

Rosenberger, Lukas; Hanlon, W. Walker; Hallmann, Carl

Working Paper

Innovation Networks in the Industrial Revolution

CESifo Working Paper, No. 11299

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Rosenberger, Lukas; Hanlon, W. Walker; Hallmann, Carl (2024) : Innovation Networks in the Industrial Revolution, CESifo Working Paper, No. 11299, CESifo GmbH, Munich

This Version is available at:

<https://hdl.handle.net/10419/305541>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Innovation Networks in the Industrial Revolution

Lukas Rosenberger, W. Walker Hanlon, Carl Hallmann

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Innovation Networks in the Industrial Revolution

Abstract

How did Britain sustain faster rates of economic growth than comparable European countries, such as France, during the Industrial Revolution? We argue that Britain possessed an important but underappreciated innovation advantage: British inventors worked in technologies that were more central within the innovation network. We offer a new approach for measuring the innovation network using patent data from Britain and France in the late-18th and early-19th century. We show that the network influenced innovation outcomes and demonstrate that British inventors worked in more central technologies within the innovation network than French inventors. Drawing on recently developed theoretical tools, and using a novel estimation strategy, we quantify the implications for technology growth rates in Britain compared to France. Our results indicate that the shape of the innovation network, and the location of British inventors within it, explains an important share of the more rapid technological change and industrial growth in Britain during the Industrial Revolution.

Keywords: industrial revolution, innovation network, patents, economic growth.

Lukas Rosenberger
LMU Munich / Germany
lukas.rosenberger@econ.lmu.edu

W. Walker Hanlon
Northwestern University, Evanston / IL / USA
whanlon@northwestern.edu

Carl Hallmann
Jane Street
carlhallmann2022@u.northwestern.edu

August 14, 2024

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

1 Introduction

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

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

In this paper, we harness recent advances in economic growth theory and develop new approaches to using historical data in order to improve our understanding of this critical period in economic history. We begin by proposing a novel hypothesis for Britain’s advantages during the Industrial Revolution period. Specifically, we argue that British inventors were working “at the right place” in the *innovation network*, i.e., in technologies that generated more beneficial innovation spillovers. Our idea builds on emerging literature in growth economics which finds that innovation in some technologies generates more spillover benefits than innovation in others (Acemoglu, Akcigit, and Kerr, 2016; Cai and Li, 2019; Huang and Zenou, 2020; Liu and Ma, 2023), so that how a country’s research efforts are allocated within the innovation network can substantially impact the overall rate of economic growth.

Translating these ideas into the context of the Industrial Revolution, we ask: did Britain possess a growth advantage because British inventors were more focused on technologies, such as steam engines, machine tools, or metallurgy, that generated stronger spillover benefits for other technologies and were therefore more central in the innovation network? In contrast, could it have been the case that Continental economies like France experienced slower technological progress, in part, because their inventors focused on developing technologies, such as apparel, glass, or papermaking, which were more peripheral in the

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

innovation network?²

To structure our analysis, we begin with a growth model, from Liu and Ma (2023), that incorporates an innovation network. In this network, each node is a technology type, while each edge reflects the extent to which innovations in one technology type increase the chances of further innovation in another. This model provides a framework for thinking about how the distribution of researchers across technology types affects the growth rate in the economy. It also generates specific expressions that, given the matrix of connections across technology types, allow us to quantify how different allocations of researchers across those types affects growth.

Examining these issues during the Industrial Revolution requires the development of a new empirical approach. Like previous studies, we study innovation using patent data: for Britain from 1700 to 1849, and for France from 1791-1844.³ 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.⁴ However, key pieces of data, such as patent citations and R&D expenditures, are missing in our setting.

A key challenge in our setting is measuring spillovers across technology types. The innovation literature typically uses patent citations, but these are not available in our historical setting. Instead, 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 who patent in j to subsequently patent in i . Since our approach is new, we validate it using modern data. Specifically, using U.S. patents from 1970-2014, we construct innovation networks using our approach as well as the citation-based approach used in modern studies. Comparing these networks shows that the two approaches generate networks that are extremely similar. This suggests that our method does a good job of recovering the underlying innovation network. Developing and validating this new approach, which allows the study of innovation networks much further back in history than currently possible, is one primary contribution of our study.

Using our approach, we document innovation networks in Britain and France that feature a dense central core of closely related—and mainly mechanical—technologies. One

²Hallmann, Rosenberger, and Yavuz (2021) quantify how technological leadership in invention of Britain vs. France varied across technologies, with Britain leading, besides others, in steam engines and textile technologies, and France leading, besides others, in papermaking and shoemaking. Mokyr (1990, Chapter 5) provides a historical overview on British technological lead or lag in invention relative to Continental Europe.

³Both 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.

⁴Of course, not every useful invention was patented, as (Moser, 2012) has shown. To account for this, we also generate results based on data from non-patented inventions exhibited in the 1851 Crystal Palace exhibition. These results confirm the patterns that we obtain with the patent data.

important question about our estimated networks is, do they reflect fundamental features of the underlying technologies or simply reflect the local innovation environment in each country? To answer this question, we begin by constructing a crosswalk between the two very different technology categorization systems used in the two countries, based on a set of patents that we identify that were patented in both locations. Once we have both British and French patents expressed in common technology categories, we compare the innovation networks derived from these two sets of patents. We find that the networks share a common underlying structure. Thus, our innovation matrices do not just reflect the local economic environment; instead, a significant part of each reflects an underlying ‘global’ network of spillovers across technology categories.

Next, we establish that the shape of the technology spillover network matters for innovation outcomes. As a first step, we follow existing work on modern patent data by analyzing how patenting rates vary across technology categories depending on the lagged knowledge stock in other categories, weighted by the strength of connections through the innovation matrix. Consistent with the theory, and the results in previous studies of modern data, we find a significant positive associations of patenting with the lagged network weighted knowledge stock, shrinking toward zero as lags increase. However, the lack of exogenous variation in the lagged knowledge stock means that this result could be due to common shocks that affect connected technology categories.

Thus, in the second step, we provide evidence based on a source of quasi-exogenous variation in the timing of increases in the knowledge stock at some nodes of the innovation network. Specifically, we use the unexpected arrival of “macroinventions.” These are inventions which Mokyr (1990) describes as “a radical new idea, without a clear precedent, emerges more or less ab nihilo.” We offer three different approaches to identifying macroinventions, and then examine whether the arrival of a new macroinvention in one technology category leads to a subsequent increase in patenting in downstream technology categories within the innovation network. Here, the identifying assumption is *not* that the development of a particular type of macroinvention was random, but that the timing of the arrival of a macroinvention of a particular technology type was unpredictable within the two-decade time frame of analysis. Using pooled difference-in-difference and event study analyses for a time frame of ten years before and after the arrival of each macroinvention, we show that the arrival of a macroinvention upstream in the innovation network from a technology category results in an increase in patenting within that category. In contrast, the arrival of a macroinvention downstream from a technology category is not associated with an increase in patenting in that technology category. This second result provides a useful placebo check on our analysis. Providing better-identified evidence on the causal impact of the innovation network on technology development than what is currently available is the second main contribution of our study.

Next, we look at whether there are notable differences in the allocation of British and

French inventors within the innovation network. We study this by comparing British inventors patenting in France to French inventors, and comparing French inventors patenting in Britain to the British inventors.⁵ We find that among French patents, patents by British-based inventors were significantly more central compared to the average patents by French domestic inventors, whereas among British patents, patents by French-based inventors were less central compared to the average patent by British domestic inventors. The pattern indicates that British inventors were more likely to work in central technology categories than French inventors. As more central nodes have stronger spillover connections to other technology categories, the more central locations occupied by British inventors are consistent with a greater “bang for the buck” of British innovation on the aggregate rate of technological progress.

In last part of our paper, we bring our data to the model. To do so, we need to back out the allocation of research effort across technology types. Previous studies use R&D expenditures to measure the allocation of research effort, but these are not available in our setting. To overcome this, we develop a new estimation strategy that allows us to back out the allocation of research effort, as well as other key model parameters, from observed research outputs (patents). Using these estimated allocations, and the structure of the model, we can quantify the growth impact of differences in the allocation of British and French researchers across different technology categories. Importantly, these predicted growth differences are driven only by differences in the *allocation* of researchers across technology types, and abstract from any differences in the overall level of research effort in the two countries, as well as any other factors that differentially affect growth rates outside of the innovation network. Developing this new estimation strategy, and using it to quantify the impact of differences in the allocation of researchers across technology types is the third main contribution of our paper.

Existing estimates for Britain suggest that industrial production—our focal aspect of growth since the inventions that we study were most relevant there—grew by between 3.1 and 3.5% during the first half of the nineteenth century (Broadberry, Campbell, Klein, Overton, and van Leeuwen, 2015). In France, estimates indicate growth rates of between 1.9 and 2.5% in the same period (Crouzet, 1996; Asselain, 2007).⁶ Taking the midpoint of this range, our estimates suggest that differences in the allocation of researchers explain around one-half of the observed growth difference between the two countries.

Finally, our method allows us to identify which technologies made the largest contri-

⁵We have also attempted to study whether British vs. French inventors were more central within the innovation network of a third country, using U.S. patent data. Unfortunately, this analysis is not possible because U.S. patents only become available starting in 1836 (earlier patent information was lost due to a fire) and there are too few British and French inventors patenting in the U.S. in the two decades after that to draw any clear conclusions on their relative centrality within the U.S. network.

⁶How different the British and French growth rates truly were during the period we are interested in has been a long-term subject of debate, with important contributions by Cameron (1958) and O’Brien and Keyder (2011).

bution to the overall growth difference. We find that by far the most important technology category for Britain’s growth advantage was steam engine technologies, followed by technologies related to fuels (such as coal and coke), shipbuilding, water and fluids (such as pumps), and metallurgy. These results confirm, quantitatively, the historical narrative emphasizing the importance of these industries for the British economy during the Industrial Revolution. In contrast, French research effort was relatively more concentrated in technologies related to consumer products, such as flour milling, coffee, boots and shoes, and lamps and lighting, which provided fewer growth advantages.

In sum, we show that Britain benefited from an advantageous allocation of research effort across technology types during the Industrial Revolution, and that this difference meaningfully contributed to Britain’s more rapid industrialization. As in existing work using similar methods such as Liu and Ma (2023), our analysis takes as given the differences in the allocation of inventors across technologies. Thus, our mechanism complements explanations for the British advantage during the Industrial Revolution, in particular those that can explain why British inventors were more likely than the French to work on technologies that happened to be more central within the innovation network, in particular mechanical technologies. For example, it could be that Britain’s practical Enlightenment tradition and well-developed apprenticeship system (Mokyr, 2009; Kelly et al., 2014) contributed to the British inventors’ greater ability for working on mechanical technologies, or that high wages and access to cheap coal steered British inventors to focus on labor-saving mechanical devices (Allen, 2009a).⁷ Thus, the contribution of our paper lies in demonstrating *that* Britain was at the right place in the innovation network at the right time, and that this provided Britain with a meaningful growth advantage, while relying on existing work to explain *why* it was there but France was not.

In addition to improving our understanding of one of the most important questions in economic history, our study also contributes to work by growth economists on the importance of innovation networks. Relative to studies in this area (cited above), we offer four main contributions. First, we offer new ways of using data that can help researchers study innovation networks further back in history, when standard tools such as systematic patent citations and R&D spending are unavailable. This opens up the possibility of studying the influence of innovation networks in different contexts or over longer periods. Second, our analysis of macroinventions provides additional, more causal, evidence that innovation networks matter for technology development. Third, we develop a new estimation approach that can be used to back out the allocations of research effort when information on R&D expenditures are unavailable or unreliable. Fourth, our application

⁷A stable institutional environment and well-developed patent system may have contributed in shifting inventors from technologies that can be protected by secrecy toward technologies as mechanical devices that are easily reverse engineered and thus profit the most from patents (Moser, 2005). However, as both Britain and France had strong patent protection, it is unclear how this mechanism could explain the *differential* focus of British vs. French inventors on mechanical devices.

demonstrates empirically the value of recent theoretical advances integrating innovation networks into economic growth models.

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, Alessandro, Tartari, and Tranchero (2021), Nuvolari, Tortorici, and Vasta (2023), Hallmann et al. (2021), and Hanlon (2023). Relative to this extensive literature, we are the first to study the role of innovation networks in influencing inventive activity during the Industrial Revolution and the first to quantify the contribution of one specific mechanism to Britain’s growth advantage.

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, while Section 7 provides evidence that the structure of the network has a causal effect on innovation rates. Section 8 shows that British inventors tended to operate in more central nodes of the innovation network. Finally, Section 9 uses the structure of the model to quantify the implications of these differences for growth.

2 Historical context

The classical Industrial Revolution is typically thought of as starting around the 1770s and lasting up to the middle of the nineteenth century. Led by Britain, this was a period in which Western European countries experienced industrialization and the initiation of sustained modern economic growth. There is broad agreement that technological progress played a central role in this process.⁸

From historical evidence, we know that Britain was a leader in many of the characteristic technologies of the first Industrial Revolution: steam engines, spinning machines, metalworking, railroads, etc. One reflection of this lead was the effort that Continental countries, particularly France, (as well as Americans) put into obtaining British technologies in these sectors (Harris, 2017). Another reflection was the leading position that Britain occupied in many of the industries reliant on these technologies. Broadberry (1994), for example, states that “Few would dissent from the view that in terms of technology, Britain was the leading manufacturing nation during the first half of the nineteenth century.”

However, we also know that Continental inventors were also innovative, particularly

⁸See, e.g., Landes (1969), Mokyr (1990), Mokyr (1999), Allen (2009a).

those in France, the natural comparison country for Britain during this period.⁹ “There is,” wrote Harris (2017) (p. 560), “no suggestion that Frenchmen were inherently less inventive than Britons.” Instead, the most notable difference between the two countries is in the types of technology in which invention took place. Lagging behind in mechanical technologies, French inventors instead made important advances in papermaking, glassmaking, fine ceramics, and chemical processes (bleaching, soda, dyes, etc.).¹⁰ We are interested in how these differences in the technologies in which inventors in the two nations were working may have influenced their growth prospects.

The different distributions of inventive output were the result of numerous underlying factors. Britain’s success in mechanical technologies, for example, has been attributed to several sources. Some credit the country’s rich endowments of coal and metallic ores with providing a “focusing device” (Harris, 2017). Others note the superiority of the British apprenticeship system for transmitting mechanical knowledge (Kelly et al., 2014; Mokyr, Sarid, and Van Der Beek, 2022; Kelly et al., 2023). Yet others argue that high wages encouraged the development and application of labor-saving mechanical inventions (Allen, 2009a,b; Voth et al., 2022). Historical patterns of industry distribution must have had some influence; Britain’s established wool and cotton textile industries must have contributed to ongoing British invention in those sectors, in contrast to silk, where the French industry, and French inventors, were relatively more important. In chemical technologies, where French inventors made major contributions, France’s scientific predominance surely played a vital role.¹¹

Of course, these factors not only influenced the allocation of researchers across different technology types, they may have also had direct impacts on growth. The theoretical tools that we employ allow us to abstract from those direct growth effects, in order to isolate the impact of the allocation of researchers on growth operating through spillovers in the innovation network. Thus, we seek to isolate and quantify the contribution of the mechanism we emphasize, while not denying that many other mechanisms may have also been influencing relative growth in the two countries over the period of interest.

In fact, there are numerous factors that influenced relative growth rates in Britain and France during the period 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.¹² While temporarily retarding growth, the

⁹Leading 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.

¹⁰See, e.g., Harris (2017), p. 540-543, which discusses all of these important French innovations with the notable exception of papermaking. Also p. 560.

¹¹Harris (2017), p. 540.

¹²Most actual battles were fought outside of France, and a portion of the funding for them was extracted from conquered territory.

disruptions also ushered in important institutional changes, such as the adoption of a patent system and the introduction of the Code Napoleon, which have been shown to have contributed to subsequent economic growth (Acemoglu, Cantoni, Johnson, and Robinson, 2011). The Napoleonic Blockade also cut Britain off from many of its closest trade partners during this period, while providing an umbrella of protection that allowed some French industries to modernize (Juhász, 2018).

The net result of these and other factors, as well as the mechanism that we emphasize, was more rapid growth in industrial production in Britain than France. The best available estimates suggest that industrial production in Britain grew by between 3.1 and 3.5% during the first half of the nineteenth century (Broadberry et al., 2015), while in France the rate was probably between 1.9 and 2.5% (Crouzet, 1996; Asselain, 2007). Calculating the growth rates over the exact same time period, the British industry grew 0.6 to 1.2 percent per year faster than the French industry during 1815–1860s.¹³ Within this tumultuous historical period, our aim is to isolate and analyze the contribution of one potentially important but heretofore unstudied mechanism to the observed differences in relative industrial growth in the two countries.

3 Theory: Growth with Innovation Networks

This section presents a theory of growth with innovation networks. Our aim is to study how different allocations of research effort across technology types affect the growth rate of an economy through the innovation network, taking the allocation of research effort as given. Our model is based on recent work by Liu and Ma (2023), who introduce a matrix of spillovers across technology sectors into a continuous-time closed-economy endogenous growth framework.

3.1 Preferences and Production

The model features a representative consumer with utility at time t that is a function of discounted log consumption c_s in period t and all future periods:

$$V_t = \int_t^{\infty} e^{-\rho(s-t)} \ln c_s ds \quad .$$

¹³Summarizing the work by French economic historians, Asselain (2007) reports three best estimates for French industrial growth during the Industrial Revolution: 2.5% during 1815–1850 (Crouzet) and between 1.9% (Lévy-Leboyer) and 2.5% (Toutain) during 1820/24–1860/64 per year. For Britain, we calculated the corresponding industry growth rates using the latest data by Broadberry et al. (2015) as 3.5% during 1815–1850 and 3.1% during 1820/24–1860/64 per year.

Consumption is a Cobb-Douglas aggregation of goods from K different sectors,

$$c_s = \sum_{i=1}^K c_{it}^{\beta_i} \quad ,$$

where the β_i parameters give the consumption shares for each sector i and $\sum_i \beta_i = 1$.

Within each sector i , there is a continuum of varieties of intermediate products, denoted by v , which can be supplied in a countably infinite set of quality levels. The highest quality level available for any variety is given by $q_{it}(v)$. Only the highest quality version of each variety is used in the production process. We denote the quantity of a variety of quality q produced as $x_{it}(v|q)$ and total production (and consumption) of goods from sector i is:

$$\ln c_{it} = \int_0^1 \ln (q_{it}(v)x_{it}(v|q)) \, dv$$

Given some available quality level, production in a sector depends only on the number of workers allocated to that sector: $x_{it}(v|q) = l_{it}(v)$ where l_{it} is the quantity of labor employed in sector i .

3.2 Innovation

Following Liu and Ma (2023), we define the knowledge stock available in a sector i at time t as q_{it} , where $\ln q_{it} = \int_0^1 \ln q_{it}(v) \, dv$. These knowledge stocks are the state variables in the model. The knowledge stock in a sector improves through the efforts of researchers, r_{it} , working on developing new technologies in that sector at time t . The innovation production function is given by:

$$n_{it} = r_{it} \eta_i \chi_{it} \quad \text{where} \quad \chi_{it} = \prod_{j=1}^K q_{jt}^{\omega_{ij}} \quad . \quad (1)$$

In this expression, n_{it} is the set of new ideas in sector i generated in time t , η_i is a parameter that determines the productivity of research effort in sector i , and χ_{it} reflects the impact of spillovers across the innovation network that improve the chances of generating new innovations in sector i . These spillovers depend on the stock of knowledge in every other sector and a matrix of ω_{ij} parameters, the key parameters in our study, that determine the extent to which existing ideas in sector j increase the chances of producing new ideas in sector i . In order to obtain balanced growth across sectors, we need that $\sum_j \omega_{ij} = 1$ for all i . We denote the $K \times K$ matrix of these parameters as Ω , which we refer to as the *innovation network*.

New ideas translate into incremental quality improvements according to:

$$\dot{q}_{it}/q_{it} = \lambda \ln(n_{it}/q_{it}) \quad (2)$$

where the inclusion of q_{it} in the denominator on the right-hand side implies that improving quality becomes more difficult as the quality level rises. This formulation is intuitive in that it reflects the idea that improvements become more difficult once the “low-hanging fruit” has been harvested. It also plays an important functional role in the model, because it means that the continually increasing stock of existing knowledge, which generates a corresponding increase in useful knowledge spillovers, does not generate explosive growth.

3.3 Resource constraints

The number of production workers in the model is fixed at \bar{l} and the number of researchers at \bar{r} . Thus, $\sum_i l_{it} = \bar{l}$ and $\sum_i r_{it} = \bar{r}$. These assumptions abstract from the potentially important possibility that changes in the productivity of research activities may cause more workers to shift into research. This is a desirable feature for our purposes, since it emphasizes that the growth differences predicted by the theory are not due to differences in overall research effort in the two countries. Instead, the model will help us isolate the growth impact of differences in the allocation of research effort across technologies.

3.4 Key results

The model provides several results that we will use in our empirical analysis. The first result is related to how the innovation network determines the relationship between the current stock of knowledge in one sector and the rate of innovation in other sectors. We can derive this relationship from Equations 1 and 2. We obtain:

$$\ln n_{it} = \ln \eta_i + \ln r_{it} + \lambda \sum_{j=1}^K \omega_{ij} \left(\int_0^\infty e^{-\lambda s} \ln n_{j,t-s} ds \right) \quad (3)$$

Later, we will use this expression to structure our investigation of whether our estimated innovation network matters for innovation outcomes as well as to estimate the allocation of research effort.

A second key result relates the allocation of research effort across sectors to the growth rate:

Proposition: On a balanced growth path with researcher allocation vector \mathbf{r} , the aggregate stock of knowledge and consumption in each sector grows at the rate $g(\mathbf{r}) = \lambda(\mathbf{a}' \times \ln \mathbf{r}) + \lambda(\mathbf{a}' \times \ln \boldsymbol{\eta})$, where λ is a parameter, $\boldsymbol{\eta}$ is a vector of parameters, and \mathbf{a} is the dominant left-eigenvector of the innovation network Ω with an associated eigenvalue of one and satisfying $\sum_{i=1}^k a_i = 1$.

Proof: This proposition is a straightforward restatement of Lemma 2 in Liu and Ma (2023), so readers are referred there for the proof.

The difference in growth rates between the UK and France as a function of the vector of research effort allocations in the two economies is then given by,

$$g(\mathbf{r}^{UK}) - g(\mathbf{r}^{FR}) = \lambda \mathbf{a}' \times (\ln \mathbf{r}^{UK} - \ln \mathbf{r}^{FR}) \quad .$$

This expression can be used to quantify the differences in growth implied by different allocations of research effort in the two economies. Note, however, that growth differences may be due to either (i) differences in the overall amount of research effort in the two economies or (ii) differences in how that effort is allocated across technology types. We are specifically interested in the second of these, i.e., in how differences in the allocation across technology types interacts with the structure of the innovation network to influence growth. To separate these two effects, we decompose the $\ln \mathbf{r}^{UK} - \ln \mathbf{r}^{FR}$ term into two parts. Let,

$$\ln \mathbf{r}^{UK} - \ln \mathbf{r}^{FR} = \boldsymbol{\zeta} + \theta$$

where $\boldsymbol{\zeta}$ is a mean-zero vector reflecting differences in the log allocations of research effort in the two economies and θ is a constant term that captures overall differences in log research effort. The impact of differences in the allocation of log research effort on growth in the two economies can now be expressed as:

$$g(\mathbf{r}^{UK}) - g(\mathbf{r}^{FR}) = \lambda \mathbf{a}' \times \boldsymbol{\zeta}. \quad (4)$$

Equation 4 allows us to quantify the impact of differences in the allocation of log research effort on growth in the two economies given estimates of λ and $\boldsymbol{\zeta}$. In Section 9 we describe how we estimate these values.

As a point of reference, it is also useful to identify the growth-maximizing allocation of researchers across sectors.

Corollary 2: The allocation of researchers that maximizes the rate of technology growth, \mathbf{r}^* , solves $\argmax_{\mathbf{r}} \quad \mathbf{a}' \ln \mathbf{r}$ subject to $\mathbf{r} \geq 0$ and $\mathbf{1}' \mathbf{r} / \bar{r} = 1$. The solution to this problem is the vector \mathbf{a} .

This result tells us that the growth-maximizing allocation of researchers will involve a greater share of research effort put toward developing technologies that are more central in the innovation network.¹⁴

3.5 Discussion and extensions

In the theoretical setup above, we treat the innovation matrix Ω as fixed and exogenous. However, over long periods of time this matrix is likely to be evolving. Some of this

¹⁴This is not the same as the welfare-maximizing allocation, since the welfare-maximizing social planner will be willing to sacrifice some future growth in order to increase current consumption because future consumption is discounted.

evolution may be driven by factors outside of the model, but it is also possible that the network evolves as an endogenous function of the knowledge stock. In an extension of their model, Liu and Ma (2023) show that when the innovation matrix is modeled as an endogenous function of the knowledge stocks, their results continue to hold as a first-order local approximation around the balanced growth path. What this means for us is that even if the matrix is evolving endogenously, our results will still be valid as a first-order approximation. In Appendix D.3 we examine the evolution of the innovation matrix empirically and find that, while it is evolving over time, it was doing so quite slowly. Given this, it is reasonable to treat the matrix for main period of our analysis as fixed.

Another important simplification in the model above is that it does not incorporate the possibility that countries may receive innovation spillovers from other countries. However, it is relatively easy to extend the model to incorporate two economies, allowing new ideas produced in one to generate innovation spillovers for the other.¹⁵ To incorporate international spillovers, we modify the innovation production function so that innovation spillovers are a Cobb-Douglas aggregate of the domestic knowledge stock, q_{jt} and the foreign knowledge stock q'_{jt} with the relative weights of these two components determined by the parameter ι :

$$n_{it} = r_{it} \eta_i \chi_{it} \quad \text{where} \quad \chi_{it} = \prod_{j=1}^K \left(q_{jt}^{\iota} * q_{jt}^{\iota(1-\iota)} \right)^{\omega_{ij}}$$

This yields the following analog of Eq. 3:

$$\ln n_{it} = \ln \eta_i + \ln r_{it} + \iota \lambda \sum_{j=1}^K \omega_{ij} \left(\int_0^\infty e^{-\lambda s} \ln n_{j,t-s} ds \right) + (1 - \iota) \lambda \sum_{j=1}^K \omega_{ij} \left(\int_0^\infty e^{-\lambda s} \ln n'_{j,t-s} ds \right) \quad (5)$$

Later, we will use Eq. 5 to study how our quantitative results are affected by allowing for international spillovers through the innovation network.

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

¹⁵Note that this is not the only way that one economy can influence another. Economies might also adopt technologies from each other, which is different from sharing innovation spillovers, and they may also trade outputs. However, while both of those types of connections will affect welfare, our focus is on the rate of production of new technology. For that, innovation spillovers are the main dimension along which innovation in one economy will affect the other.

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.¹⁶ 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.¹⁷ 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.¹⁸

The French patent data used in our analysis were digitized by the French National Patent Institute (INPI). The data include patent number, patent title, inventor name, inventor occupation, inventor address, and additional details such as the type of patent and the patent term. French patents are divided into three main types: patents of invention, the standard format for new inventions; patents of importation for inventions originating abroad; and patents of improvement for modifications of existing patents. Our analysis focuses on the first two types, as they are the categories that represent truly new inventions.¹⁹ 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.²⁰

4.2 Linking inventor’s patents

Our approach requires that we identify all of the patents produced by an inventor. However, doing so is not trivial given that neither the British nor the French patent data include

¹⁶For 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).

¹⁷Most inventors were located in Britain, though a small number filed patents from a foreign address. In addition, 1,350 patent–inventor observations were “communicated from abroad.” In these cases, the location and name of the original inventor is unknown.

¹⁸See Hanlon (2023) for additional details about the data.

¹⁹Patents of improvement provided a cheap way to modify an inventor’s existing patent, but they did not extend the term of the original patent.

²⁰Another difference between the French and British patent systems is that inventors could choose in France to apply for patents of different lengths: 5, 10, or 15 years. Longer patent terms required higher fees. See Hallmann et al. (2021) for additional details about the data.

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.²¹

4.3 Concordant technology categories

Another critical challenge for our analysis is establishing concordance between the French and British technology categories. The difficulty is that the two nation’s patent offices employed structurally different systems of classifying patents into technology categories during our period. To build a mapping between these two different systems, we begin by identifying 1,148 patents that were filed in both countries—and thus classified by both patent offices into their national classification system. Using these bi-national patents, we construct a list of how British *subcategories* correspond to French *subcategories* (keywords). Aggregating these subcategories into distinct categories, we obtain 123 concordant technology categories into which we can reliably classify patents from both countries. Appendix A provides further details.

4.4 Production network connections

When analyzing the effect of the innovation network on invention, we would like to control for the influence of the production network operating through input–output (IO) connections between industries (Bloom, Schankerman, and Van Reenen, 2013). To measure the production network, we use the IO table for Britain in 1907 by Thomas (1984). The earliest detailed input–output matrix available for the British economy, this IO table gives us a matrix of upstream and downstream connections between 33 industries.²² The key challenge for creating production network controls, however, is how to map input–output industries to patenting technology categories.²³

²¹The 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.

²²The IO matrix by Thomas (1984) contains 41 sectors. We exclude the four service sectors (*Laundry, Public utility, Distributive services, Other services*), aggregate the four chemical industries into one because of difficulties in matching unique occupations (*Chemicals, Soap and candle, Oils and paint, Explosives*), and exclude the *Motor and Cycle* industry because it did not yet exist during our period. Horrel, Humphries, and Weale (1994) provide an input–output matrix for the British economy in 1841 which is much less detailed.

²³No 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

We introduce a novel approach based on inventors’ occupations for mapping industries and technology categories. A substantial fraction of the occupations reported by patenting inventors can be unambiguously associated with specific industries. Given the classification of patents into technology categories, we can use unambiguous occupations to construct a probabilistic mapping from technology categories to industries. Specifically, we construct a set of weights φ_{in} reflecting the ratio of patents in technology category i linked to industry n to the total number of patents in category i linked to any industry. Combined with the input–output matrix, the mapping allows us to construct controls reflecting how technology categories are connected upstream and downstream through the production network. Appendix B provides further details.

4.5 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.²⁴ 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.²⁵ 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.

4.6 Mapping the data to the model

Before proceeding with the analysis, we need to decide how to map our patent data to the model. In the model, n_{it} is the increment of new ideas generated in category i in time t . While we clearly want to relate this to the number of new patents generated in the category and period, we have some latitude in terms of the functional form for this relationship. It may seem natural to set $n_{it} = Patents_{it}$, as is done in some previous studies. However, this functional form is unappealing in our setting, because it means that the spillovers of category-year cells with zero patents, which are frequent in our data, will be undefined.

different industries.

²⁴This exhibition was followed by others, such as the Exposition Universelles in Paris (1855, 1867) or the American Centennial Exhibition in Philadelphia in 1876. We focus only on the Crystal Palace because it is the first international exhibition that occurred in close proximity to the period covered by our patent data.

²⁵The original dataset includes 30 categories, but we omit category 30, which is for fine arts exhibits. For some French exhibits, it also provides secondary or tertiary categorizations for French exhibits. But as most exhibits have only one category, we focus on each exhibit’s first category.

Since those spillovers appear on the right-hand side of Eq. 3 for all other technology categories, we would end up with undefined terms in every equation. That would not only create problems for operationalizing the model, it is also intuitively unappealing, because it seems more natural that spillovers from a category-year cell where no patent was produced should generate zero spillovers, rather than being undefined.

A more attractive functional form is $n_{it} = Patents_{it} + 1$. Intuitively, this amounts to assuming that there is some baseline level of unpatented and therefore unobserved knowledge production going on. Adding a baseline knowledge production equal to one is attractive in our setting because it means that categories where no patent was produced generate no spillovers. We adopt this functional form throughout the remainder of our analysis. However, in subsection 9.4 we discuss how it affects our quantitative results. Note that adopting this functional form does not represent a departure from the model, because the model makes no statement about how the patent data relates to the objects in the model. However, it does mean that any parameters estimated using the model must be interpreted in the context of this assumption.

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

5.2 Novel inventor-based method

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

To overcome this lack of data, we introduce a novel method that follows a similar intuition as the citation-based method but instead exploits information contained in the

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

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

$$\omega_{ij} = \frac{\sum_k P_{kij}}{\sum_k P_{ki}} \quad (6)$$

where P_{kij} is the number of patent sequence pairs by inventor k , with the first patent filed in technology category j and the next patent filed in i , and where P_{ki} is the total number of patents by inventor k in technology category i which pair with an earlier patent (which can be either in i or in another technology category g).²⁶

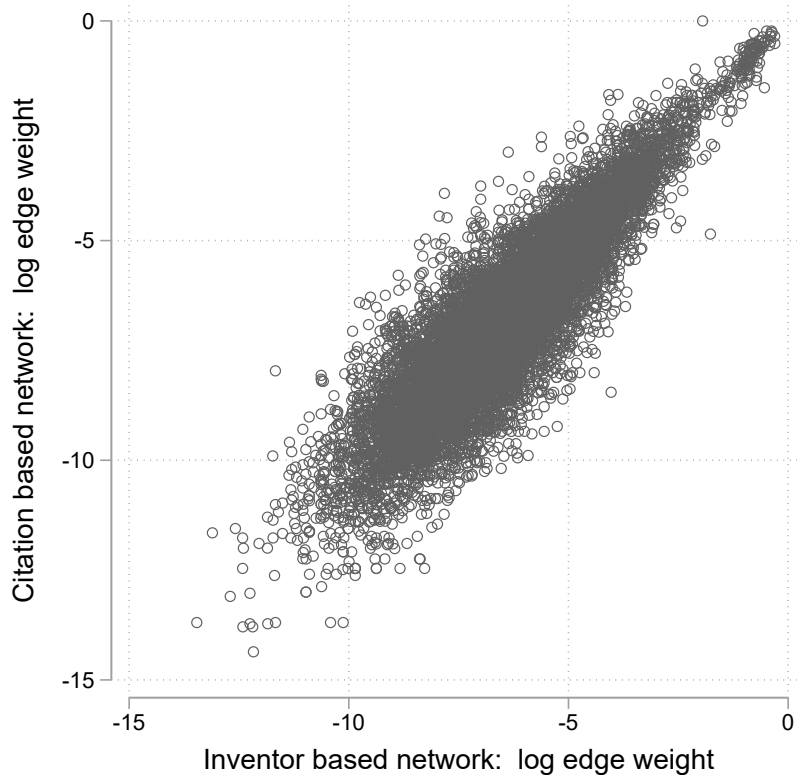
We can calculate ω_{ij} values using either British patents, French patents, or both. In our main analysis, we will rely on a matrix generated using both sets of patents after they have been mapped to our common set of concordant technology categories. This joint matrix can be constructed in several ways; either will require judgment about what relative weight should be attached to patents from each country. As patents in the two countries are the product of different patent systems and institutional environments, it is not obvious how to determine an objectively correct weighting. Thus, we choose the simple approach: giving each patent system equal weight. Specifically, we simply average the (row-normalized) edge values across Britain and France to obtain the joint matrix.

5.3 Validating the inventor-based method

Whether the inventor-based method provides a useful measure of the innovation network is ultimately an empirical question. To provide some confidence in our new approach, we use modern patent data to compare the innovation network generated by the inventor-based method to the innovation network generated by the standard citation-based method. We use data on U.S. patents from 1970–2014 from PatStat. As described in more detail in Appendix C, we generate a citation-based innovation matrix using a standard approach

²⁶One complication in our setting is that some British patents are categorized into multiple technology categories. To deal with this, we generalize P_{ki} and P_{kij} to be the weighted count of pairs of patents by inventor k , with weights corresponding to the inverse of the number of categories a patent is listed in. This is only one possible solution to the problem; another alternative is to focus on the modal technology category and throw a coin for cases without distinct mode (e.g. Jaffe, Trajtenberg, and Henderson, 1993).

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



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

taken from previous studies. Our inventor-based innovation matrix is obtained using the approach described above. Once we have the two matrices, we can compare either the edge values or the centrality of the nodes in the two matrices.

Figure 1 presents a scatterplot comparing the edges (ω_{ij} terms) obtained from the standard citation-based method against our inventor-based method. It is evident that the two approaches give very similar results. The corresponding regression of citation-based against inventor-based log edge weights estimates the slope as .994 (95% confidence interval [.984, 1.003]) with an associated R^2 of .79. Thus, our inventor-based approach provides a very close approximation to the network generated using the citation-based approach commonly used in existing studies.

6 Innovation networks in the Industrial Revolution

Having established the validity of our method, we now apply it to British and French patent data in the Industrial Revolution. A first glimpse of the innovation network is shown in Figure 2, which provides a visualization of the innovation network based on the

joint matrix constructed from British and French patents. In the figure, each technology category is a node, the size of node reflects the number of patents filed in that category, and the location of the node is determined by the strength of connections between that node and every other node in the network as determined by a multidimensional scaling algorithm.

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

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

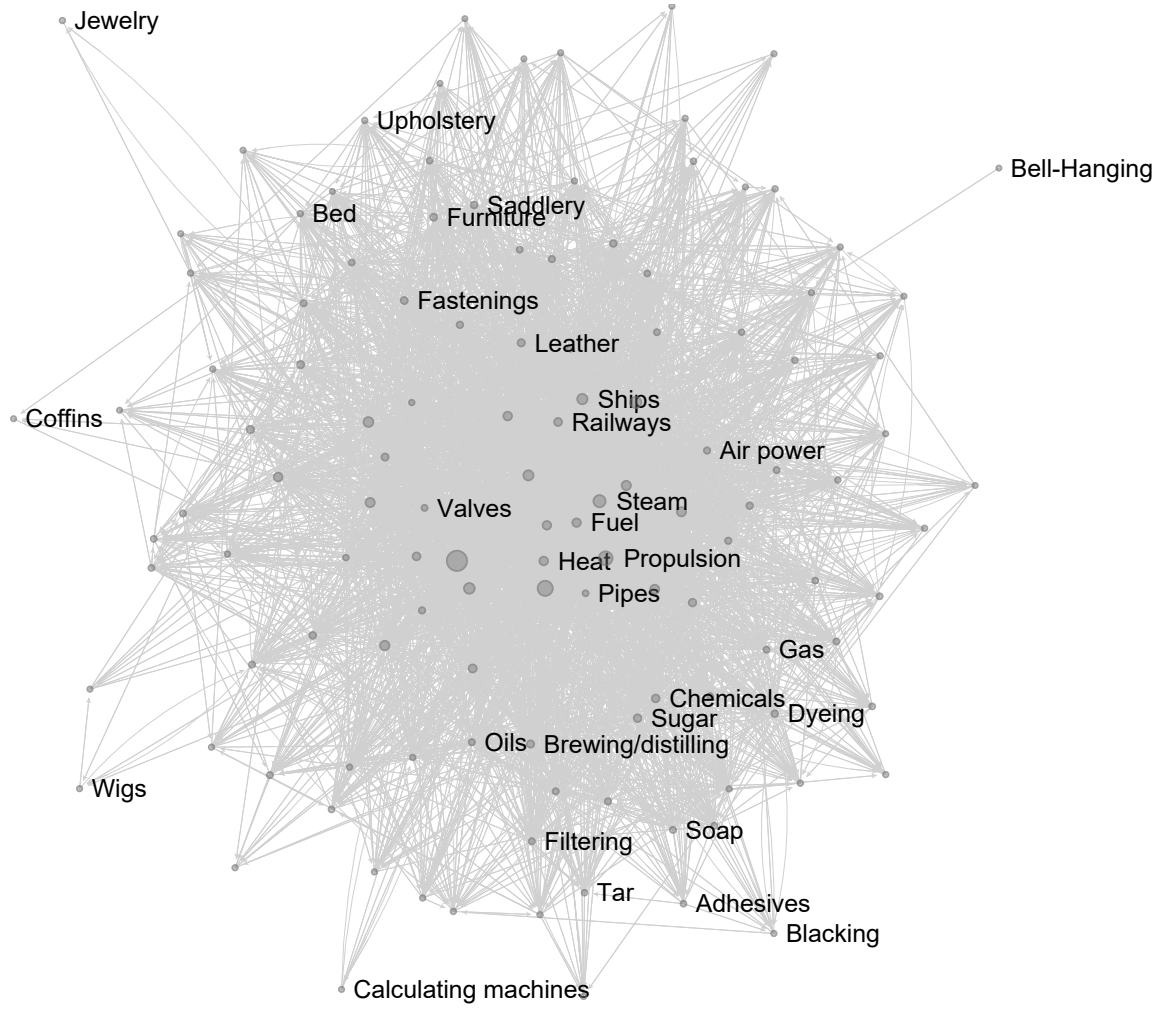
To assess the similarity of the British and French networks, we begin with two separate innovation networks, one constructed using only French patents and another constructed using only British patents, but both expressed in terms of the common technology categories. We then apply the following regression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

Figure 2: The joint innovation network visualized



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

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

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

Observations are network edges connecting nodes (technology categories) *i* and *j*. 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

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 weight, i.e., the ω_{ij} values, and include those nodes with zero measured connection. Note that we have a large number of nodes, over 15,000, relative to the number of patents in our data. As a result, many nodes may have zero measured connection 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 nodes where we observe some positive connection ($\omega_{ij} \geq 0$) in both the British and the French realization of the innovation network. This compares the strength of the connections at the intensive margin.

Node centrality An alternative approach to assessing matrix similarity is to focus on the centrality of the network nodes, which provides a useful way to summarize the shape of the network. In the network literature, there exist different metrics for node centrality; which is most appropriate depends on the application. Our theoretical results highlights the importance of *eigenvector centrality* in determining outcomes.

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

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

To summarize, across all of these various approaches, we find strong evidence of similarity in the French and British innovation networks. This tells us that at least part of the observed innovation network is likely to be due to some underlying “global” factors common to both countries, as assumed in the theory.

7 Effect of the network on innovation

In this section, we examine the effect of the network on innovation. In the first step, we follow existing studies on modern innovation networks by running panel regressions using lagged values of the network-weighted knowledge stock based on Eq. 3. Identification in this approach relies on the assumption of no common shocks to connected technology categories, which can be difficult to accept. To address this concern, we introduce in the second step a novel approach that uses the unexpected arrival of important inventions in certain technology categories to isolate variation in knowledge stocks.

7.1 Effect of knowledge stocks on patenting

Equation 3 expresses the log number of new ideas in a particular technology category i and year t as a function of the log knowledge stock in other categories that generate spillovers for technology i through the innovation matrix. This expression has been used by existing studies, such as Liu and Ma (2023), to provide evidence that the innovation network has an impact on innovation outcomes. We operationalize this using the following regression specification,


$$\ln(n_{it}) = A_i + B_t + \sum_s \beta_s \sum_j \omega_{ij} \ln(n_{j,t-s}) + \epsilon_{it} \quad (7)$$

where A_i and B_t are, respectively, technology category and year fixed effects, and s is the number of lags of the knowledge stock included in the regression.

Figure 4 presents the estimated effect of the lagged British knowledge stock on British patenting, for lags from one to eight years. The network proximity weighted lagged knowledge stock is significantly and positively associated with patenting rates at shorter lags. The association decreases quickly over time, consistent with the exponential decay that we would expect given the model. As the finding is fairly similar to those obtained by

studies using modern data, it appears that our novel network measures are representing the innovation network well.²⁷ In Appendix E.1, we show that similar decay patterns are obtained for the upstream knowledge stock if we control for the lagged *downstream* knowledge stocks. Knowledge stocks in downstream categories, in contrast, have a negative effect on patenting, and coefficients are considerably smaller in magnitude.

Figure 4: The effect of the lagged knowledge stock on patenting rates



The figure presents estimated coefficients and 95% confidence intervals for PPML regressions based on Eq. 7 applied to all British patents by domestic inventors. The innovation matrix is based on all patents (British and French). Patents appearing in multiple (N) technology categories count as only a fraction ($1/N$) of a patent in each of category.

7.2 Effect of macroinventions on patenting

One important concern with the approach above is that there may be common shocks to connected technology categories, which would result both in greater knowledge stocks in some categories as well as higher rates of patenting in other connected technologies, but not as a result of spillovers through the innovation network.

To provide a stronger test of the role of innovation networks, we use the arrival of unexpected macroinventions in certain technology categories as a source of quasi-exogenous variation in knowledge stocks. Macroinventions are ideal for this exercise because (1) they represent substantial increments to existing knowledge and (2) they are thought to be largely unpredictable. Mokyr (1990), for example, described macroinventions as “inventions in which a radical new idea, without a clear precedent, emerges more or less *ab nihilo*.” According to Crafts (1995, p. 596), “Technological history suggests that seeking for socio-economic explanations of macroinventions is likely to be a fruitless pursuit.”

²⁷See, e.g., Liu and Ma (2023) Figure A.9.

The key identifying assumption is that the exact timing of arrival of macroinvention is unpredictable within the analysis window. The analysis does *not* assume that the technology category in which a macroinvention appeared was random. To illustrate the variation harnessed in our analysis, take the example of steam engines. After Thomas Newcomen introduced the atmospheric engine in 1712, there were consistent efforts to improve the efficiency of the design. Thus, it was likely that a major advance would occur in the area of steam engines at some point in time. However, it took until 1769 that James Watt invented the separate condenser. From the historical accounts, there is no apparent reason why that idea may not have occurred earlier—and it may well have occurred many years later if genius had not struck Mr. Watt.

In this part of the analysis, we focus only on British patents, where we have access to three different approaches to identifying macroinventions. Our first approach relies on a set of 406 patents that were the first patent in a particular technology subcategory.²⁸ This set of patents is likely to be unexpected since each patent opened up a new technology *subcategory*, but it is unclear how impactful they were. Our second approach is based on a list of 65 British macroinventions provided by Nuvolari et al. (2021).²⁹ Nuvolari et al. (2021) provide evidence that these patents were particularly impactful. As a third measure, we use the intersection of the two sets, which generates a small set of five patents that are likely to be both impactful and unexpected.³⁰ This may seem like a small set of experiments to work with but recall that we can examine the impact across all other technology categories for each event. We call these three alternative macroinvention definitions the “First list”, “Nuvolari list”, and “Intersection list.”

Empirical specification We structure the dataset as a stacked panel. We define ‘event’ e as a year t in which at least one macroinvention occurred in technology category j . For each event, we construct a sub-panel dataset with four five-year periods τ : Two periods before the event year ($t - 10$ to $t - 6$ and $t - 5$ to $t - 1$) and two after ($t + 1$ to $t + 5$ and years $t + 6$ to $t + 10$), excluding the year of event t itself. The cross-sectional dimension of the panel is the set of all technology categories $i \neq j$. Thus, the level of observation will be macroinvention–event e by period τ by technology category i cells. We then look at whether technology categories that receive more spillovers from the technology category

²⁸We exclude patents from this list before 1750 since they may appear to be the first patent in their subcategory only because our data began in 1700.

²⁹These are identified using a wide variety of sources, including contemporary citations to patents compiled by Bennett Woodcroft and the British Patent Office, biographies of famous inventors such as the Oxford Dictionary of National Biography, and modern histories of technology such as Bunch & Helleman’s *History of Science and Technology*. Nuvolari et al. (2021) define macroinventions as the top 0.5 percentile of patents in a composite citation score that is based on all of the sources they review.

³⁰The five macroinventions on the intersection list are Jedidiah Strutt’s stocking rib machine (patent 722), Alexander Cumming’s flush toilet (patent 1105), Joseph Bramah’s beer pump (patent 2196), George Stephenson’s half-lap joint for railway tracks (patent 4067), and Charles Macintosh’s rubberised waterproof cloth (patent 4804).

j where the macroinvention arrived exhibit relatively more patenting in the post-arrival periods.³¹

Our “stacked difference-in-difference” specification is,

$$\ln(n_{ie\tau}) = \beta^{UP} \omega_{ie}^{UP} \cdot post_{e\tau} + \beta^{DOWN} \omega_{ie}^{DOWN} \cdot post_{e\tau} + X_{ie\tau}\Gamma + \gamma_{ie} + \eta_{e\tau} + \epsilon_{ie\tau} \quad (8)$$

where $\ln(n_{ie\tau})$ is the log number of patents plus one in technology category i in time period τ of event e (though results are robust to using log patents), ω_{ie}^{UP} is the (British) innovation network edge from the macroinvention category to the focal category i , ω_{ie}^{DOWN} is the (British) innovation network edge from technology category i to the macroinvention category, $post_{e\tau}$ an indicator for the periods after the arrival of the macroinvention, $X_{ie\tau}$ a set of control variables (interacted with post indicator) defined later, γ_{ie} a set of technology-category-by-event fixed effects, $\eta_{e\tau}$ a set of period-by-event fixed effects, and $\epsilon_{ie\tau}$ an error term that may be correlated across time within a technology category.

The main coefficient of interest in this regression is β^{UP} , which tells us whether patenting in a technology category changed after the arrival of a macroinvention upstream in the innovation network. We expect to observe a positive and statistically significant β^{UP} coefficient. We are also interested in β^{DOWN} , which tells us how patenting in a technology category responded to the arrival of a macroinvention downstream in the innovation network. If our identification strategy is working, then we should not see a positive coefficient on β^{DOWN} . In fact, downstream macroinventions may actually reduce innovation if the macroinvention attracts inventors to redirect their effort to the technology where the macroinvention arrived.

This specification effectively averages coefficients from separate difference-in-difference regressions by stacking the panels to obtain common β^{UP} and β^{DOWN} coefficients. Given the distribution of the dependent variable, which is censored at zero and skewed, we estimate Eq. 8 using Pseudo-Poisson Maximum Likelihood (PPML) regressions (Correia, Guimarães, and Zylkin, 2019, 2020).

We might worry that the arrival of a macroinvention affects innovation in other technology categories through input-output connections, rather than through spillovers across the innovation network (Bloom et al., 2013). To deal with this concern, we include controls for upstream and downstream IO linkages to the macroinvention category.

Results Table 2 presents results for equation 8. Panel A presents results based on an innovation matrix calculated using only British patents (since this analysis focuses only on British patents as the outcome). Columns 1 and 2 present results using the first patent list. Columns 3 and 4 present similar results using the Nuvolari et al. list, while Columns 5 and 6 present results using the intersection of the two lists. In each set of results, we

³¹If there were multiple macroinventions in the year t , then we sum the proximity across all macroinventions and omit from the panel any category where a macroinvention occurred.

begin with our baseline specification and then add in controls for input-output connections. The results show that the arrival of a macroinvention upstream significantly increases patenting in connected downstream technology categories. In contrast, we find no evidence that innovation increased when a macroinvention arrived downstream of a technology category, and in fact that appears to be associated with a decrease in innovation. We also test for the equality of the upstream and downstream coefficients, at the bottom of the panel, and find that they are statistically differentiable for all specifications except those using the intersection list as the outcome variable, where our estimates are imprecise. The inclusion of controls for input-output connections does not substantially affect the results.

It is interesting to note that the size of the estimated effects in the Nuvolari list and the First list are substantially smaller than the effects estimated for the Intersection list, though the latter are much less precise. This makes sense given that the Intersection list includes patents that are both novel and important.

In Panel B, we limit the analysis to patents filed in Britain after 1791, the year when our French patent series begins. We do this for comparability with Panel C, where we address the potential concern that our results may be influenced by the fact that we are using British patents both as our outcome measure and to construct our innovation matrix connections (ω_{ie} terms). In Panel C, we replace connections based on a matrix produced using only British patents with connections from an innovation matrix based only on French patents. We continue to find clear results indicating that the arrival of a macroinvention upstream increased patenting in a technology category, while a macroinvention downstream did not.

In Appendix E.2 we also show that our results are robust to other potential concerns. For example, we show that our results are robust to dropping all other patents by inventors who generated macroinventions from the data, as well as to using $\log(Patents_{ij})$ instead of $\log(Patents_{ij} + 1)$ as the dependent variable.

Appendix E.2 also includes results from an event study specification. These show that prior to the arrival of a macroinvention there was no differential patenting trends in technology categories that had stronger upstream connections to the macroinvention category. After the arrival of the macroinvention, we observe an immediate increase in the number of patents in technology categories that had a stronger upstream connection to the macroinvention category. We also observe no evidence that having a stronger downstream connection to the macroinvention category was associated with an increase in patenting.

Table 2: Macroinventions regression results

	Dep var: Ln (British patents + 1)					
	First list		Nuvolari list		Intersection list	
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: All years, British innovation network</i>						
$\omega_{ie}^{UP} \times \text{post}$	0.440*** (0.167)	0.434** (0.170)	0.660** (0.300)	0.704** (0.305)	1.475 (1.323)	1.652 (1.274)
$\omega_{ie}^{DOWN} \times \text{post}$	-0.217 (0.142)	-0.216 (0.143)	-0.710** (0.295)	-0.626** (0.294)	-0.194 (0.584)	-0.109 (0.578)
IO controls		✓		✓		✓
N (Obs = category–event–period)	34296	34296	18392	18392	2084	2084
p upstream = downstream	0.0038	0.0044	0.0035	0.0042	0.26	0.23
<i>Panel B: 1791–1849, British innovation network</i>						
$\omega_{ie}^{UP} \times \text{post}$	0.413** (0.180)	0.396** (0.187)	0.803** (0.337)	0.844** (0.340)	3.280*** (1.168)	3.290*** (1.217)
$\omega_{ie}^{DOWN} \times \text{post}$	-0.220* (0.115)	-0.237** (0.115)	-1.143*** (0.282)	-1.081*** (0.281)	-0.164 (0.487)	-0.170 (0.486)
IO controls		✓		✓		✓
N	20524	20524	11524	11524	1396	1396
p upstream = downstream	0.0068	0.0071	0.000085	0.00010	0.0086	0.0099
<i>Panel C: 1791–1849, French innovation network</i>						
$\omega_{ie}^{France-UP} \times \text{post}$	0.295** (0.127)	0.293** (0.127)	0.353*** (0.134)	0.357*** (0.133)	1.743* (0.916)	1.655* (0.886)
$\omega_{ei}^{France-DOWN} \times \text{post}$	0.008 (0.071)	0.006 (0.073)	-0.176 (0.141)	-0.146 (0.137)	0.179 (0.377)	0.092 (0.373)
IO controls		✓		✓		✓
N	8008	8008	4036	4036	224	224
p upstream = downstream	0.054	0.055	0.0076	0.0094	0.12	0.11

PPML regressions, implemented using *ppmlhdfc* command. There are four periods per event, two before ($[t - 10, t - 6]$ and $[t - 5, t - 1]$) and two after ($[t + 1, t + 5]$ and $[t + 6, t + 10]$). p upstream = downstream is the p -value of a Wald test (χ^2) for equality of upstream vs. downstream coefficients. Standard errors clustered by technology category in parentheses. All regressions include category-event and event-period fixed effects.
 * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

8 Centrality of inventors by country

The previous section provides evidence that the shape of the innovation network matters for technological progress. In this section, we provide a reduced-form analysis of whether

there are systematic differences between Britain and France in terms of the distribution of patenting activity across technology categories, which could have implications for their rate of technology growth. In the next section, we will take the model more seriously, use it to estimate the differences in the allocation of research effort between the two countries, and then quantify the differential growth outcomes that those allocation differences imply.

To generate a fair comparison between the centrality of British and French inventors within the innovation network, we compare, in both countries, foreign inventors to domestic inventors using the domestic innovation network. If we find that foreign inventors were always patenting in more central categories, in Britain as in France, differences in centrality could be due to a foreign inventor selection effect. If we find that only inventors from one country are more central, however, we can rule out such selection effect. For example, using French patents and the French network, we estimate

$$Centrality_{pkt} = \beta_{UK} UK_k + \beta_{foreign} OtherForeign_k + \phi_t + \epsilon_{pkt} \quad (9)$$

where $Centrality_{pkt}$ is the eigenvector centrality of the technology category associated with patent p patented by inventor k in year t , UK_k is an indicator for whether inventor k reported a UK address when filing the patent in France, and $OtherForeign_k$ is an indicator for whether the inventor listed some other location outside of France as their address, for example in the USA, or the patent type is “of unspecified origin” (*communication* in British patents, *importation* in French patents).

Results In Columns 1-2 of Table 3, we look at the centrality of British and other foreign inventors, relative to French inventors, within the French patent system. Columns 1 shows that foreign inventors tended to patent in more central technology categories within the French innovation network, while Columns 2 shows that this was true for both British and foreign inventors, but less true for inventors from other countries. There are two potential explanations for these patterns. Either it may be that British inventors specifically worked in more central categories, or it may be that foreign inventors always systematically selected into more central technology categories. To tell the difference between these stories, we need to compare the results in Columns 1-2 to those in Columns 3-4. There, we see that foreign inventors (with the exception of those from the US) tended to patent in less central technology categories relative to British inventors within the British patent system. These results tell us that, relative to French inventors, British inventors tended to patent in more central technology categories within the innovation network (and US inventors even more so).

Additional results based on exhibition data To provide further support for these conclusions, we have also looked at the centrality of British exhibitors relative to French exhibitors within the innovation network based only on exhibition data. Presented in Appendix F, these results show that more central nodes within the exhibition-based innovation

Table 3: Centrality of foreign inventors within the British and French innovation network

	Dep var: Log eigenvector centrality			
	French patents		British patents	
	(1)	(2)	(3)	(4)
Foreign origin	0.095*** (0.023)		-0.070* (0.037)	
British inventor		0.090** (0.037)		
French inventor				-0.257* (0.147)
US inventor		0.411*** (0.103)		0.274** (0.109)
Other foreign inventor		0.055 (0.052)		-0.063 (0.175)
Year FE	✓	✓	✓	✓
<i>N</i> (Obs = inventor–patent)	13173	13173	8767	8767
<i>R</i> ²	0.010	0.010	0.018	0.019

Observations are inventor–patents in France (col 1, 2) or in Britain (col 3, 4) 1791–1843. Eigencentality is calculated as left eigenvector in line with the theory. Foreign origin patents in France include “importation patents” and patents with a foreign address. Foreign origin patents in Britain include “communicated patents” and patents with a foreign address. The sample sizes are slightly smaller than the total number of inventor–patent observations in each country because we cannot assign a few patents to common technology categories. Furthermore, French domestic patents from territories occupied under Napoleon (up to 1813) are excluded. OLS regressions, robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

network had a higher ratio of British to French exhibitors.

9 Quantification of growth effects

From the theoretical results, we know that the growth-maximizing allocation of researchers involves a larger share working in more central technology categories. The results in the previous section indicated that British inventors tended to work in technology types that were more central within the innovation network. However, these results alone do not imply that the observed more central allocation of British inventions translates into a higher growth rate. In order to take that next step, we need to use the theory in order to estimate the underlying allocation of research effort before we can quantify how differences in that allocation affect growth.

The starting point for our quantification exercise is Eq. 4, which tells us that we can calculate the growth difference between two economies given estimates of the λ parameter, the vector α' which can be obtained from the innovation network, and a vector ζ describing the relative allocation of research effort in the two economies across each technology category. In order to implement this approach, we therefore need estimates of λ and ζ .

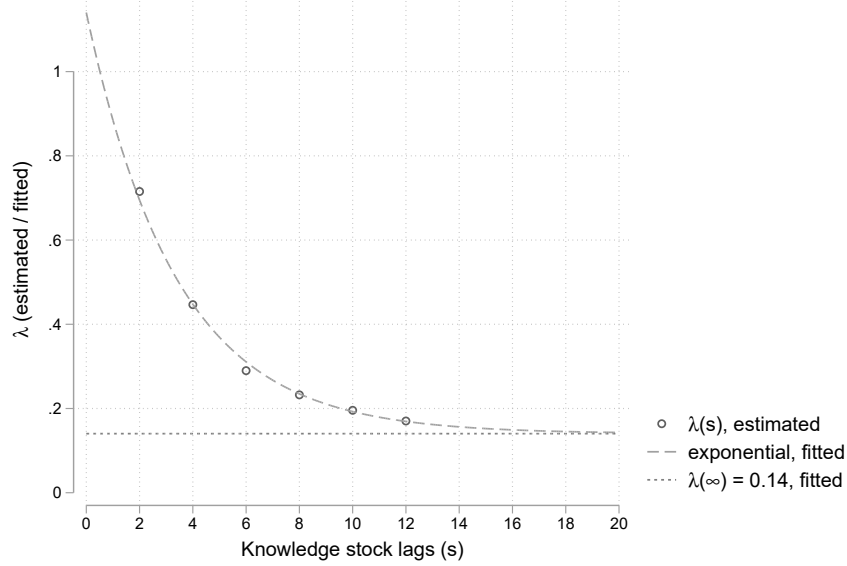
Previous studies have measured the distribution of research effort using information on R&D expenditures (Liu and Ma, 2023). However, such data are not available in the historical setting that we study. To deal with this, we develop a new method that uses the structure of the model to back out research effort allocations from the observed research outputs (patents) while accounting for the influence of the knowledge stock on researcher productivity. While this method allows us to overcome the lack of data on R&D expenditure in our setting, it may also be useful in others settings where R&D expenditure data are less reliable due to accounting practices, tax treatments, and other factors.

The starting point for our method is Eq. 3, which relates the number of patents produced in a country in a particular period (observed in the data) to the researchers working in that area as well as the available knowledge stock in that country. If we take the difference in this expression across the two countries that we study, we obtain

$$\begin{aligned} \ln n_{it}^{UK} - \ln n_{it}^{FR} &= \ln r_{it}^{UK} - \ln r_{it}^{FR} \\ &+ \lambda \left[\sum_{j=1}^K \omega_{ij} \left(\int_0^\infty e^{-\lambda s} \ln n_{j,t-s}^{UK} ds \right) - \sum_{j=1}^K \omega_{ij} \left(\int_0^\infty e^{-\lambda s} \ln n_{j,t-s}^{FR} ds \right) \right] \end{aligned} \quad (10)$$

We are interested in using this to estimate average differences in log research effort between the two countries across our study period. Thus, treating the difference in log research

Figure 5: Estimated, fitted, and asymptotic λ values



The estimated λ are obtained from different implementations of Eq. 11 that include up to s lags of the knowledge stocks. The exponential is fitted via NLS as $\lambda(s) = 0.14 + e(-0.295 \times s)$, indicating that λ converges to $\lambda_\infty = 0.14$ as the number of included lags s goes to infinity.

effort as constant over time, we can rewrite this as the following regression equation

$$\ln n_{it}^{UK} - \ln n_{it}^{FR} = \xi_i + \lambda \left[\sum_{s=1}^v e^{-\lambda s} \left(\sum_{j=1}^K \omega_{ij} \ln n_{j,t-s}^{UK} - \sum_{j=1}^K \omega_{ij} \ln n_{j,t-s}^{FR} \right) \right] + \epsilon_{it} \quad (11)$$

where the ξ_i fixed effects will reflect average (time invariant) differences in log research allocations.


Note that we need to estimate both the ξ_i fixed effects and the λ parameter. Because of the way that λ enters into this expression, we estimate Eq. 11 using Non-linear Least Squares.³²

We also need to choose the number of lags of the knowledge stock to include (v). The choice of lags turns out to be important for our estimate of λ , though it hardly effects the researcher allocations that we estimate.³³ Figure 5 describes how our estimate of λ changes as we increase the number of lags of the knowledge stock that we include. As we increase the number of lags, the λ estimate appears to be asymptotically approaching a stable value. Fitting an exponential line to the estimated λ values in Figure 5 suggests that as the number of lags included approaches infinity the value of λ is approaching 0.14. We will use this value in our main quantification results. This value is fairly similar to the λ

³²Note that we do not apply PPML as we did in Section 7 because here our dependent variable follows an approximately symmetric distribution around zero.

³³Specifically, the correlation of the researcher allocation vectors estimated using lags ranging from two to twelve are always above 0.99.

Figure 6: Comparing estimated researcher allocations and observed research output



The figure plots the log difference in research output, $\ln n_i^{UK} - \ln n_i^{FR}$, against the estimated log difference in research effort, $\xi_i = \ln r_i^{UK} - \ln r_i^{FR}$ for each technology category i and a fitted regression line. ξ_i is estimated by implementing Eq. 11 using Non-Linear Least Squares. The results in this figure are based on estimates using eight lags of the knowledge stock, but the number of lags has very little impact on the estimated researcher allocations.

parameters obtained in modern studies. For example, Acemoglu, Akcigit, Alp, Bloom, and Kerr (2018) estimate a value of 0.13, while Liu and Ma (2023) obtain a smaller estimate of 0.05.³⁴

Figure 6 plots the estimated differences in log research effort (ξ_i values, x-axis) against the difference in log research output (patents, y-axis) rates across the two economies. Here we are including eight lags in our calculation of the knowledge stock, but the researcher allocations are not very sensitive to this choice. We can see that differences in the allocation of research effort are similar to differences in research output, but that they are not identical. The differences are due to the fact that the knowledge stocks differ in the two economies, which affects the productivity of research effort. This demonstrates why it is important to use the model to back out research effort, rather than simply plugging in observed research output.

9.1 Implied growth differences

We now use our model and estimates to quantify the growth implications of the differences in research allocations in the two economies. We begin by decomposing our estimate of

³⁴One reason that the λ value in our setting may be somewhat larger than in modern settings is the high cost of patenting and the difficulty in enforcing patent rights (Khan, 2005), which we would expect to generate a larger inventive step size.

Table 4: Quantification results

			Low	Med	High
Estimated growth diff. (pp.)	0.45	Observed growth diff. (pp.)	0.6	0.9	1.2
		Share explained (%)	0.75	0.5	0.38

The growth difference is estimated as $g(\mathbf{r}^{UK}) - g(\mathbf{r}^{FR}) = \lambda \mathbf{a}' \times \boldsymbol{\zeta}$ (using Eq. 4). To be conservative, we set $\lambda = \lambda_\infty = 0.14$, use the joint network for calculating \mathbf{a} , and the demeaned $\hat{\boldsymbol{\zeta}}$ from a NLS regression of Eq. 11 using eight lags of the knowledge stocks. The observed growth difference gives lower, upper, and mid-point values based on industrial production growth estimates for the UK and France during ca 1815–1860 (see footnote 13 for details).

the difference in log research effort, $\hat{\boldsymbol{\xi}}$ into a mean-zero vector representing the differences in allocations across technology types, $\hat{\boldsymbol{\zeta}}$, and a common level shift in overall research effort $\hat{\theta}$. We can then plug $\hat{\boldsymbol{\zeta}}$ into Equation 4 to obtain the predicted difference in growth rates given the differences in the allocation of log research effort across technologies and the innovation network, independent of any overall difference in research effort.

On the left side of Table 4, we present our preferred estimate of the growth effect resulting from differences in the allocation of research effort in the UK compared to France. This uses our estimate of the asymptotic λ value as the number of lags of the knowledge stock increases (0.14). As the estimated growth difference is increasing in the λ parameter, the reported results are conservative relative to what we would obtain when using fewer lags of the knowledge stock.

As discussed in the historical background, the existing literature estimates that the British industrial production grew between 0.6 and 1.2 percent per year faster than the French industrial production during this period. The right side of Table 4 shows that the effect of differences in the allocation of researchers across sectors between the UK and France can account for between 0.38 and 0.75 of the difference in the growth of industrial production in these two economies.

Our method is robust to unobserved differences in the knowledge content of patents in the UK vs. patents in France as long as those differences are constant across technology types. Such differences could be due to factors such as differences in the cost of patenting or the probability that a patent is enforceable in the two locations. To see this, suppose that the knowledge content of an observed patent in the UK is proportional to the knowledge content of a French patent up to a constant γ , so $n_i^{UK} \equiv \gamma n_i^{FR}$. If we plug this into Eq. 11 above, we can see that the only impact of a proportional difference in the knowledge content of patents is a level shift in the estimated ξ_i terms. When we decompose the $\boldsymbol{\xi}$ vector into $\boldsymbol{\zeta}$ and θ , this level shift will end up in θ . As a result, the growth implications of differences in the allocation of log research effort will be unaffected.

Table 5: Quantification results with international spillovers

	Included lags of knowledge stock					
	2	4	6	8	10	12
λ	0.346 (0.147)	0.401 (0.082)	0.254 (0.057)	0.264 (0.050)	0.207 (0.042)	0.140 (0.035)
X	0.511 (0.106)	0.416 (0.068)	0.266 (0.041)	0.254 (0.039)	0.203 (0.031)	0.151 (0.024)
ι (implied)	1.238	1.019	1.024	0.982	0.991	1.038

Non-linear Least Squares. Standard errors in parentheses.

9.2 With international spillovers

The results in the previous subsection treat France and the UK as if they were isolated economies. However, it is possible that new ideas produced in one economy could affect technology development in the other through the innovation network. To see how this affects our results, we estimate the following equation, which is a modified version of Eq. 11 that incorporates international spillovers (as described in Eq. 5):

$$\ln n_{it}^{UK} - \ln n_{it}^{FR} = \xi_i + (2\iota - 1)\lambda \times \left[\sum_{s=1}^v e^{-\lambda s} \left(\sum_{j=1}^K \omega_{ij} \ln n_{j,t-s}^{UK} ds - \sum_{j=1}^K \omega_{ij} \ln n_{j,t-s}^{FR} ds \right) \right] + \epsilon_{it} \quad (12)$$

Estimating this equation allows us to separately identify both ι , which governs the importance of international spillovers, and λ . Table 5 presents the results of this exercise, where $X = (2\iota - 1)\lambda$. The main take-away from these results is that ι is very close to one, which implies that international spillovers are essentially irrelevant in the context we study. Note that this does not mean that France and Britain were not utilizing new technologies developed in the other country. Instead, it simply tells us that the rate at which one country produced new ideas did not benefit substantially from ideas produced in the other. This makes sense. Long-distance communication was challenging at this time. Technological progress was also accelerating. As a result, by the time new ideas produced in one country arrived in the other they were often behind the frontier of knowledge. Of course, there are likely to be a number of exceptions to this. However, our results suggest that, by and large, the spillovers generated by new ideas produced in one country did not substantially affect innovation rates in the other.

9.3 Contributions by technology category

One valuable feature of our approach is that we can separate out the growth impact by technology category. Table 6 illustrates this for those sectors that contribute the most to the growth advantage of either Britain or France.

Table 6: Which technology categories were the most important for each country?

Top British allocation difference	ξ	Top French allocation difference	ξ
Agricultural Produce	0.68	Flour milling	0.63
Steam; Steam-Engines & Boilers	0.62	Coffee, Cocoa, Chocolate, & Tea	0.59
Hinges, Joints, & Pulleys	0.54	Boots, Shoes, Clogs, Pattens, etc	0.58
Fuel	0.48	Games, Exercises, & Amusements	0.5
Brewing, Distilling, & Rectifying	0.46	Lighting; Lamps & Luminaries; Matches	0.39
Top British growth contribution	p.p. diff.	Top French growth contribution	p.p. diff.
Steam; Steam-Engines & Boilers	0.42	Lighting; Lamps & Luminaries; Matches	0.17
Fuel	0.14	Heat, Heating, Evaporating, Concentrating	0.1
Ship-Building, Rigging, & Working	0.13	Flour milling	0.08
Water & Fluids	0.12	Boots, Shoes, Clogs, Pattens, etc	0.07
Metals & Metallic Substances	0.1	Spinning & Preparing For Spinning	0.07

The top panel of this figure presents (left side) the five technology categories where Britain had the largest relative allocation of research effort, as indicated by the ξ values estimated using Eq. 11, and (right side) the top five categories where France allocated more research effort. The bottom panel of the graph presents the top five technology categories in terms of the contribution each made to the country's growth. These are estimated by comparing the growth difference between the two countries from our main results and the counterfactual growth difference obtained when setting the ξ_i value for that technology category to zero with all other ξ_i values held constant.

In the top panel of Table 6, we list the top five categories in which Britain had the largest relative allocation of research effort compared to France (left side), and the top five in which France had the largest allocation relative to Britain (right side) (these are the ξ values estimated using Eq. 11). It is notable that British research effort was particularly high in steam engines as well as other industrial areas such as fuel, and agricultural technologies, while French efforts were generally more focused on consumer products.

The most interesting results are in the bottom panel, which shows the five sectors that made the greatest contribution to Britain's growth advantage (left side) and the five sectors that contributed the most to France's relative growth (right side). These contributions are the difference between the baseline growth rate difference and a counterfactual difference that we obtain by setting ξ_i to zero for one technology at a time, as if there was no difference in relative research effort in that technology.³⁵ Note that these contributions depend on both

³⁵This exercise essentially amounts to equalizing the number of researchers in the two countries in one category, but not reallocating those researchers across the remaining categories, so that relative effort in the remaining categories is unchanged. An alternative approach would be to reallocate the researchers proportionally across the other categories. Both approaches are valid, but they require slightly different

the difference in relative research allocations in a category and the importance (eigenvector centrality) of that category within the innovation network.

The results in the bottom panel of Table 6 show that steam engine research made the greatest contribution to Britain’s growth advantage, followed by fuels, shipbuilding, water and fluid technologies (pumps, etc.), and metallurgy. This list fits the historical narrative emphasizing the importance of Britain’s preeminence in mechanical technologies well; these are all areas where historical evidence suggests that Britain had a technology lead, and they are all technologies thought to have played an important role in the Industrial Revolution. Steam engine technology was particularly important; Britain’s much larger allocation of research effort to that area contributed 0.42 percentage points to the country’s growth advantage over France. It is interesting to contrast this with technologies related to agricultural produce, a category where Britain allocated even more research effort in relative terms (see top-left panel), but which made less of a contribution to the country’s growth advantage. This reflects the more peripheral position of agricultural produce technologies in the innovation network, compared to steam engines.

9.4 Additional discussion of quantitative results

The quantitative results presented above are obtained using $n_{it} = Patents_{it} + 1$, as discussed in Subsection 4.6. This mapping is intuitive given the model structure, because it means that a technology category in which has zero patents generates zero spillovers through the innovation network. However, it is natural to wonder how this setup affects our results. To examine this, we have looked at how our estimates change as we modify this approach. Table 7 shows that as we add a larger underlying rate of idea generation to the number of patents, the growth differences shrink. This makes sense, because as patents become relatively less important compared to the underlying unobserved flow of ideas, differences in research allocations are becoming less pronounced.

One message from these results is that, relative to the approach of setting $n_{ij} = Patents_{ij}$ used in most existing work, our approach will generate relatively smaller estimates of the growth difference. For this reason, we think of our baseline results as relatively conservative estimates of the growth differences induced by differences in the allocation of research effort across technologies between the Britain and France.

10 Conclusions

Did it matter that in the early decades of the Industrial Revolution many British inventors worked in technology types, such as steam engines or textile machinery, rather than technologies such as papermaking or chemicals? We argue that the answer is that, yes,

interpretations. We think the approach we have chosen is slightly more intuitive.

Table 7: Results for alternative mappings from patents to n_{it}

	$n_{it} = Patents_{it} + \dots$						
	0.001	0.01	0.1	0.5	1	1.5	2
Estimated $\lambda_{8\text{ lags}}$	0.206	0.222	0.232	0.232	0.231	0.229	0.221
Growth diff (p.p.)	2.172	4.305	2.614	1.3	0.747	0.469	0.3

We estimate Eq. 11 including 8 lags of the knowledge stock for a range of different values that, before taking logs, we add to $Patents_{it}$.

it mattered. Specifically, we show that the allocation of British inventors within the innovation network differed in fundamental ways from the allocation of inventors in the most natural comparison country, France, and that this allocation had a meaningful impact on the difference in technology growth rates in the two countries. To make this argument, we bring together frontier theoretical tools, rich historical patent data, and a novel approach to measuring the innovation network in a historical setting.

Our results enrich our understanding of the factors that contributed to Britain’s industrial dominance during the Industrial Revolution. They also contribute to a broader literature looking at the importance of innovation networks in economic growth, by providing direct evidence on the role that the structure of the innovation network played during an important period of economic history.

In addition to helping us better understand the nature Britain’s advantages during the early decades of the Industrial Revolution, our findings may also shed light on why these advantages slipped away in the late-nineteenth and early-twentieth centuries. It seems likely that the structure of the innovation network was slowly evolving over the nineteenth century, with the rising importance of chemical and electrical technologies that characterized the Second Industrial Revolution. This change in the technology space away from the mechanical technologies may help explain why Britain found it increasingly difficult to maintain its position as industrial leader. One interesting direction for future work is assessing the extent to which slow-moving changes in the underlying innovation network may have undermined Britain’s advantages and contributed to the erosion of British leadership in the late nineteenth and early twentieth century.

References

- ACEMOGLU, D., U. AKCIGIT, H. ALP, N. BLOOM, AND W. KERR (2018): “Innovation, Reallocation, and Growth,” *American Economic Review*, 108, 3450–3491. 33
- ACEMOGLU, D., U. AKCIGIT, AND W. R. KERR (2016): “Innovation Network,” *Proceedings of the National Academy of Sciences*, 113, 11483–11488. 1, 16
- 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. 8
- ACEMOGLU, D., S. JOHNSON, AND J. ROBINSON (2005): “The Rise of Europe: Atlantic Trade, Institutional Change, and Economic Growth,” *American Economic Review*, 95, 546–579. 1
- ALLEN, R. C. (2009a): *The British Industrial Revolution in Global Perspective*, Cambridge University Press. 1, 5, 6, 7, 19
- (2009b): “The Industrial Revolution in Miniature: The Spinning Jenny in Britain, France, and India,” *The Journal of Economic History*, 69, 901–927. 7
- ASSELAIN, J.-C. (2007): “Le Project Française d’ Histoire Économique Quantitative: Ambitions et Résultats,” *Economies et Sociétés (Serie ‘Histoire Economique Quantitative’)*, *Association Française de Cliométrie*, 36, 567–609. 4, 8
- BLOOM, N., M. SCHANKERMAN, AND J. VAN REENEN (2013): “Identifying Technology Spillovers and Product Market Rivalry,” *Econometrica*, 81, 1347–1393. 14, 26
- 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
- BROADBERRY, S., B. M. S. CAMPBELL, A. KLEIN, M. OVERTON, AND B. VAN LEEUWEN (2015): *British Economic Growth, 1270–1870*, Cambridge: Cambridge University Press. 4, 8
- 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, 7
- 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. 1
- CAMERON, R. E. (1958): “Economic Growth and Stagnation in France, 1815–1914,” *The Journal of Modern History*, 30, 1–13. 4
- CORREIA, S., P. GUIMARÃES, AND T. ZYLKIN (2019): “Verifying the existence of maximum

- likelihood estimates for generalized linear models,” . 26
- (2020): “Fast Poisson estimation with high-dimensional fixed effects,” *The Stata Journal*, 20, 95–115. 26
- 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. 7
- CRAFTS, N. F. (1977): “Industrial Revolution in England and France: Some Thoughts on the Question,” “Why Was England First?,” *The Economic History Review*, 30, 429–441. 1, 7
- (1995): “Macroinventions, Economic Growth, and Industrial Revolution in Britain and France,” *The Economic History Review*, 48, 591–598. 1, 7, 24
- CROUZET, F. (1996): “La Première Révolution Industrielle,” in *Histoire de La France Industrielle*, ed. by M. Lévy-Leboyer, Larousse. 4, 8
- FERNIHOUGH, A. AND K. O’ROURKE (2021): “Coal and the European Industrial Revolution,” *Economic Journal*, 131, 1135–1149. 1
- GRILICHES, Z. (1990): “Patent Statistics as Economic Indicators: A Survey,” *Journal of Economic Literature*, 28, 1661–1707. 14
- HALLMANN, C., L. ROSENBERGER, AND E. E. YAVUZ (2021): “Invention and Technological Leadership during the Industrial Revolution,” . 2, 6, 13
- HANLON, W. W. (2023): “The Rise of the Engineer: Inventing the Professional Inventor during the Industrial Revolution,” . 6, 13, 14, 61
- HARRIS, J. (2017): *Industrial Espionage and Technology Transfer*, London: Routledge. 6, 7
- HORREL, S., J. HUMPHRIES, AND M. WEALE (1994): “An Input-Output Table for 1841,” *The Economic History Review*, 47, 545–566. 14
- HUANG, J. AND Y. ZENOU (2020): “Key Sectors in Endogenous Growth,” . 1
- JAFFE, A. B., M. TRAJTENBERG, AND R. HENDERSON (1993): “Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations*,” *The Quarterly Journal of Economics*, 108, 577–598. 17
- JUHÁSZ, R. (2018): “Temporary Protection and Technology Adoption: Evidence from the Napoleonic Blockade,” *American Economic Review*, 108, 3339–76. 8
- 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, 5, 7
- (2023): “The Mechanics of the Industrial Revolution,” *Journal of Political Economy*, 131, 59–94. 1, 7
- 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, 33
- LANDES, D. S. (1969): *The Unbound Prometheus*, London: Cambridge University Press. 6, 19
- LIU, E. AND S. MA (2023): "Innovation Networks and R&D Allocation," NBER Working Paper 29607. 1, 2, 5, 8, 9, 10, 12, 16, 23, 24, 31, 33, 49
- 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
- MIURA, H. (2012): "Stata Graph Library for Network Analysis," *Stata Journal*, 12, 94–129. 51
- MOKYR, J. (1990): *The Lever of Riches: Technological Creativity and Economic Progress*, Oxford University Press. 2, 3, 6, 19, 24
- (1999): "Editor's Introduction: The New Economic History and the the Industrial Revolution," in *The British Industrial Revolution*, ed. by J. Mokyr, 2 ed. 6
- (2009): *The Enlightened Economy*, New Haven and London: Yale University Press. 1, 5
- MOKYR, J., A. SARID, AND K. VAN DER BEEK (2022): "The Wheels of Change: Technology Adoption, Millwrights and the Persistence in Britain's Industrialisation," *The Economic Journal*, 132, 1894–1926. 7
- MOSER, P. (2005): "How Do Patent Laws Influence Innovation? Evidence from Nineteenth-Century World's Fairs," *American Economic Review*, 95, 1214. 5, 6, 15, 60
- (2012): "Innovation without Patents: Evidence from World's Fairs," *Journal of Law and Economics*, 55, pp. 43–74. 2, 6, 15, 60
- NEWMAN, M. E. J. (2010): *Networks: An Introduction*, Oxford University Press. 51
- NICHOLAS, T. (2011): "Cheaper Patents," *Research Policy*, 40, 325–339. 6
- NORTH, D. C. AND B. R. WEINGAST (1989): "Constitutions and Commitment: The Evolution of Institutional Governing Public Choice in Seventeenth-Century England," *The Journal of Economic History*, 49, 803–832. 1
- NUVOLARI, ALESSANDRO, 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, 25
- NUVOLARI, A. AND V. TARTARI (2011): "Bennet Woodcroft and the Value of English Patents, 1617-1841," *Explorations in Economic History*, 48, 97–115. 6
- NUVOLARI, A., G. TORTORICI, AND M. VASTA (2023): "British-French Technology Transfer from the Revolution to Louise Philippe (1791-1844): Evidence from Patent Data," *The Journal of Economic History*, 83, 833–873. 6
- O'BRIEN, P. AND C. KEYDER (2011): *Economic Growth in Britain and France 1780-1914*, New York: Routledge. 4
- POMERANZ, K. (2000): *The Great Divergence*, Princeton: Princeton University Press. 1
- 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. 14
- VOIGTLÄNDER, N. AND H.-J. VOTH (2006): “Why England? Demographic Factors, Structural Change and Physical Capital Accumulation during the Industrial Revolution,” *Journal of Economic Growth*, 11, 319–361. 1
- VOTH, H.-J., B. CAPRETTINI, AND A. TREW (2022): “Fighting for Growth: Labor Scarcity and Technological Progress During the British Industrial Revolution,” Mimeo. 1, 7
- WOODCROFT, B. (1854): *Titles of Patents of Invention, Chronologically Arranged*, London: George Edward Eyre and William Spottiswoode. 43

A Concordant technology categories

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

A.1 Starting sample: Bi-national patents

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

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

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

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

Out of 916 “British patents” in France during 1844–1852, we can assign 855 (93%) to the original British patent as listed in Woodcroft (1854, Vol.2). The remainder is lost due to alterations of the patent title (e.g. shortening during translation) or the usage of patent agents in Britain (in which case the British patent lists a different inventor than the French “British” patents).³⁶

³⁶Curiously, these 855 French patents correspond to 808 British patents: some British patents were split into two when patented abroad.

A.2 Establishing concordance

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

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

To make the French classification correspond to the British, we assign French technology subcategories (keywords) to the new set of 133 (potentially concordant) categories. In a few cases, we cannot resolve ambiguity and do not assign keywords to any category. For example, one bi-national patent was classified into the French keyword “Yarn and fabric singeing” and into two British categories, “Spinning and Preparing For Spinning” and “Cloth Fulling, Dressing, Cutting, and Finishing.” Yet, both seem equally relevant—and too distinct to warrant aggregation. In sum, we can reliably assign French keywords to 123 concordant technology categories.³⁷

³⁷To 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 British IO matrix construction

This appendix provides some additional details and discussion of the methods used to construct input-output links between technology categories. Note that these links are primarily used in our analysis of the impact of macroinventions, which uses only British data. Thus, our focus is on construct input-output controls for that context.

The main challenge in constructing these matrices is mapping technology categories to the industries available in the IO matrix. To do so, we first try to match the occupation found in the patent data to IO industries. This is done through a manual review of the roughly 7,000 occupation titles listed in British patents from 1700-1849. A subset of these occupations unambiguously match to industries present in the IO table. Note that this does not always mean that the patented invention is associated with that industry; our assumption is that on average individuals working in a particular industry are likely to be invention technologies associated with that industry.

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.³⁸

Once we have a mapping from occupations to industries, the mapping from technology categories to industries is straightforward given that occupations are associated with patents which are classified into technology categories.³⁹ We can use this mapping, together with the information included in the IO matrix, to construct measures of the upstream and downstream links between different technology categories.

The resulting probabilistic mapping from technology categories to industries appear quite reasonable. To illustrate this, table B.2 lists, for the alphabetically first 30 technology categories i , the most important IO industry n (highest weight ϕ_{in}). In cases where an industry exists that is broadly similar to the technology category, this industry receives the highest weight: For example, the *Agriculture* technology category depends most on the *Agriculture, Forestry, etc* industry. This holds similarly for the technology categories *Building Processes*; *Building Materials*; *Boots, Shoes, etc*; and *Chemical Products*. Highly

³⁸Some professions are ambiguous even if qualified by a topic, for example “coal merchant” or “cloth merchant” because we do not know if this occupation worked in industry or in the excluded distribution services. One exception to the rule are composite occupations like “woollen manufacturer and merchant” because there the “manufacturer” clearly indicated involvement in production.

³⁹Two 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.1: Information used for matching input–output industries to occupations

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

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

specialized technologies are mapped with high precision into one industry (and the “correct” one), as in the cases of *Bell-Hanging* (to Non-ferrous metals), *Blackening* (to Boot and Shoe), or *Calculating Machines* and *Combs* (to Mixed Manufacture). Furthermore, technology categories that one would expect to have more diverse industrial applications place a low weight on the top industry, as for example the *Prevention of Accidents* technology category.

Table B.2: Most important input–output industry by technology category

Technology category	Input–output industry	φ_{in}
Accidents, Prevention Of	Non-ferrouse metals	0.15
Adhesive Substances	Chemicals	0.62
Aerated Liquors, Mineral Waters, etc	Chemicals	0.62
Aerial Conveyances	Furniture	1
Agricultural Produce	Food processing	0.31
Agriculture	Agriculture, Forestry, etc	0.41
Air & Wind ; -Air & Gas Engines & Windmills	Misc. Manufactures	0.28
Baths & Bathing-Machines	Misc. Manufactures	0.48
Bed & bedding	Furniture	0.3
Bell-Hanging	Non-ferrouse metals	1
Blacking	Boot and shoe	1
Bleaching, Washing, & Scouring	Textile finishing	0.39
Boots, Shoes, Clogs, Pattens, etc	Boot and shoe	0.62
Bottles, Vessels, & Jars, Covers & Stoppers	Misc. Manufactures	0.39
Brewing, Distilling, Rectifying	Drink	0.55
Bridges, arches, viaducts, aquaducts	Building, etc.	0.46
Brushes	Misc. Manufactures	0.81
Building & Relative Processes	Building, etc.	0.5
Building Materials-Burning Lime	Building, etc.	0.38
Buttons, Buckles, Studs, Dress-Fastenings	Misc. Manufactures	0.47
Calculating-Machines; Apparatus for Teaching	Misc. Manufactures	1
Candle Manufacture	Chemicals	0.91
Casks & Barrels	Drink	0.57
Casting	Iron and Steel	0.56
Chemical Products	Chemicals	0.86
Clocks, Watches, Chronometers, Timekeepers	Misc. Manufactures	0.98
Coaches & Other Road Conveyances	Iron and Steel	0.22
Coffee, Cocoa, Chocolate, & Tea	Food processing	0.48
Combs For The Hair	Misc. Manufactures	1
Condensing	Chemicals	0.67

The tables lists by technology category the most important IO industry, focusing on the alphabetically first 30 technology categories. The weights φ_{in} mapping IO industry n into technology category i are normalized to one by technology category.

C Validating our approach using modern data

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

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

Our inventor-based innovation network is constructed using the approach shown in Eq. 6. Our citation-based network is generated using the approach used in Liu and Ma (2023) as well as other modern studies:

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

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

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

D Additional results: Innovation network during the Industrial Revolution

D.1 Network plots

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

Figure D.1: Fully labeled joint innovation network



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

D.2 Network comparison

Table D.3 regression results comparing node centrality in the British network to that in the French network. We consider four centrality measures: Eigenvector, degree, closeness, and betweenness. Across all measures, British and French node centrality appear strikingly similar. This holds with small variations for both level and logarithmic specifications.

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

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

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

D.3 Stability of innovation network over time

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

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

$$\omega_{ijt} = \bar{\omega}_{ij} + \epsilon_{ijt} \quad (13)$$

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

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

$$\omega_{ijt} = \rho\omega_{ijt-1} + e_{ijt} \quad (14)$$

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. 14 and subtract out ω_{ijt-1} to obtain:

$$\omega_{ijt} - \omega_{ijt-1} = (\rho - 1)\omega_{ijt-1} + e_{ijt} \quad (15)$$

Using Eq. 15, we can obtain an estimate of ρ . Note that this is essentially a panel data analog of the setup for the Dickey-Fuller test. If we define $\delta = \rho - 1$, then we testing whether the model is a random walk amounts to testing the null hypothesis that $\delta = 0$. Alternatively, the model in Eq. 13 would suggest that we should observe $\delta = -1$.

Table D.4 presents results based on Eq. 15 using innovation networks constructed using patents from subsets of our study period. In Columns 1-3, we compare the innovation network based on patents from 1834-43 to a network based on patents from ten years before, 1824-33. Column 1 looks across the full matrix, while Columns 2 and 3 separate the diagonal and off-diagonal elements. In all three cases, we can reject both the random walk model ($H_0: \delta = 0$) and the white noise model ($H_0: \delta = -1$), though the fact that we obtain estimates that are fairly close to zero suggests that the model is closer to the one in Eq. 13 than Eq. 14. In Column 4, we restrict the period used to construct the dependent variable to include the same number of patents as the period used to construct the explanatory variable, which is less than ten years because the number of patents was growing rapidly over time. In Column 5, we show similar results using earlier periods. Overall, these results strongly reject the random walk model, and though they do suggest that $\rho > 0$, we find evidence that ρ is not large. Overall, these findings suggest that the matrix we observe is fairly close to one characterized by a fixed underlying structure.

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

	Dep var: Edge change $\Delta : \omega_{i,t} - \omega_{i,t-1}$				
	1834–1843			1840–1843	1824–1833
	(1)	(2)	(3)	(4)	(5)
	All edges	Diagonals	Off-diagonals	Same # pats p. period	Earlier period
Edge $\omega_{i,t-1}$: 1824–1833	-0.836*** (0.030)	-0.702*** (0.100)	-0.970*** (0.008)	-0.841*** (0.031)	
Edge $\omega_{i,t-1}$: 1790–1823					-0.799*** (0.037)
N (Obs = edges)	10500	97	10403	8649	9595
R^2	0.56	0.44	0.68	0.44	0.32


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

E Additional results: Effect of the network on innovation


E.1 Lagged knowledge stocks analysis

Here, we document the robustness of Figure 4 in the main text. Specifically, Figure E.2 replicates the baseline figure controlling additionally for downstream knowledge stocks. As is apparent, the general decay pattern for the upstream knowledge stocks is robust. The effects at low lags are large, effects at increasing lags are smaller, and effects at large lags ($t - 8$) are indistinguishable from zero. Notably, the effect size is shifted upwards across lags compared to the uncontrolled version. This is in line with the observation that the downstream knowledge stocks have a negative effect on patenting (which is much smaller in magnitude). The negative effect is also converging towards zero, though slower.

Figure E.2: The lagged effect of the knowledge stock on patenting rates



(a) Effects of upstream knowledge stocks



(b) Effects of downstream knowledge stocks

The figure presents coefficients and confidence intervals (95% in bold, 99% in thin) for PPML regressions based on Eq. 7, using British patents by domestic inventors. The innovation matrix is based on all patents (British and French). Patents appearing in multiple (N) technology categories count as only a fraction ($1/N$) of a patent in each of category.

E.2 Macroinvention analysis appendix

Here, we document the robustness of the macroinvention results from section 7.2. To begin, Table E.5 provides a more complete set of results corresponding to our preferred specification, in Panel A, Columns 2, 4, and 6 of Table 2, including the coefficients on the input-output connection controls.

Table E.5: Macroinvention robustness: IO coefficients

	Dep var: Ln (British patents + 1)		
	First list	Nuvolari list	Intersection list
	(1)	(2)	(3)
$\omega_{ie}^{UP} \times \text{post}$	0.434** (0.170)	0.704** (0.305)	1.652 (1.274)
$\omega_{ie}^{DOWN} \times \text{post}$	-0.216 (0.143)	-0.626** (0.294)	-0.109 (0.578)
$\omega_{ie}^{IO-UP} \times \text{post}$	-0.156 (0.376)	0.002 (0.682)	-1.986 (2.807)
$\omega_{ie}^{IO-DOWN} \times \text{post}$	0.308 (0.514)	-1.361 (0.903)	-1.276 (1.713)
Category \times event FE	✓	✓	✓
Period \times event FE	✓	✓	✓
N (Obs = category–event–period)	34296	18392	2084
No. clusters (category)	121	121	121
No. events (macroinv.–year)	84	45	5
Pseudo R^2	0.264	0.264	0.257
p upstream = downstream (ω)	0.0044	0.0042	0.23
p upstream = downstream (IO)	0.56	0.30	0.82

PPML regressions, implemented in Stata using *ppmlhdfc*. The specification is the same as in Table 2 (columns 2, 4, and 6). p upstream = downstream (ω) is the p -value of a Wald test (χ^2) for equality of upstream vs. downstream coefficients in the innovation network. Analogously, p upstream = downstream (IO) is the p -value for equality of upstream vs. downstream coefficients in the production network. Standard errors clustered by technology category in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table E.6 presents results using the macroinventions but dropping all other inventions by the macroinventors. This ensures that our downstream spillover results are not driven by subsequent patents by the inventors of the macroinventions. We can see that these estimates are essentially unchanged relative to our main results.

Table E.7 shows robustness of the baseline results to an alternative specification where we use the log number of patents on the left-hand side of our estimating equation, rather

Table E.6: Robustness – excluding patents by macroinventors

	Dep var: Ln (British patents + 1)					
	First list		Nuvolari list		Intersection list	
	(1)	(2)	(3)	(4)	(5)	(6)
Upstream shock $\times \omega_{ie} \times \text{post}$	0.426** (0.185)	0.416** (0.190)	0.719** (0.288)	0.766*** (0.292)	1.440 (1.324)	1.606 (1.276)
Downstream shock $\times \omega_{ei} \times \text{post}$	-0.246 (0.165)	-0.241 (0.168)	-0.706** (0.302)	-0.617** (0.300)	-0.172 (0.596)	-0.095 (0.590)
Category \times event FE	✓	✓	✓	✓	✓	✓
Period \times event FE	✓	✓	✓	✓	✓	✓
IO controls		✓		✓		✓
N (Obs = category–event–period)	31292	31292	18368	18368	2084	2084
No. clusters (category)	121	121	121	121	121	121
No. events (macroinv.–year)	84	84	45	45	5	5
Pseudo R^2	0.265	0.265	0.264	0.264	0.256	0.256
p upstream = downstream	0.0093	0.012	0.0021	0.0025	0.28	0.24

PPML regressions, implemented in Stata using *ppmlhdfc*. Compared to the baseline estimation, the number of observations is reduced because for more category–events, the outcome (patents) is zero in each event–period, and such category–events are separated by a fixed effect. p upstream = downstream is the p -value of a Wald test (χ^2) for equality of coefficients on upstream vs. downstream shock times post, respectively. Standard errors clustered by technology category in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table E.7: Macroinvention robustness: Ln (patents) instead of ln(patents + 1)

	Dep var: Ln (British patents, unweighted)					
	First list		Nuvolari list		Intersection list	
	(1)	(2)	(3)	(4)	(5)	(6)
$\omega_{ie}^{UP} \times \text{post}$	0.448*** (0.149)	0.454*** (0.151)	0.565* (0.296)	0.586* (0.300)	3.275** (1.610)	3.379** (1.646)
$\omega_{ie}^{DOWN} \times \text{post}$	-0.203 (0.123)	-0.201 (0.123)	-0.825*** (0.281)	-0.798*** (0.281)	-0.395 (0.629)	-0.361 (0.630)
Category \times event FE	✓	✓	✓	✓	✓	✓
Period \times event FE	✓	✓	✓	✓	✓	✓
IO controls		✓		✓		✓
N (Obs = category–event–period)	22197	22197	12232	12232	1321	1321
No. clusters (category)	118	118	118	118	109	109
No. events (macroinv.–year)	84	84	45	45	5	5
Pseudo R^2	0.208	0.208	0.208	0.208	0.193	0.193
p upstream = downstream	0.0011	0.00090	0.0025	0.0024	0.053	0.050

PPML regressions, implemented in Stata using *ppmlhdfc*. The specification differs from the baseline by using $\ln(x)$ rather than $\ln(x + 1)$ as outcome. This reduces the sample size because observations with zero patents ($x = 0$) will necessarily be excluded. p upstream = downstream is the p -value of a Wald test (χ^2) for equality of coefficients on upstream vs. downstream shock times post, respectively. Standard errors clustered by technology category in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

than our preferred $\log(\text{Patents} + 1)$.⁴⁰ Using as outcome the log of patents furthermore reduces the sample size, as observations with zero patents ($x = 0$) will necessarily be excluded.

We also estimate a more demanding “stacked event study” specification,


$$\ln(n_{ie\tau}) = \sum_{\tau} \beta_{\tau} \omega_{ie} \cdot \mathbb{1}(\tau) + X_{ie\tau} \Gamma + \gamma_{ie} + \eta_{e\tau} + \epsilon_{ie\tau} \quad (16)$$

where, different than before, we estimate β_{τ} flexibly for periods $\tau \in \{1, 3, 4\}$ —period $\tau = 2$ being the omitted reference period—, and also interact the controls $X_{ie\tau}$ with period indicators.

Figure E.3 presents the results.

⁴⁰Implementing this specification also requires us to abandon weighting British patents, which we do in the baseline to adjust for some patents being listed in more than one technology category. When specifying the outcome as $\ln(x)$ and (down-)weighting patents with more than one category, the outcome might become negative in some cases and the model could not be estimated via PPML.

Figure E.3: Macroinvention event study results



The figure presents estimated coefficients and confidence intervals (95% in bold, 99% in thin) for PPML regressions of log patents on the exposure to a macroinvention that occurs upstream or downstream of a technology category (interacted with period indicators). The five-year period just before the macroinvention is the omitted reference category. The regression controls for upstream and downstream IO connections to the macroinvention category (interacted with period indicators). See the text for the definition of macroinvention lists.

F Innovation network based on exhibition data

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

As explained in the main text section 4.5, 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.⁴¹

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

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

$$\omega_{ij}^{Exhibit} = \frac{E_{kij}}{E_i + E_j}$$

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

Note that this approach is very similar to that applied to the patent data, except that the resulting exhibition network is undirected. The reason why we can’t distinguish E_{ij} from E_{ji} is that both exhibits are at the same time; we do not have information in which sequence the exhibitor created them. As a result, $\Omega^{Exhibit}$ is symmetric.⁴²

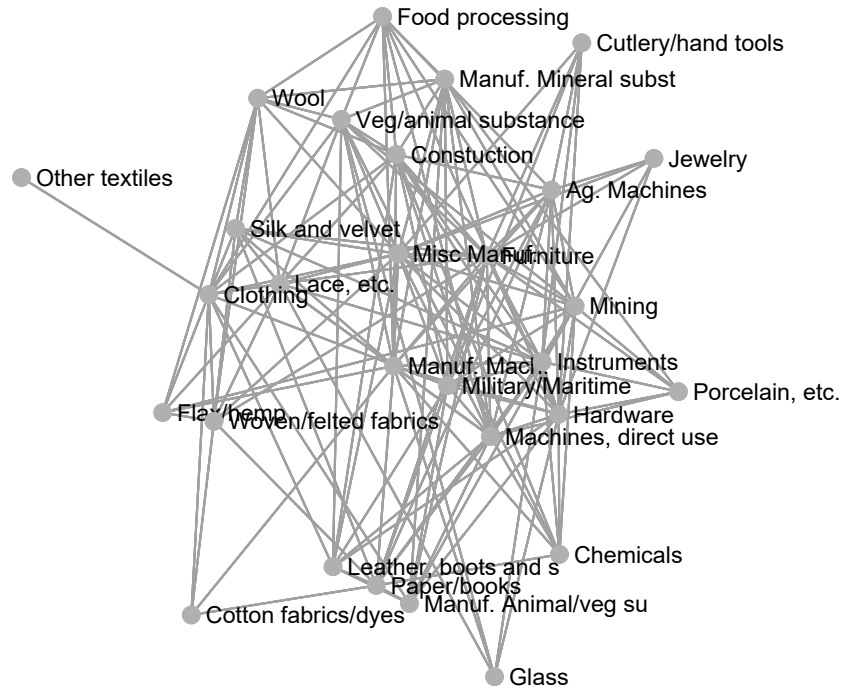
Figure F.4 presents the innovation network obtained by applying this method to the exhibition data.⁴³ The network features a dense central core, located just below and to the right of the center of the plot. The core includes technology types such as *Machines for direct use* (including steam engines and locomotives), *Manufacturing machines*, *Military/Maritime* (including shipbuilding), *Instruments*, *Mining*, and *Hardware*. Similarities between this

⁴¹The original dataset includes 30 categories, but we omit category 30, which is for fine arts exhibits. For some French exhibits, it also provides secondary or tertiary categorizations for French exhibits. But as most exhibits have only one category, we focus on each exhibit’s first category.

⁴²The interpretation of the resulting exhibition-based innovation network is not as straightforward as the one obtained from a patent-based innovation network, since exhibiting in two categories might more likely reflect factors other than innovation spillovers.

⁴³Note 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.4: The technology network obtained from exhibition data



Plot is generated using multidimensional scaling.

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

To compare the exhibition-based innovation network to the patent-based innovation network more directly, we need to map patent technology categories into exhibition categories. One potential approach is to assign patent categories to exhibit categories based on their descriptions. However, this works well only for some categories but not for others. Instead, we adopt an approach that relies on inventors who both filed patents and exhibited at the Crystal Palace. Specifically, we use a list of British-based inventor-exhibitors from Hanlon (2023). The list was constructed by manually searching the British patent data for each exhibitor and also for inventors of exhibits mentioned in the exhibit descriptions, which might be different for example because ownership of a technology had been transferred. This manual review identified 634 exhibits associated with an individual that also patented, with links for 566 individuals in total (some individuals account for more than one exhibit).⁴⁴

⁴⁴Note 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.8: Comparing patent-based and exhibit-based innovation network

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

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

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

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 should also have more central exhibits within the exhibit network.

We assess the prediction of whether inventors with more central patents also exhibited more centrally in Table F.8. The unit of analysis is inventor–exhibitors, individuals who appear in both the patent and exhibition data. The dependent variable is the average eigenvector centrality of the inventor’s patents in the (British) patent network. The explanatory variable is the average eigenvector centrality of the inventor’s exhibits in the exhibition network. We observe a strong positive relationship between the centrality of an individual’s patented and exhibited inventions, which holds for different empirical specifications (centrality in levels, logs, or ranks). This result indicates strong similarities between the patent-based network and that obtained from exhibition data. Thus, we conclude that patterns observed in the patent data reflect patterns that can also be observed in data that do not rely on the filing of patents.

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.9 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*

Table F.9: Distribution of British and French exhibits

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

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

Table F.10: Country centrality within the exhibition network

	Dep var: Exhibition eigencentality		
	Level	Log	Rank
	(1)	(2)	(3)
Exhibit ratio Britain vs France	0.004* (0.002)	0.176*** (0.051)	0.452*** (0.129)
N (Obs = exhibit category)	29	28	29
R^2	0.126	0.180	0.204
Std. β	0.356	0.424	0.452

The table documents that, across exhibition categories at the London World's Fair 1851, a higher ratio of British to French exhibitors was significantly and strongly associated with higher network centrality. Eigencentality is calculated as undirected eigenvector because the exhibition network is undirected. OLS regression, robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

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.

In Table F.10, we study more systematically how the centrality of French exhibitors differs from that of British exhibitors within the exhibition-based innovation network. The unit of observation is the exhibition category. The dependent variables is a category's eigenvector centrality in the exhibition-based network. The explanatory variable is the ratio of British to French exhibitors in each category. Across different empirical specifications (levels, logs, or ranks), the table shows that British inventors tended to locate in more central technology categories within the network. The finding confirms the patterns documented in Table 3 for the patent data. Thus, British inventors appear more central than French inventors even in data covering many unpatented inventions.