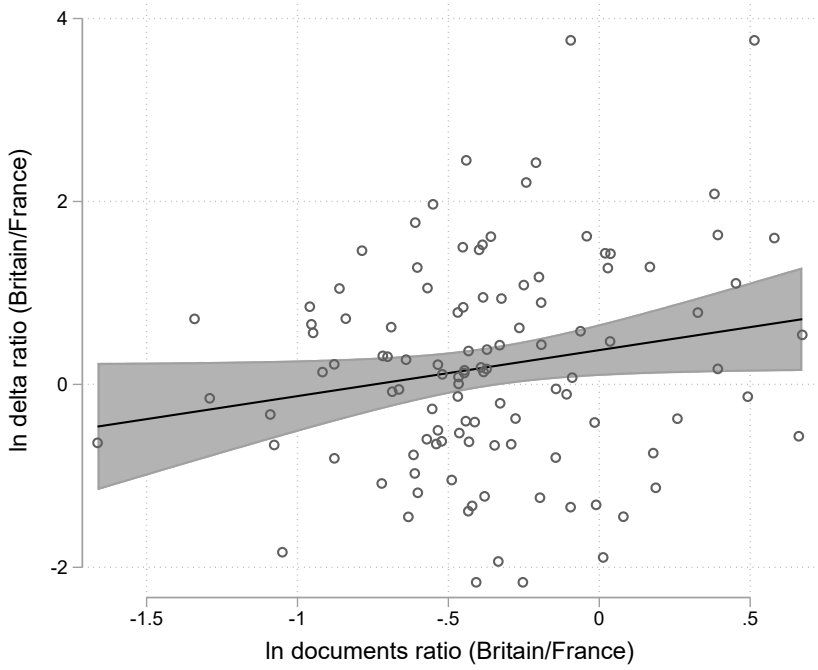


Figure H.1: Scatter plot, corresponding to Column 1 of Table 4

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 1800. From Ngrams, we recover the number of documents associated with each keyword.

H.2 Additional results

Table 4 in the main text presents our baseline analysis of the Google Ngrams data. In this subsection, we provide some additional results and robustness checks supporting those in the main text.

In Figure H.1, we present a scatterplot corresponding to the specification shown in Column 1 of Table 4 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 H.2 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 δ 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 H.2: Results estimated at the keyword-pair level

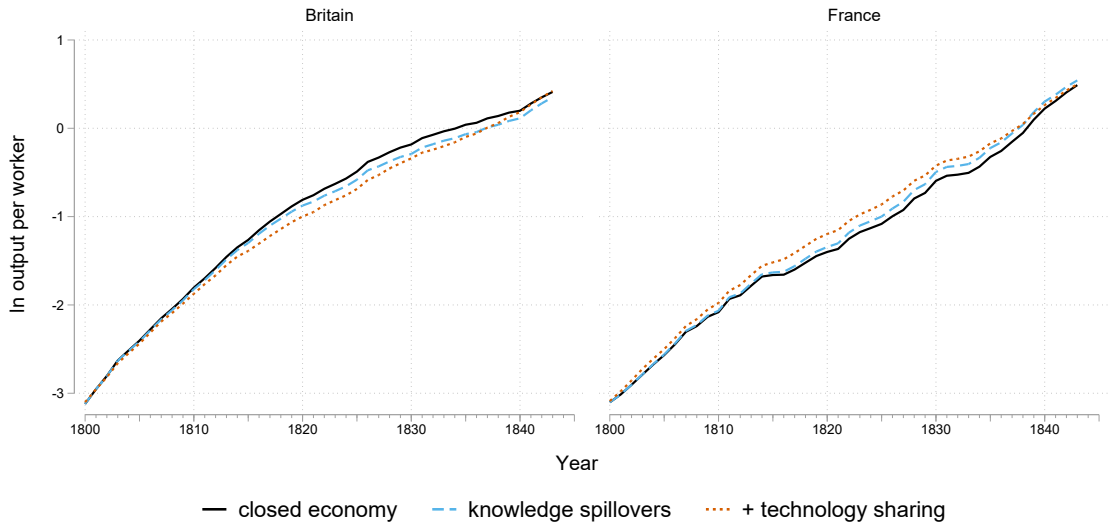
	Dep var: Ln δ^B/δ^F ratio				
	Baseline (1)	Patent counts (2)	q^{UK} 1790 (3)	Late 1816–43 (4)	Endogenous growth (5)
Panel A: Document ratio					
Ln document ratio 1750–1799	0.085	0.100*	0.117**	0.097**	0.068
	(0.054)	(0.053)	(0.058)	(0.046)	(0.043)
N (Obs = gram)	974	974	974	974	974
Clusters (category)	109	109	109	109	109
R^2	0.009	0.011	0.013	0.015	0.009
Panel B: English and French documents					
Ln English documents 1750–1799	0.088	0.095*	0.111*	0.092*	0.071
	(0.057)	(0.055)	(0.062)	(0.048)	(0.045)
Ln French documents 1750–1799	−0.084	−0.101*	−0.118**	−0.098**	−0.068
	(0.055)	(0.053)	(0.059)	(0.046)	(0.044)
N (Obs = gram)	974	974	974	974	974
Clusters (category)	109	109	109	109	109
R^2	0.009	0.012	0.013	0.015	0.009

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

I Model simulation appendix

This appendix presents additional simulation results. Figure I.1 plots log output per worker separately for the two simulated economies, corresponding to the values used to create the log output per worker ratio in main text Figure 6. We can see that growth is initially faster in the UK, but this slows down after 1820, in part because of the technology drag effect induced by $\kappa > 0$. In contrast, output increases more rapidly in France after 1820 as the innovation environment there improves.

Figure I.1: Simulated log output per worker separately



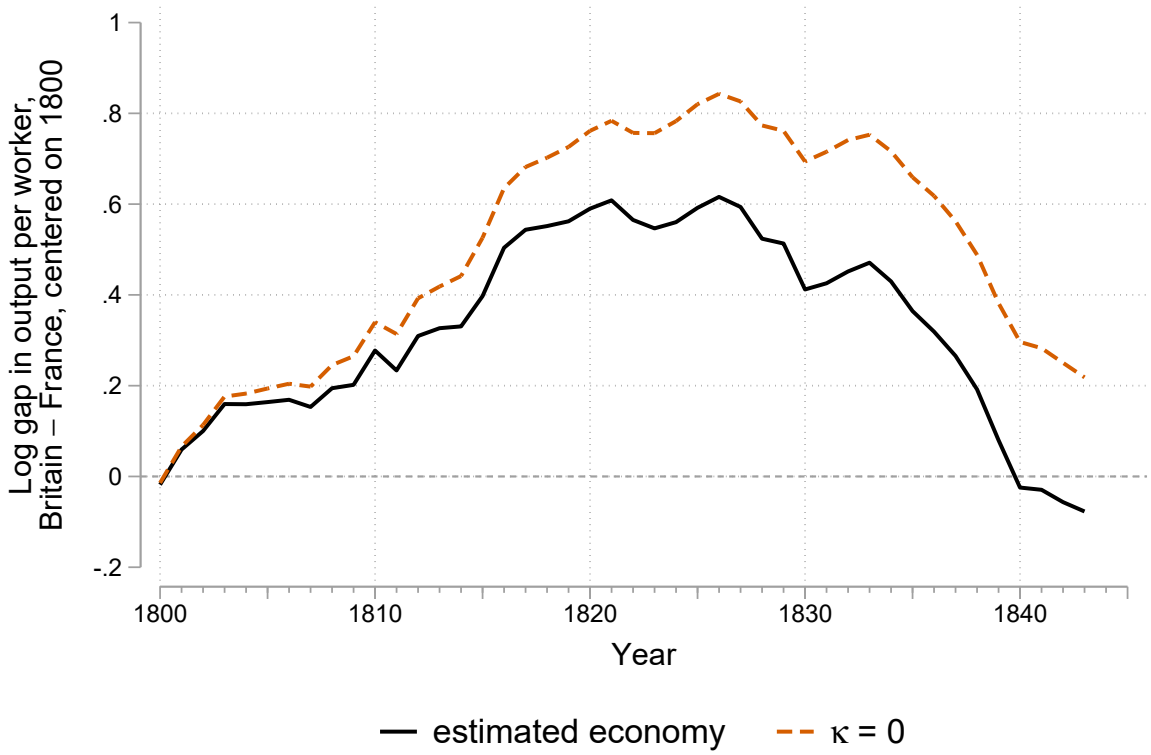
The figure plots log output per worker in our two simulated economies for the closed economy case, the case with cross-country knowledge spillovers, and the case with both knowledge spillovers and technology sharing. The model uses the estimated parameters corresponding to the specification in Column 1 of Table 2.

Figure I.2 shows simulation results for our baseline simulation and a simulation where we have set $\kappa = 0$. This shows that allowing for semi-endogenous growth has a substantial impact on our quantitative results.

In our main simulation results, we turn off the impact of the innovation network by setting all elements of the matrix to be equal. That means that each technology category generates an equal amount of spillovers for all other categories. An alternative way to think about turning off the effect of the innovation matrix is to set the matrix to the identity. In that case, each technology benefits only from its own knowledge stock and there are no cross-technology spillovers. In the endogenous growth case of $\kappa = 0$, these two effects cancel out; in the case of $\kappa > 0$, which is the empirically relevant case, the drag from the own knowledge stock as ideas are getting harder to find dominates.

Figure I.3 shows simulations for the closed economy case when we take this alternative approach. We can see that when we set the innovation network to the identity, in the gray

Figure I.2: Simulation results for endogenous and semi-endogenous growth scenarios

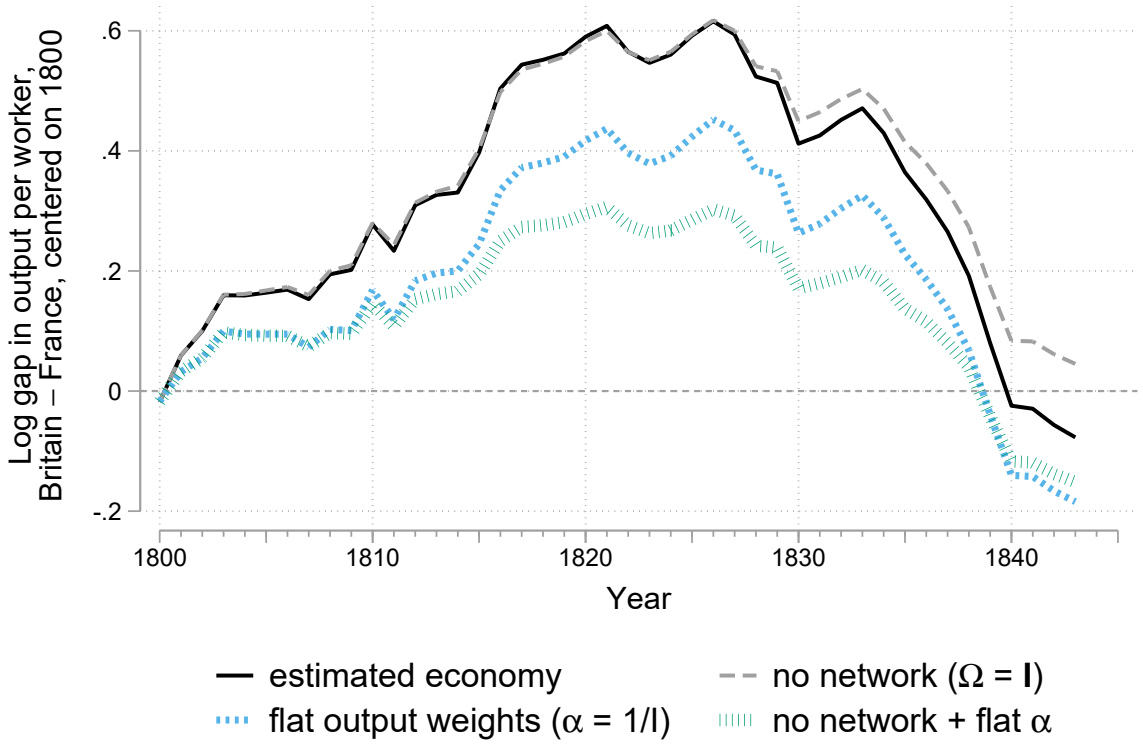


The figure plots the log output per capita gap in our two simulated economies for the closed economy case, either with our estimated κ or with $\kappa = 0$. The model uses the estimated parameters corresponding to the specification in Column 1 of Table 2.

line, it has very little impact on Britain's output advantage relative to the baseline (black line). This suggests that Britain's advantage remains when spillovers are concentrated within technologies. However, it is interesting to see that this does matter when combined with flat output weights, in the bottom line. That suggests that Britain's advantage when spillovers are confined to being within-technology is driven at least in part by the fact that British inventors had easier access to technologies that serve a larger share of demand.

We also examine robustness to the initial conditions used in the simulation. The simulation in the main text uses a uniform initial quality level across all technology types ($q = 10$). In Figure I.4 we present results for a variety of different initial conditions, focusing on the closed economy case. The heavy black line shows results using a uniform initial quality level $q = 10$ for all technologies. The simulation is indifferent to the exact level, but the assumption of uniformity is meaningful. In the remaining lines, we consider alternatives where the initial technology levels vary. To generate realistic variation, we use patents in the UK up to 1790 to generate initial q values for each technology type. We then add to these values a constant to ensure that $q > 0$ for all technologies. We can see that this impacts the output gap somewhat, but the overall patterns are very similar to the

Figure I.3: Alternative results setting the innovation network to the identity

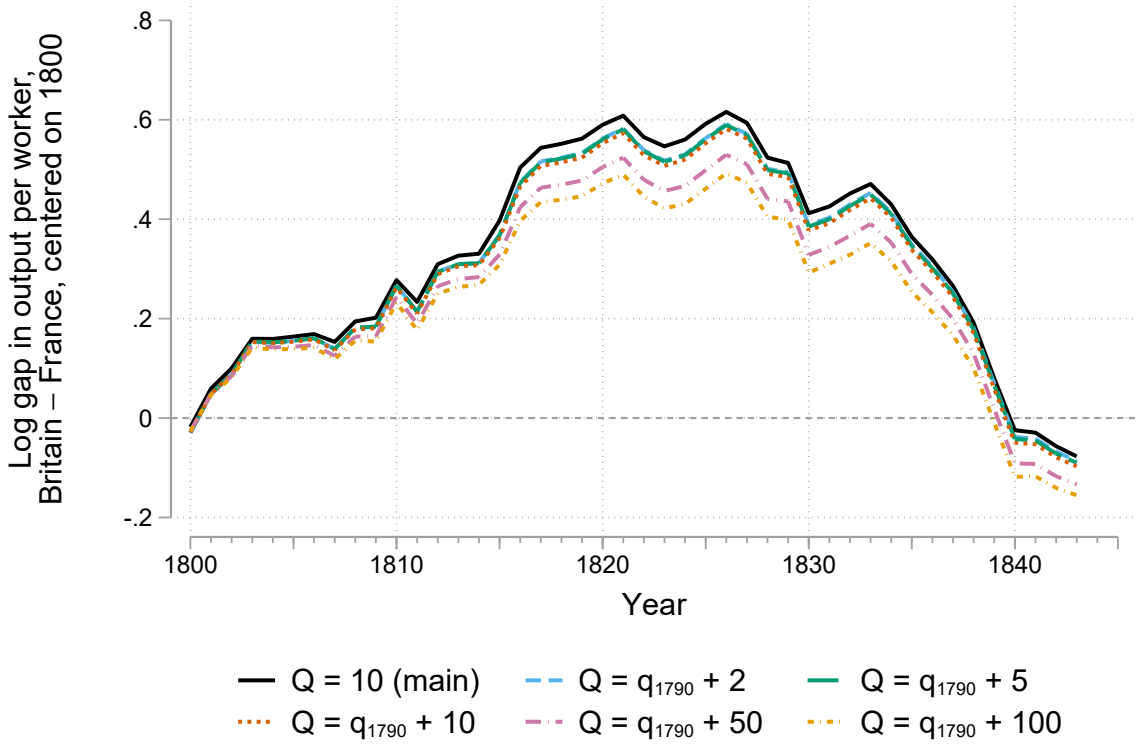


The figure plots the difference in log output per worker in two simulated economies, one with the estimated parameters of Britain and the other with the estimated parameters of France, starting from a common initial technology level. The model uses the estimated parameters corresponding to the specification in Column 1 of Table 2. The estimated economy corresponds to the closed economy scenario depicted in the main Figures 6 and 7. The counterfactual that equalizes generality in production ($\alpha_{ik} = 1/I \, \forall i$, blue short dashed line) is the same as in Figure 7. What differs here is the counterfactual for the innovation network, which here is turned off by setting the network to the identity, $\Omega = \mathbf{I}$, either in isolation (gray dashed line) or together with equalizing production demand (green hashed line).

uniform version.

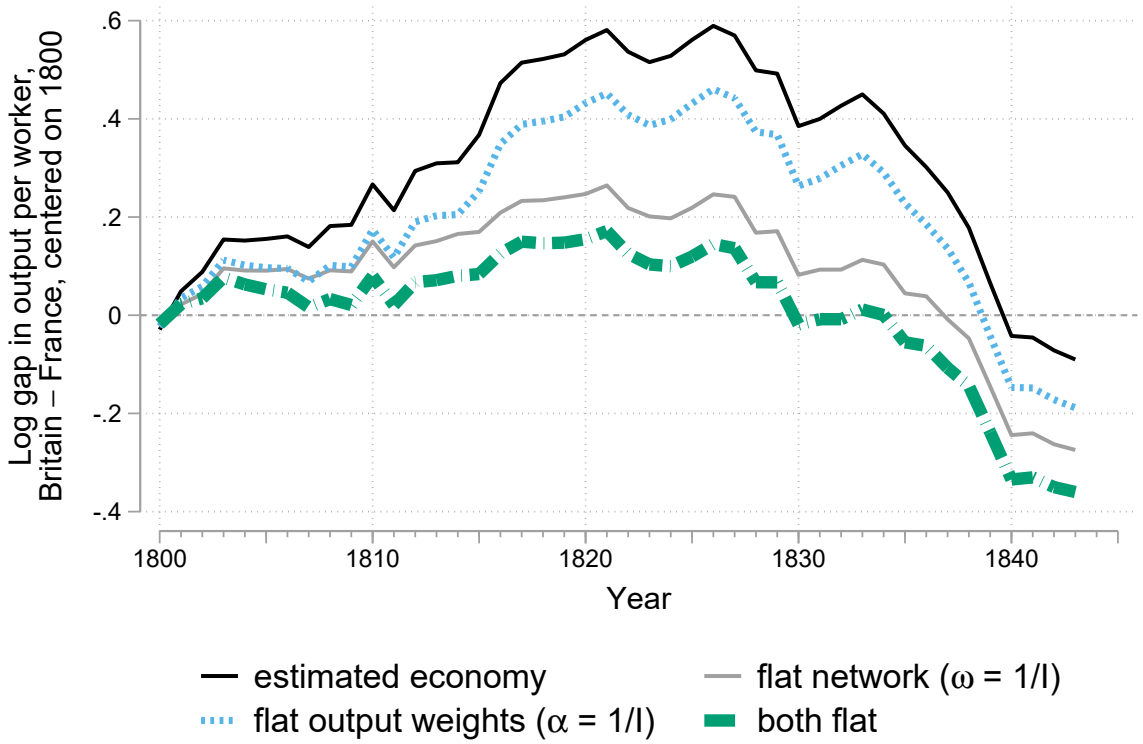
To show how initial conditions affect our counterfactuals, Figure I.5 presents counterfactual results using as the initial condition the $q_{1790} + 5$ case. It is interesting to see that under these conditions, the impact of turning off the innovation network or demand differences across technology types is even larger than in our baseline counterfactuals. The same is true in other counterfactuals where we allow for differences in initial technology levels. This suggests that our baseline counterfactual is providing a relatively conservative estimate of the importance of the distribution of innovation frictions across technology types for Britain's overall growth advantage.

Figure I.4: Simulations using alternative initial conditions



The figure plots the difference in log output per worker in the two simulated economies under different initial conditions. The top line shows results under a uniform q of 10 in all technologies, the approach used in the main text. The remainder use patents from the UK up to 1790 to generate initial quality in each technology type (q_{1790}) plus a constant to ensure that $q > 0$ for all technologies.

Figure I.5: Simulation counterfactuals using alternative initial conditions ($q_{1790} + 5$)



This figure shows simulated log output per worker gaps in different counterfactuals using initial technology levels based on UK patents up to 1790 (q_{1790}) plus a constant, 5, to ensure that all $q > 0$.

J Estimating the technology sharing parameter (η)

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

J.1 Deriving the estimating equation

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

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

We assume that a share of domestic varieties is under monopoly, while foreign varieties are produced with domestic labor under perfect competition:

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

where in the last line, we defined $Q_{kt} \equiv \sum_i \alpha_{ik} \ln q_{it}$, $Q'_{kt} \equiv \sum_i \alpha_{ik} \ln q'_{it}$, $\Phi_{kt} \equiv \sum_i \alpha_{ik} \ln(1+\lambda)\,(r_{it} \tilde{X}_{it} / (q_{it}^{1+\kappa} + r_{it} \tilde{X}_{it}))$ with $\tilde{X}_{it} = X_{it}^{\iota}\, X_{it}^{\prime(1-\iota)}$, and $A_k \equiv \sum_i \alpha_{ik} \ln \alpha_{ik} \beta_k$, and used that $\sum_i \alpha_{ik} = 1$. We then take the long difference between two periods, $\tau = T$ and $\tau = t_0$:

$$\ln y_{kT} - \ln y_{k\,t_0} = \eta\, Q_{kT} + (1-\eta)\, Q'_{kT} - \eta\, Q_{k\,t_0} - (1-\eta)\, Q'_{k\,t_0} - \eta\, \Phi_{kT} + \eta\, \Phi_{k\,t_0} + const \tag{18}$$

where the constant absorbs the change of the number of workers, $\ln \bar{L}_T - \ln \bar{L}_{t_0}$. The remaining terms on the right-hand side, Q, Q', and Φ, can be calculated from data, the measured network and demand parameters, and the estimates from Table 2.

J.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 come in two forms: One is a quantity index, the other is the value of production.[55] 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.[56]

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*.[57] As non-zero output is first reported for *Gas* and *Rubber* in the 1840 (1835–44) decade, our balanced estimation panel includes 14 industries.

J.3 Estimation

The empirical counterpart of Equation 18 is

$$\Delta \ln y_k = \quad b_0 + b_1 Q_{kT} + b_2 Q'_{kT} - b_3 Q_{k\,t_0} - b_4 Q'_{k\,t_0} - b_5 \Phi_{kT} + b_6 \Phi_{k\,t_0} + \varepsilon_k \,, \tag{19}$$

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 as defined above. Our main goal is to estimate η. Since $N = 14$, we apply two parsimonious estimation approaches: without any restriction, or with the structural constraint that $\eta = b_1 = 1 - b_2$, the coefficients on Q_{kT} and Q'_{kT} that denote the weighted sums of domestic (French) and foreign (British) knowledge stocks in T. For the analysis period, we choose $T = 1840$ and $t_0 = 1810$, which is the longest difference for which we have both patents and output available.

Table J.1 reports estimates of the technology sharing parameter using various specifications and data. In Panel A, we report unconstrained coefficients on b_1 and b_2, and in

[55] The 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). We cannot back out the actual production levels since neither the quantities of 1938 nor the period prices are reported.

[56] The early data gaps are due to the French Revolution and the Wars of the Sixth and Seventh Coalitions.

[57] Three 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.

Panel B, we use the structural constraint that $\eta = b_1 = 1 - b_2$. In column 1, we use the data and estimates corresponding to the baseline estimate in Table 2, column 1. The OLS coefficient is $b_1 = 0.724$, which is close to $b_1 = \eta = 0.693$ when we enforce the constraint. In the remaining columns, we implement this estimation for the data and parameters from alternative specifications in Table 2. 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 average implied value for η is 0.71.

Table J.1: The importance of technology sharing for production (η estimation)

	Dep var: Δy (1840–1810) (industrial output growth)			
	(1)	(2)	(3)	(4)
	Baseline	Patent counts	q_0^{UK} 1790	Endogenous
Panel A: OLS				
$Q_{k,1840}$	0.724	1.135	0.735	0.740
	(0.441)	(0.717)	(0.499)	(0.441)
$Q'_{k,1840}$	0.456	0.938	0.593	0.524
	(1.401)	(1.775)	(1.509)	(1.349)
Model-implied controls	✓	✓	✓	✓
N (industries)	14	14	14	14
R^2	0.62	0.59	0.65	0.64
Panel B: constrained, coefficients sum to one				
$Q_{k,1840}$	0.693	0.803	0.651	0.688
	(0.413)	(0.579)	(0.540)	(0.397)
$Q'_{k,1840}$	0.307	0.197	0.349	0.312
	(0.413)	(0.579)	(0.540)	(0.397)
Model-implied controls	✓	✓	✓	✓
N (industries)	14	14	14	14
F-stat	2.58	2.50	4.96	2.95
implied η	0.693	0.803	0.651	0.688

Estimates of the long-difference equation 19. The outcome is the growth of French industrial production between the 1810 decade (1803–12) and the 1840 decade (1835–44). Q and Q' are the French and British knowledge stocks. The model-implied controls are $Q_{k,1810}$, $Q'_{k,1810}$, $\Phi_{k,1840}$, and $\Phi_{k,1810}$, and the constant, with coefficients not reported. The columns correspond to the specifications of Table 2, from which the parameter and fixed-effect estimates entering the Φ terms are taken. *Baseline* (column 1): quality-adjusted patents, counting British patents from 1761 and French patents from 1791. *Patent counts* (column 2): patent counts in place of quality-adjusted patents. q_0^{UK} *1790* (column 3): only British patents filed after 1790 enter the British knowledge stock. *Endogenous* (column 4): endogenous growth model with κ constrained to zero. Robust standard errors in parentheses.