PAUL A. DAVID

Paul Allen David

24 March 1935 – 22 January 2023

elected Fellow of the British Academy 1995

by

TIMOTHY GUINNANE

GAVIN WRIGHT

Paul David was an economic historian, a major figure in redefining the methods and philosophical bases for that hybrid field. An early practitioner of cliometrics – the study of historical economies using rigorously specified models and quantitative evidence – Paul was a strong advocate of the view that economics should be an historical social science, i.e., that 'history matters'. Having strong roots in both the US and the UK, Paul's influence extended well beyond narrow research specialties, including pathbreaking contributions to macroeconomics and growth accounting, technological innovation and diffusion, historical demography, natural resources, and the historical origins of European institutions of science and technology.

Biographical Memoirs of Fellows of the British Academy, 21, 403–424
Posted 1 December 2023.

Education and dissertation

Paul David was born in New York City, the son of Henry David and Evelyn Mae Levinson David. Henry David was an accomplished historian of the American labour movement, who wrote a famous book about the Haymarket Square affair (1936). Paul studied both history and economics at Harvard (1952–56), taking Alexander Gerschenkron's graduate course in economic history during his senior year. Having won a Fulbright scholarship after graduation, Paul spent the next two years at Cambridge University, developing lasting relationships with such leading figures as R.C.O. Mathews, M.M. Postan and H.J. Habakkuk. His Cambridge experience had a lasting influence, endowing Paul with greater historical sensibilities than most cliometricians.

Returning to Harvard as a graduate student, Paul joined Gerschenkron's seminar, famous as a launching site for what was then known as the New Economic History, now more commonly called Cliometrics: the application of rigorously specified models and quantitative methods to the study of economic history. Though not a cliometrician himself, Gerschenkron had high academic standards and fostered a sense of competitive camaraderie among the members. The result was a remarkable cluster of distinguished alumni.

Paul was very much part of this milieu, and his dissertation project was perhaps the most ambitious. By the time he received a job offer from Stanford in 1961, Paul had prepared a 900-page manuscript on the economic history of Chicago, complete with tables, formulas, appendices, and addenda. Yet the thesis was not submitted at that time, and ultimately was 'set aside'. The reasons for this turn of events are far from clear, but lurking in the background was the word among Gerschenkron's students that 'nothing was ever quite good enough.' As Paul well knew, Gerschenkron insisted that Albert Fishlow invest an additional two years revising his completed thesis on antebellum railroads, before granting his approval. Whatever the full explanation might be, similar patterns recurred throughout Paul's career: ambitious projects were planned, and lengthy drafts written, yet somehow the completed books never appeared. Paul was always in pursuit of his next new idea.[1]

Noncompletion of the dissertation did not mean that Paul was inactive. The late 1960s and early 1970s were some of his most creative and productive years, marked by original contributions in growth accounting, learning effects, technology and technological diffusion, and in developing the case for a more historical version of cliometrics. By 1969, when an offer from Oxford beckoned, Stanford promoted Paul to a tenured position, and Gerschenkron strongly supported the move (Dawidoff 2002: 267). Paul subsequently received his PhD from Harvard in 1973, based on his impressive portfolio

[1] These paragraphs draw upon Dawidoff (2002), a biography of Gerschenkron written by his grandson.

of published articles. In essence, the collection comprised Paul's first book, published in 1975.

Accounting for US economic growth

One of the New Economic History projects of longest standing was compiling the quantitative record of economic growth for the United States, a good example of collective learning.[2] Owing to the research of Simon Kuznets and his student Robert E. Gallman, a consensus had developed that acceleration of US growth predated the Civil War. But was this acceleration gradual or episodic, propelled by a Rostovian 'take-off' in the 1840s? These questions were stymied by an absence of systematic data prior to the federal census of 1839: hence the label for the pre–1840 era as a 'statistical dark age'.

Paul's contribution was to assemble fragmentary estimates of the gainfully occupied labour force and productivity by sector, to create what he called 'controlled conjectures' about the likely course of real product growth during these decades (David 1967a; 1967b). His conclusion was that although antebellum growth was clearly higher than during the colonial era, there was no sharp break in the growth trend at any point. Although later treatments lean more towards a finding of acceleration during the 1850s, in effect Paul opened a broad new range of considerations, which subsequent scholars have had to address.

Paul himself returned to the topic thirty years later, to adjust a feature of the exercise that had nagged him all along: part of his growth estimate derived from a productivity 'gap' between farm and nonfarm sectors. Although such a gap is common in historical cases of development, American economic historians teach their students that the challenge for manufacturing was the high opportunity cost of labour, set by the prospect of self-employment in an expanding farm sector. Revisiting the 1839/40 sectoral productivity figures from this perspective, Paul produced a new set of growth estimates which remain unpublished (1996). The upshot was even more remarkable than before: more than half of the 0.9 per cent average annual growth during the antebellum era was attributable to an increase in manhours per member of the population, raising major questions about the welfare content of this growth.

These papers showed Paul's capacity for pursuing issues across many decades, maintaining and updating his data files even while his primary academic efforts were directed elsewhere. On the growth-accounting front, the next phase was in collaboration with Paul's Stanford colleague Moses Abramovitz, already known as a chronicler of long-term trends in the US economy. Writing in the 1950s, Abramovitz, along with

[2] See for example, W.N. Parker ed. (1960) and D.S. Brady ed. (1966).

Robert Solow, established the ‘stylized fact’ that most US growth was attributable not to expansion of inputs but to the ‘residual’, often labelled ‘technological progress’ but more appropriately seen (in Moe’s famous phrase) as ‘a measure of our ignorance’. Abramovitz & David (1973a; 1973b) showed, however, that this pattern only held for the 20th century, not for the 19th. As US growth per capita rose from 1.1 per cent to 1.6 per cent per year across the century, most of this acceleration was attributable to growth of inputs: labour, land and capital. Thus, growth was far from the ‘balanced growth’ path posited in theoretical models, and the ‘rise of the residual’ should be seen not as a ‘stylized fact’ but as a distinct historical development.

Paul carried this analysis further, in a solo-authored article published only in a relatively obscure conference series on public policy (1977). US growth acceleration was depicted as a ‘grand traverse’ from one steady-state path to another, propelled by an increase in the real gross savings rate from about 10 per cent before 1834 to 28 per cent during 1889–1898. This pronounced capital-deepening was not, however, attributed to a national increase in Thrift, but to a capital-deepening bias in the Progress of Invention, reflected in a long-term fall in the relative price of investment goods. On the supply side, the rise in measured savings was seen less as an increase in forward-looking or acquisitive behaviour, but more as a portfolio shift away from unconventional and largely non-market forms of investment, primarily land-clearing, child-raising and slave-rearing. This elegant synthesis, acknowledged as a Voltairian ‘fable’ though with deep quantitative documentation, was shown to be broadly consistent with the views of late–19th century observers such as Böhm-Bawerk, Sidgwick and Taussig regarding the course of economic progress. With hindsight, the article suggested, one explanation for the pattern is that rapid progress in mechanical engineering had not yet been joined by advances in chemical or electrical engineering.

Abramovitz and David were careful to stipulate that the small relative size of the conventional residual did not imply that the 19th century saw only limited ‘technological progress’; instead, the growth accounting exercise served to underscore the shortcomings of total-factor-productivity as an index of technology. This critique of conventional usage was reinforced by research with Gavin Wright on the historical origins of American ‘resource abundance’ (David & Wright 1997). The project grew from concerns during the 1980s that the US was losing ‘technological leadership’ to Japan and other nations. But what was this ‘leadership’ and where did it come from historically? An initial study, using trade content as an indicator, found that the most distinctive feature of US manufacturing exports around the turn of the 20th century was intensity in nonrenewable natural resources: minerals (Wright 1990). Indeed, the US was the world’s leading producer of virtually every one of the world’s major industrial minerals during this era, a consideration which at first blush might call into question the entire notion that the country was a technological ‘leader’ in any true sense.

Closer study revealed, however, that dominance in minerals was not based on geological ‘endowment’ but emerged historically, roughly between 1870 and 1910. The driving forces were intense exploration, enabled by an accommodating legal environment on the public domain; investment in the ‘infrastructure of public knowledge’, largely through the work of the US Geological Survey; and the training of mining engineers, not in one national academy but in competing state mining schools. Contrary to intuition, the American minerals sector presented many traits associated today with the ‘knowledge economy’.[3]

Influenced by contemporary concerns, attention thus turned to the US growth and productivity record in comparative historical context. For a volume to mark the opening of Stanford’s Landau Economics Building, Abramovitz & David (1996) offered a synthesis of what might be called a ‘Stanford’ interpretation of the historical basis for American productivity leadership: less formal than growth accounting, but perhaps for that reason, more nuanced and comprehensive. The authors pointed to a ‘fortunate concordance between America’s own exceptional economic and social characteristics and the nature of the dominant path of technological progress and labor productivity advances’ (25). That path was natural resource-intensive, tangible capital-intensive, and scale-dependent in its elaboration of mass-production and high-throughput modes of business organisation. This unique ‘congruence’ between emerging technologies and US national traits imposed limitations on the ability of other nations to grow rapidly simply by importing American technologies. In the post-Second World War era, however, these inter-country differences steadily declined: mineral deposits were discovered at new locations around the world; transportation costs and trade barriers fell, offering larger-scale export markets to many countries; technological progress reduced the shares of primary products in final costs. These trends set the stage for the postwar ‘convergence’ process, in which national growth rates were systematically and inversely related to the gap with the US level – but which the authors prefer to call ‘catch-up’ because of the distinctive historical context. The larger message would seem to be that US productivity leadership arose from distinctive historical circumstances rather than unique national virtues or divine endowment. Hence the ‘waning of American exceptionalism’ need not be seen as either a symptom of decline or a reason for alarm.

The final chapter in the Abramovitz-David collaboration was their entry in Volume III of *The Cambridge Economic History of the United States* (2000). The essay consolidates the growth-accounting evidence, the primary theme being the transition from extensive to intensive sources of growth from the 19th to the 20th century, from expansion

[3] This characterisation of minerals is complemented by the work of Olmstead & Rhode (2011) for farmland. They show that the expansion of American agriculture into geographic areas with unfamiliar soils, weather and precipitation levels required extensive adaptations in crop varieties and farming practices.

of inputs to the acquisition and exploitation of knowledge. Echoing earlier writings, the authors stress that this view does not imply that the 'pace of technological progress' was faster in the 20th century, but its biases and institutional sources shifted, to codified technologies explicitly derived from science, structured forms of research and development, and more highly educated workers. This account left the productivity slowdown of the most recent period (1966–1989) as something of a puzzle, to which the authors suggested at least the form of an answer in their discussion of growth driven by a series of 'general-purpose technologies', each of which required time for adaptation and complementary innovations. The chapter constitutes the most accessible yet rigorous presentation of the Abramovitz-David analysis.[4]

Technological innovation and technology diffusion

Paul's first study of technological diffusion grew out of his dissertation research on Chicago and was published in a *Festschrift* for Gerschenkron in 1966. After calling attention to the substantial role of farm implements and machinery in midwestern industrialisation, Paul then posed the puzzle of the twenty-year lag between invention and adoption of the McCormick reaper. The main technical breakthroughs occurred in the early 1830s, yet the first major wave of adoptions occurred only midway through the 1850s. Paul's answer set frictions, information issues and risk aversion to the side, because of one basic consideration: adoption only made economic sense if the farmer planted a large enough acreage to justify the large fixed cost of purchasing a reaper. Average planted acreage surpassed this 'threshold' level between the 1830s and the 1850s for three main reasons: farm size grew as land-clearing proceeded; the price of the reaper declined; and wages of harvest labour increased during the Crimean-war boom years. When the distribution of farm sizes crossed over the threshold, the result was the familiar sigmoidal 'take-off' of reaper adoptions.

The midwest reaper paper was paired with one on Victorian Britain (1971), explaining the *limited* adoption of McCormick's machine in that country during the 1850s and 1860s – indeed a monograph-length comparative study was planned, though not further pursued. In the British case, the simple threshold calculations were complicated by the need for extensive investment in drainage and reconfiguration of the cultivated areas: hence the phrase 'technical interrelatedness'. Such investments were further discouraged by the prevalence of farm tenancy as opposed to owner occupancy, introducing thorny additional co-ordination and distributional considerations. The upshot was that such mechanisation-cum-investments would not have repaid loans at 3.5 per cent, at cost conditions prevailing before 1875.

[4] A poignant account of the chapter's completion may be found in Abramovitz's memoir (2001).

Subsequent research has not been particularly favourable to the simple application of the threshold model to antebellum reaper diffusion. Using the records of the McCormick company, Olmstead (1975) identified large numbers of joint reaper purchases by two or more households, and argued further that the machines became more popular over time largely because they became not only cheaper but more reliable. Olmstead & Rhode (1995) counted many more small-scale purchasers and found that rapid diffusion was fostered by local markets and co-operative exchanges, questioning the entire individualist framework of the threshold model. Nonetheless, despite these findings, census data for 1860 show a clear positive relationship between implement-intensity and farm acreage, so that the notion of coevolutionary adaptation between technology and the family-farm social unit remains viable.

Perhaps more importantly, the reaper paper provided a template for what came to be called 'equilibrium' models of diffusion, which ultimately became central to Paul's reputation and impact. The core proposition was that even if agents are well-informed and eager to pursue profit opportunities, new technologies will generally be adopted only with a time lag, reflecting the need for complementary investments and adaptations before they become profitable. Thus, technology diffusion deserves attention equal to invention, for understanding history as well as for designing growth policy. Paul published theoretical pieces on aspects of this phenomenon, with various co-authors, during the 1980s. The breakthrough came in 1990 and 1991, with the appearance of a pair of articles proposing an analogy between modern Information Technology and historical experience with electrification. The longer version was detailed and nuanced; the shorter version became one of Paul's most influential works.

The issue of the day was the 'Solow paradox', epitomised by the Nobel laureate's remark: 'We see computers everywhere, except in the productivity statistics.' New technologies were ubiquitous, yet productivity growth stagnated, beginning in the 1970s. Paul pointed out that the same pattern occurred with electrification: the main technological innovations dated from the 1880s, yet widespread adoptions and the associated productivity surge did not happen until the 1920s. Reasons for the delay included the need for investment in an infrastructure of central station generators, and at the micro level, the need for new single-story factories, whose layout could be designed to streamline the flow of materials. These conditions were met with the investment boom of the 1920s. Later analysis confirmed that electrification was the General Purpose Technology driving productivity growth in manufacturing across the decade, reducing capital-labour ratios in virtually every sector (David & Wright 2003).

'The Dynamo and the Computer' acquired cult status because it seemed to predict the IT-driven productivity surge of the late 1990s. Of course, Paul never intended to advance a 'forecast' in any precise sense, and indeed he cautioned 'against the dangers of embracing the historical analogy too literally' (1990: 360). With the aid of hindsight, however, the parallels between the two historical episodes are indeed striking and instructive. In both cases, mere ‘adoption’ did little to advance productivity. The major gains came only when the new technology was deployed to restructure and reconceptualise work assignments and responsibilities. In the 1990s, IT was associated with diffusion of a new set of practices known as high-performance work organisation (HPWO), including job rotation, pay for knowledge, autonomous teams, total-quality management, and quality circles (Cohen *et al.* 2001). Also in both cases, the channeling of new technology in labour-saving directions was fostered by higher wages and tighter labour market conditions. In other words, as Paul was at pains to emphasise for years afterwards, his contribution should not be read as propounding a law of history, but as an effort to better understand why historical patterns of productivity growth often occur in alternating periods of surge and pause.

Path dependence: the search for historical economics

Economic history seems always to struggle for respect and, sometimes, for survival. In one such effort, W.N. Parker organised a session at the 1984 meetings of the American Economic Association, inviting such luminaries as Kenneth Arrow, James Tobin, Robert Solow and Sir Arthur Lewis to make the case for economic history. Asked to participate, Paul devoted his presentation to a set of ideas that he had been discussing with W.B. Arthur and other Stanford colleagues for some time: history seen as a stochastic path dependent process, in which remote past events can exert lasting influence on subsequent outcomes. The paper was organised around one compelling example: the persistence of the QWERTY typewriter keyboard layout into the electronic age, despite the existence of demonstrably better-performing alternatives. As with Dynamo and Computer, the longer version (1986) told the story in full detail; the short version (1985) became one of the most cited works in all of economics.

The gist of the story was that the QWERTY keyboard represented a chance solution to an engineering design problem in early typewriters, when the speed of typist tended to outrun the efficacy of machine, causing jamming of the keys. The intent therefore was to spread the most frequently used keys, slowing rather than speeding the pace of typing. Several keyboards competed for market share at the time, but with the advent of ‘touch-typing’ and its incorporation into the curricula of leading schools, QWERTY emerged as the standard. Clearly intending the case not as an eccentric anomaly but more broadly as a metaphor for historical economics, Paul identified three core elements at work: technical interrelatedness (the need for compatibility between the keyboard ‘hardware’ and the ‘software’ of touch-typist memory); economies of scale, not just conventional economies for an individual firm, but decreasing-cost conditions for the system as a whole; and quasi-irreversibility of investment in mastering any particular keyboard layout. The article ended with 'a message of faith and qualified hope … that the absorbing delights and quiet terrors of exploring QWERTY worlds will suffice to draw adventurous economists into the systematic study of essentially *historical* dynamic processes …' (1985: 47).

The response was electric, with much appreciation of the article's novelty and promise, but also including a fair amount of scepticism. Some critics saw QWERTY as just another example of 'market failure', a stalking horse for government intervention. Others acknowledged the possibility of such suboptimal outcomes but argued that the associated efficiency losses could not be large, or else self-interested actors would find a remedy.[5] Paul's response to this criticism was erratic. He devoted his 1989 presidential address to the Economic History Association to the theme of path dependence but never submitted the paper for publication. The same was true for his Marshall Lectures at Cambridge in 1992 and his Tawney Lecture in Glasgow in 2005. When Paul did reply in writing, the pieces were erudite and cogent, but also complex and didactic. Paul's 2007 article in *Cliometrica* listed more than thirty post-QWERTY works on path dependence, none in mainstream journals and no one of them packing the concise rhetorical punch of the original.

The strongest demonstrations of the value of the path dependence concept, however, are not theoretical possibilities but compelling examples from economic history, some written by Paul himself, others by his students or by scholars pursuing the same vision. In his formalisation of path dependence and in subsequent writing, Arthur (1989) offered several illustrations of outcomes driven by historical accident and increasing returns. The convention for 'clockwise' circular motion is one example. American cars driving on the right-hand-side is another, where the indeterminacy of the outcome is neatly shown by the fact that not all countries arrived at the same answer. The same can be said of railway gauge, where the strongest imperative for new construction is to choose the *same* gauge as those on incoming and outgoing lines, as opposed to the most efficient gauge. Historical bases for diverse regional and national outcomes receive detailed scrutiny from Paul's student Douglas Puffert (2009).

Another of Arthur's cases also comes from automobile history. In 1890 there were three competing models: steam, gasoline, and electricity, each with supporting engineers and enthusiasts. According to Arthur, the matter was only settled by an outbreak of hoof-and-mouth disease in 1914, which caused the removal of horse troughs and greatly increased the cost of steam-driven cars. The fate of early electric cars receives rigorous analytical attention from Paul's student David Kirsch (2000). In another example of competing networks, David & Bunn (1988) showed that the timing of a 'gateway

[5] The most persistent critics were S.J. Liebowitz and S.E. Margolis (1990; 1994).

technology' enabling compatibility between Alternating and Direct Current systems (the rotary converter) was critical for the resulting US configuration, which contrasts with outcomes in Britain and on the European Continent.

Were all these cases of minor economic significance? Hardly. The problem with this claim is that the implications of taking one path at a fork of history may be so far-reaching that they do not readily lend themselves to measurement as efficiency losses from conventional 'market failure'. Often the first task of the historian is to persuade the reader that an alternative historical path was actually conceivable. This can be a real challenge, because the path taken often seems in hindsight to be the only one possible, especially to economists who are experts at *ex post facto* rationalisation. But accepting this non-trivial challenge is indeed part of the research agenda suggested by path dependence as a concept. To be sure, when the shortcomings of a particular path are recognised, both technical and policy remedies may be undertaken. Interpreting these efforts as 'path-constrained amelioration' would be a different way of writing economic history and constitutes a second component of Paul's proposed agenda. These reflections pertain to technologies and technological systems, which have received most attention in this literature. They apply even more to extensions into the realm of 'conventions, organizations, and institutions', as Paul clearly envisioned (David 1994).

Historical demography

While engaged in these high-visibility debates, Paul also pursued what would seem to be an entirely separate research agenda on historical demography. One longstanding thread was the history of fertility control. Paul saw a natural connection with his broader interest in the adoption of technology: contraception can be seen as another technology, and the 19th and 20th centuries saw several purely technological changes in this area. As in other areas, Paul contributed both to devising new empirical methods and to efforts to understand a substantive problem from the past.

When Paul became interested in this topic in the 1970s, scholars from several disciplines had come to realise that many populations (including most of western Europe and North America) had begun to reduce fertility within marriage sometime in the 19th or early 20th centuries. Prior to this 'fertility transition', people in these societies controlled fertility mostly by delaying marriage or by not marrying at all. The fertility transition represents a dramatic change in the way human societies reproduced, and it arguably played a central role in shaping the economies we see today. Paul's research on the United States implied that by the turn of the 20th century, large minorities of married women were controlling fertility and aimed to have two children only (David & Sanderson 1987). While most scholars agreed that the fertility transition happened, there was (and still is) no consensus on its underlying causes. Economic historians looked for ways to understand how couples would adjust their fertility to changing economic opportunities and constraints. Most historical demographers (many trained as sociologists), on the other hand, saw the fertility transition as a reflection of changing ideas, such as 'secularization' or new thinking about the role of women in society.[6]

Paul was a keen student of this literature, characteristically. In contrast to many others who approach a new field, he focused on addressing unresolved research questions posed by the work of earlier scholars. Working with several different co-authors, he contributed to three overlapping areas. First, just how do we know that couples at the time sought to control their fertility? There is some evidence about intentions from court cases, letters, diaries and the like, but we lack the comprehensive surveys of fertility intentions that are common today. The best evidence remains indirect: families got smaller. Second, how did couples reduce fertility? While we know something about the efficacy of the 'rudimentary' contraceptive methods available in the late 19th century, we lack anything like a survey of practices that would enable us to state with confidence the precise role of any method. Partly because of these evidentiary problems, a few scholars remain sceptical of the idea that the methods available in the 19th century could achieve the fertility reductions we observe. This gap accounts for claims that fertility reductions reflected the spread of ailments (such as venereal diseases) that would reduce fecundity. The third issue comprises most of the debate: *why* did couples begin to desire smaller families? The 'why' question rests on how we answer the first two questions, however, because knowing why people reduced fertility relies critically on knowing what kinds of couples adopted fertility-control measures, and when they first did so.

Lurking in the background is the concept of 'natural fertility'. Natural fertility has a precise definition to a demographer: it is the absence of parity-specific fertility control. Parity is the number of births a woman has experienced, so in a natural fertility population, the probability of the next birth does not depend on the number of children already born. Economists and others often confuse natural fertility with large families, or fail to appreciate the concept's precision, so much ink has been wasted on scholars talking past one another. One of Paul's intellectual virtues was to recognise that one could disagree with a concept's usefulness but still appreciate its central role in an intellectual tradition. Natural fertility does not imply much about the number of children a woman has. It just means that she is indifferent as to the number of children she has. In some natural fertility populations, couples have, on average, seven or more children. This case contributes to the misunderstanding that natural fertility is high fertility. There are also natural fertility populations with not much more than two children per couple. This case arises

[6] Coale & Watkins (1978) summarises the famous 'Princeton project', which stresses the role of changing ideas. Guinnane (2011) summarises the economic arguments. Hirschman (1994) provides a comprehensive review.

when couples take measures that space births very widely. So long as the chance of the next birth does not depend on parity, even this extremely low fertility can be 'natural'.

Natural fertility plays an important role in debates about the fertility transition because most scholars view the transition as a shift from natural fertility to something else, and because, as a result, many important yardsticks for historical fertility focus on departures from natural fertility alone. The most straightforward departure from natural fertility reflects 'stopping' behaviour, which refers to a woman who ceases bearing children while still physically able. Until the onset of this stopping, the woman would appear to fit the idea of natural fertility. (Demographers tend to refer to 'stopping' as 'family limitation', which perhaps deepens the miscommunication, since stopping is not the only way to limit family size.) The polar opposite approach, 'spacing', refers to a situation where a woman has a small number of children whose births are widely spaced; this practice, as noted, still constitutes natural fertility. Most historical demographers believe that historical fertility control took the form of stopping behaviour. To some extent, this reflected the lesson of modern fertility surveys in developing countries, which demonstrated the priority of stopping over spacing. Paul also recognised that while 'stopping' clearly reflected an attempt to have a smaller family, apparently deliberate 'spacing' could be interpreted in alternative ways. Finally, many historical demographers doubted that birth-spacing to reduce family size could be achieved with rudimentary contraceptive methods.

The most widely-used methods for inferring deliberate fertility control in the historical record rely on comparing fertility in an actual population (the 'target' population of interest) to fertility in a population thought not to be practising contraception (the 'natural fertility' population). The best-known of these 'indirect' approaches is Coale-Trussell's '*M*' & '*m*' (Coale & Trussell 1978). Most scholars interpret '*m*' as reflecting the degree of departure from natural fertility in the target population. Paul's contribution to the measurement issue reflected an appreciation of how indirect methods work. The United States in 1900 added a question to the federal census of population, asking every currently-married woman her age at marriage, how many children she had borne, and how many currently survived. The United Kingdom (which at that time included Ireland) added a similar question to its census in 1911. Using either published tables or the later samples from the US census manuscripts, one could construct tables showing the complete parity distribution for each age-at-marriage and marital-duration bin. Coale-Trussell and similar approaches rely on more summary information, the age-specific marital fertility rates for the two populations.

Paul's Cohort Parity Analysis (CPA) shares the underlying notion of Coale-Trussell: it compares two populations, one thought to be natural fertility.[7] CPA rests on a simple

[7] The complete method underlying CPA appears in David *et al.* (1988). Other, related publications present parts of the model aimed at specific audiences or demonstrate its use for different empirical questions (David & Sanderson 1988; 1990).

notion. Consider the example of women who married at age 25–29 and have been married for 15–19 years. The parity distribution for the natural-fertility population reflects the underlying distribution of fecundity alone (fecundity of the couple; given the data, we can only observe the fecundity of the partners together). The parity distribution for the target population, on the other hand, reflects variations in fecundity as well as attempts to both space and stop births. Comparing the target to the natural-fertility population implicitly identifies the role of fecundity in the outcomes we observe among the target group. The remaining variation for the target population reflects a mixture of spacing, stopping, and couples who made no effort to control fertility. CPA provides upper- and lower-bound estimates for the fraction of each group who have initiated fertility-control measures.

One sees today a baffling tendency for scholars to assert that effective contraception started with the invention of the contraceptive pill, which became legal in the United States in 1960. If true, this would raise questions about the data discussed above. The idea probably reflects two failures of historical imagination. First, the rudimentary methods available to couples in the 19th century do not appeal to those who, today, have safer, more convenient options. Second, we have to consider the context, which is to say the goals of those making reproductive decisions at the time. Women use contraception today in part to time births around educational and professional events: between graduation from law school and the bar exam, for example. Women also use contraception when they do not want any children with a particular partner. The historical population of interest consisted of married women who largely lacked this concern over the tight timing of births. Nearly all presumably wanted at least some children. They just wanted to avoid having the eight or nine that they could otherwise expect.

Paul surveyed the (admittedly sparse) evidence on how 19th-century couples avoided pregnancies, and then used simulation methods to ask whether these methods could actually achieve the reduced fertility we observe around the turn of the 20th century (David & Sanderson 1986). The answer to the latter question is a decisive yes: there is no need to make claims about venereal disease and widespread sterility to understand the smaller families we see. The more subtle part of this work required understanding the limitations of these methods and how couples might rely on them. A (small and imperfect) survey of married US women starting in the 1890s showed that the condom, douching, and withdrawal (or *coitus interruptus*) were the staple methods. The data on coital frequency also shows that these couples had sexual intercourse less often than appears to be the case in modern surveys; this, too, might have reflected fertility goals. The vulcanisation of rubber had led to cheaper (but not cheap) rubber condoms and cervical barriers by the 1870s. Generations taught the virtues of condoms in an age of HIV-AIDs will not be shocked by the role of condoms, but most sex-education courses today strive to convince young people that douching and withdrawal are so ineffective as to be useless. Here is the core of Paul's insight: the 'per-trial failure rate' associated with most rudimentary methods seems appallingly high to generations who have ready access to better methods, but, used consistently over a lifetime, even contraceptive methods with these failure rates can result in family sizes much smaller than those that would come from a regime of no contraception. This result relies on a simple model and medical parameters that may not be perfect for the historical population in question. But they make a wonderfully simple and important point. The model implies that a woman would bear 9.4 children in her lifetime if she used no contraception. If she used a method that had a per-trial failure rate of even 25 per cent, she could expect 5.3 children. This failure rate is consistent with observed studies of couples who relied on douching, withdrawal, or condoms. More careful use of any of these methods would imply 3.35 births. Rudimentary methods can fully account for the observed fertility decline.[8]

This research led to another insight. Reproductive medicine in the later 19th century agreed that there was a 'safe period' during every woman's menstrual cycle, and some advice turned on trying to avoid sexual intercourse during the other times. Unfortunately, medical knowledge concerning the timing of a safe period was wrong-headed until the 1930s; women trying to practise a 'rhythm method' based on this false information would find themselves with large families (David 2013). One might think that this misinformation would frustrate any attempt at relying on systematic periodic abstinence. Paul noted that this was not actually the case. Indulging his taste for cute titles, he developed this insight into a model of what he called 'the Bayesian adaptive rhythm (BAR) method'. This is, as the title suggests, a Bayesian model. The core idea is that a woman starts her reproductive life with the wholly incorrect prior about her 'safe period'. She maintains the idea that there is a safe period but learns more through trial-and-error. Each 'trial' consists of a month of sexual relations timed to the days when she thinks, given her beliefs, she is least likely to become pregnant. Women who become pregnant this way adjust their priors. Those who do not, maintain their beliefs (an example of the Bayesian 'stick with the winner' principle). Simulations show that BAR could reduce births considerably relative to a benchmark in which the woman uses no contraception.[9]

The research on contraceptive methods ties back to the question of whether historical fertility transitions relied on spacing or stopping. Research independent of Paul and his co-authors identified cases where spacing was empirically important. Paul's CPA results

[8] See David and Sanderson (1986: Table 7.8). The outcomes depend on the model's assumptions about the number of times the woman has intercourse in each 24-day period. The figures in the text are for 6 times per cycle, which is higher than implied by the historical evidence the paper discusses.

[9] See Table 5.1. The model is similar to that used in David & Sanderson (1986). To take the example of a woman who has intercourse 6 times per cycle: without BAR, she would have 10.2 children, and with BAR, 5.52.

for Ireland also indicate that many Irish birth-controllers were spacers. These results met with considerable scepticism. When faced with obvious spacing (that is, unusually long birth intervals), some demographers reply that yes, those are long intervals, but they reflect some other goal, not the desire to have a smaller family. This 'real reason' interpretation reflects a basic difference between economics and other social sciences. Economic models of the demand for children incorporate a range of costs and benefits, viewing the ultimate demand as reflecting the balance of those two forces. For example, extended breastfeeding is a mild contraceptive and also improves child health. To some historical demographers, long intervals show a desire for healthier children, not fewer children. (Put differently, since one cannot really know why birth intervals are long, the argument goes, it is not acceptable to assume they reflect fertility desires.) In an economic framework, on the other hand, couples consider the trade-offs; they understand that extended breastfeeding probably means fewer children.

A second source of scepticism about historical spacing reflects a different understanding of the implications of rudimentary fertility control. Common sense might tell us that reliance on a method with a high failure rate would make it hard to space births according to some pre-determined plan. The notion here is that birth-spacing requires each birth to occur on a precise schedule. However, as Paul showed (see below), a simple model shows that high failure rates make birth-spacing more attractive than stopping. Again, what matters is the goal. Couples were not trying to have a child in every third January; rather, they were trying to reduce the size of their total brood. Consider a couple who married at the wife's age 24 and wanted to have only three children. If they relied on stopping alone, they might well find themselves with these four children when the wife was as young as 33. What then? If they really want to have no more children, and the methods available to them had high failure rates, they might well face years of complete sexual abstinence, until the wife was no longer capable of pregnancy. A better plan would be to practise 'careful love' from the start, seeking to reduce the chance of pregnancy in each month, but not attempt to set it to zero. Spacing in this way allows adjustment to unexpected pregnancies and does not imply the need for complete abstinence later in life.

Paul's research thus implied that fertility transitions might well have relied on spacing to a serious degree. But one can acknowledge an identification problem: how do we know, empirically, that couples who space their births do not have some fertility problem that means their fertility does not reflect deliberate decisions? (Parity-dependent control, on the other hand is, with the right data, hard to confuse with anything else.) There are ways to address this problem, and they boil down to asking whether couples in an apparent natural-fertility population react to changes that would not matter to them if they were not trying to control their final brood size. More formally, suppose one has the kind of birth-history data that can be generated from family reconstitution studies. If a couple is not practising contraception, then the hazard rate for the next live birth will not depend on the current number and characteristics of surviving children; natural fertility implies that couples are indifferent over the number of surviving children they have (after conditioning, perhaps on the age at which they marry).

This insight leads to some straightforward empirical tests. Paul and his co-authors were not the first to suggest this approach, but it forms a natural outgrowth of his other research on fertility. David & Mroz (1989a; 1989b) develop a simple dynamic model of family-building. The model's control variable is the hazard rate for a live birth; the couple selects the contraceptive intensity that maximises lifetime welfare. The model generates several possible falsifications of the natural-fertility hypothesis. The simplest is the so-called 'replacement effect': if the probability of a live birth rises when a child dies, it looks like an attempt to compensate by having extra children (this approach requires, of course, econometric methods that seek to remove any direct effect of a child's death of the mother's fertility and must bear in mind the possible endogeneity of child death). More subtle examples avoid the complications that arise from the mortality channel. For example, if it appears that a couple has preferences over the sex of their children, then this, too, violates the natural fertility hypothesis. While this contribution is mostly methodological, it concludes that there was significant deliberate birth-spacing in pre-Revolutionary France.

The institutions of science and technology

In the 1990s, Paul shifted his geographic center, dividing his appointment between Stanford and All Souls College, Oxford, but spending the majority of his time in Europe. In the process, he moved on to even larger conceptual and historical terrain: the nature and origin of research communities and knowledge generation in science and in technology. In early articles co-authored by Partha Dasgupta, the authors argued that scientific research serves to increase the stock of knowledge, while technological 'R&D' is devoted to earning rents from that stock (1987, 1994). The two organisational modes thus have a complementary relationship, but very different structures and norms of behaviour. Specifically, the scientific communities stress *priority* of discoveries in assigning rewards, thereby fostering the norm of *disclosure*, which in turn facilitates rapid verification or correction. These norms of 'open science' contrast with the emphasis on proprietary knowledge and secrecy in commercially oriented research.

The preceding analysis, and others like it, were characterised by Paul as 'functionalist' accounts, explaining an institutional arrangement in terms of its efficacy in mitigating market-failure problems in the real world. Ever the historian, Paul objected that these rationalisations amounted to 'fairy tales'. Did a group of wise men convene on an appointed day, analyse the structure of the problem, and prescribe an institutional remedy? Clearly not. Until we grasp the 'historical origins' of open science, Paul argued, claims for its 'logical origins' will not be adequate for understanding its historical role or its status in the present day. Devoting nearly twenty years of research and academic interchange to this topic, Paul offered a new analytical narrative: what we know as 'open' or 'public' science grew out of the post-Renaissance European patronage system, when kings and nobles competed for sponsorship of philosophers, musicians, artists – and the new breed of scientists – both for 'ornamental' and utilitarian reasons. As the new science became more mathematical in the 17th and 18th centuries, patrons were impelled to draw upon internationally renowned figures to confirm the stature and verify the claims of candidates. In time, this process of competition and 'common agency contracting' gave rise to the cultural ethos associated with the 'Republic of Science', European feudalism's great gift to the economic vigour of capitalism in the modern age (David 1998).

The open science project grew initially out of Paul's largely theoretical work with Dasgupta but, as it unfolded, drew upon seminars and conversations across two continents and multiple disciplines, appropriate for the subject matter: economists, economic historians, historians of science and technology, and sociologists. Once more, book-length treatment was promised but never completed. Once again, the lengthy manuscript was condensed into an abbreviated version for the *American Economic Review* (1998). Although that brief statement did not gain the attention of the two earlier iconic papers, it paved the way for the fullest elaboration of the case, the hundred-page article published in the new journal *Capitalism and Society* in 2008. The essay is not primarily cliometric, though it does contain two detailed tables compiling the involvement in practical arts by scientists and mathematicians in the 16th and 17th centuries. But it is historically learned yet framed in terms recognisable to economists: principal-agent problems, common agency games, etc. It draws upon repeated game theory to show that the right set of rewards and punishments can elicit 'cooperative behavior among potentially rivalrous researchers' (2008: 72).

Although an impressive accomplishment in its own right, 'open science' was also a springboard into Paul's engagement with science and technology policy on both sides of the Atlantic. His historical agenda was unabashedly 'presentist', an effort to 'offer insights from the past that carry implications for the present and the future' (2008: 8). For Paul, the upshot of the history is that the institutions of open science are a 'fragile bequest from the past', increasingly vulnerable to encroachment from the realm of commercial R&D. One channel for such intrusion is direct engagement of faculty members in privately funded research, jeopardising their willingness or capacity to release data or research results promptly. A second channel is the extension of proprietary motives to the universities themselves, claiming property rights in research results generated on campus facilities. Writing for a European audience, Paul warned against emulation of American practices emerging from the Bayh-Dole Act of 1980, which have been shown to increase delays in public access to new research results (David 2004; 2007b). Paul's larger concern, clearly, was that obliterating the distinction between the realm of public science and the realm of private technology would diminish the scope for the exploratory scientific research that has played such a large role in human advancement over the centuries.

Paul's outlook was not altogether pessimistic. His preferred response to the impending 'tragedy of the anticommons' was construction of 'research commons' in the form of common-use licensing of intellectual property by researchers, initiated on a 'bottom-up' basis but with encouragement from public officials and funding agencies. Paul was an active participant in many such efforts, including Science Common, launched as a project of Creative Commons in 2005, with the goal of bringing to the world of scientific endeavours the benefits of openness and sharing that have made Creative Commons licences a success in the arts and cultural fields. Paul served as a member of the Scientific Advisory Board of Science Commons (David 2010).

The legacy of Paul A. David

Few scholars could hope to match the range and breadth of Paul's accomplishments and influence. Although he might have reached an even larger readership if his book projects had been completed, Paul's impulse to move on to new and ever-more ambitious projects was irrepressible and may account for the steady stream of creative and original reformulations that form his academic legacy. The concept of path dependence and the QWERTY metaphor will long be associated with Paul's name. His formulations of diffusion and other dimensions of technological innovation have left an indelible mark on that field.

Also inspiring is the fact that, even while moving on to new intellectual terrains, Paul saw connections among virtually all branches of his work. Rudimentary contraceptive practices constituted a technology, and learning from experience with those practices was similar to the process of adaptation to a new technology, albeit with noisy flow of information. Perhaps most dramatically, Paul viewed the institutions and culture of Open Science as a powerful illustration of path dependence, the lasting influence of the past on subsequent historical events. The same example also illustrates that path dependence does not imply that systems are locked in to particular outcomes for the rest of time. Historically-driven structures can be fragile and quickly undone. Paul's message was and continues to be that deeper understanding of that history can help us to cope with contemporary problems and policies.

References

Abramovitz, M. (2001), 'Days gone by: A memoir for my family'. https://web.archive.org/web/20070609121742/http://www-econ.stanford.edu/abramovitz/abramovitzm.html

Abramovitz, M. & David, P.A. (1973a), 'Reinterpreting Economic Growth: Parables and Realities', *American Economic Review*, 63: 428–39.

Abramovitz, M. & David, P.A. (1973b), 'Economic Growth in America: Historical Parables and Realities', *De Economist*, 121: 251–72.

Abramovitz, M. & David, P.A. (1996), 'Convergence and deferred catch-up: productivity leadership and the waning of American exceptionalism', in Landau, R., Taylor, T. & Wrights, G. (eds), *The Mosaic of Economic Growth* (Stanford CA, Stanford University Press).

Abramovitz, M. & David, P.A. (2000), 'American macroeconomic growth in the era of knowledge-based progress: The long-run perspective', in Engerman, S.L. & Gallman, R.E. (eds), *The Cambridge Economic History of the United States, Volume III, The Twentieth Century* (Cambridge UK, Cambridge University Press), 1–92.

Arthur, W.B. (1989), 'Competing technologies and lock-in by historical small events', *Economic Journal*, 99: 116–31.

Brady, D.S. (ed.) (1966), *Output, Employment and Productivity in the United States after 1800* (Studies in Income and Wealth, National Bureau of Economic Research, 30; New York, Columbia University Press).

Coale, A.J., & Trussell, T.J. (1978). 'Finding the two parameters that specify a model schedule of marital fertility', *Population index*, 203–13.

Coale, A.J., & Watkins, S.C. (1978). *The Decline of Fertility in Europe* (Princeton, Princeton University Press).

Cohen, J., Dickens, W. & Posen, A., 'Have the new human resource management practices lowered the sustainable unemployment rate?', in Krueger, A.B. & Solow, R.M. (eds), *The Roaring Nineties* (New York, Russell Sage Foundation).

Dasgupta, P. & David, P.A. (1987), 'Information disclosure and the economics of science and technology', in Feiwel, G. (ed.), *Arrow and the Ascent of Modern Economic Theory* (New York, New York University Press).

Dasgupta, P. & David, P.A. (1994), 'Towards a new economics of science', *Research Policy*, 23: 487–521.

David, H. (1936), *The history of the Haymarket affair: A study in the American social revolutionary and labor movements* (New York, Farrar & Rinehart Inc.).

David, P.A. (1966), 'The mechanization of reaping in the antebellum Midwest', in Rosovsky, H. (ed.), *Industrialization in Two Systems* (New York, John Wiley and Sons), 3–39.

David, P.A. (1967a), 'The Growth of Real Product in the United States before 1840: new evidence, controlled conjectures', *Journal of Economic History*, 27: 151–97.

David, P.A. (1967b), 'New light on a statistical dark age: U.S. Real Product growth before 1840', *American Economic Review*, 55: 294–306.

David, P.A. (1971), 'The landscape and the machine: Technical interrelatedness, land tenure and the mechanization of the corn harvest in Victorian Britain', in McCloskey, D.N. (ed.), *Essays on a Mature Economy: Britain after 1840* (London, Methuen), 145–214.

David, P.A. (1975), *Technical Choice, Innovation and Economic Growth: Essays on American and British experience in the nineteenth century* (New York, Cambridge University Press).

David, P.A. (1977), 'Invention and Accumulation in America's Economic Growth: a nineteenth-century parable', in Meltzer, K. & Meltzer, A.H. (eds), *International Organization, National Policies and Economic Development* (Carnegie-Rochester Conference Series on Public Policy; Amsterdam, North-Holland).

David, P.A. (1985), 'Clio and the economics of QWERTY', *American Economic Review*, 75: 322–7.

David, P.A. (1986), Understanding the economics of QWERTY: The necessity of history', in Parker, W.N. (ed.), *Economic History and the Modern Economist* (London, Basil Blackwell).

David, P.A. (1990), 'The dynamo and the computer: An historical perspective on the modern productivity paradox', *American Historical Review*, 80 (May 1990): 355–61.

David, P.A. (1991), 'Computer and dynamo: The modern productivity paradox in a not-too-distant historical mirror', in Organisation for Economic Cooperation and Development, *Technology and Productivity: The Challenge for Economic Policy* (Washington DC, OECD Publication and Information Center).

David, P.A. (1994), 'Why are institutions the "carriers of history"? Path dependence and the evolution of conventions, organizations and institutions', *Structural Change and Economic Dynamics*, 5: 205–20.

David, P.A. (1998), 'Common agency contracting and the emergence of "open science" institutions', *American Economic Review*, 88 (May): 15–21.

David, P.A. (2004), 'Can "Open Science" be protected from the evolving regime of IPR protections?', *Journal of Institutional and Theoretical Economics*, 160 (March): 9–34.

David, P.A. (2007a), 'Path dependence: a foundational concept for historical social science', *Cliometrica*, 1: 91–114.

David, P.A. (2007b), 'New science, new industry and new institutions? Second thoughts on innovation and Europe's universities', in Malerba, F. & Brusoni, S. (eds), *Perspectives on Innovation* (Cambridge, Cambridge University Press), 151–73.

David, P.A. (2008), 'The historical origins of "open science": an essay on patronage, reputation and common agency contracting in the Scientific Revolution', *Capitalism and Society*, 2: 1–103.

David, P.A. (2010), 'Mitigating "anticommons" harms to science and technology research', *The AIPO Journal: Analysis of Intellectual Property Issues*, 1: 59–73.

David, P.A. (2013), 'Path-dependent learning, and the evolution of beliefs and behaviors: implications of Bayesian adaptation under computationally bounded rationality', in *The evolution of economic diversity* (Routledge), 93–140.

David, P.A. & Bunn, J.A. (1988), 'The economics of gateway technologies and network evolution: Lessons from electricity supply history', *Information Economics and Policy*, 3: 165–202.

David, P.A., Mroz, T.A., Sanderson, W.C., Wachter, K.W. & Weir, D.R. (1988), 'Cohort parity analysis: Statistical estimates of the extent of fertility control', *Demography*, 25: 163–88.

David, P.A., & Mroz, T.A. (1989a), 'Evidence of fertility regulation among rural French villagers, 1749–1789: A sequential econometric model of birth-spacing behavior (part 1)', *European Journal of Population/ Revue Européenne de Démographie*, 1–26.

David, P.A., & Mroz, T.A. (1989b), 'Evidence of fertility regulation among rural French villagers, 1749–1789: A sequential econometric model of birth-spacing behavior (part 2)', *European Journal of Population/ Revue Européenne de Démographie*, 173–206.

David, P., & Sanderson, W. (1986), 'Rudimentary contraceptive methods and the American transition to marital fertility control, 1855–1915', in *Long-term factors in American economic growth* (University of Chicago Press), 307–390.

David, P.A. & Sanderson, W.C. (1987), 'The emergence of a two-child norm among American birth-controllers', *Population and Development Review*, 1–41.

David, P.A. & Sanderson, W.C. (1988), 'Measuring marital fertility control with CPA', *Population Index*, 691–713.

David, P.A. & Sanderson, W.C. (1990), 'Cohort parity analysis and fertility transition dynamics: Reconstructing historical trends in fertility control from a single census', *Population Studies*, 44(3): 421–45.

David, P.A. & Wright, G. (1997), 'Increasing Returns and the Genesis of American Resource Abundance', *Industrial and Corporate Change*, 6: 203–45.

David, P.A. & Wright, G. (2003), 'General purpose technologies and surges in productivity: Historical reflections on the future of the ICT revolution', in David, P.A. & Thomas, M. (eds), *The Economic Future in Historical Perspective* (British Academy Centenary Monograph; Oxford, Oxford University Press for the British Academy), 135–66.

Dawidoff, N. (2002), *The Fly Swatter: how my grandfather made his way in the* world (New York, Pantheon Books).

Guinnane, T.W. (2011), 'The historical fertility transition: A guide for economists', *Journal of economic literature*, 49(3): 589–614.

Hirschman, C. (1994), 'Why fertility changes', *Annual review of sociology*, 20(1): 203–33.

Kirsch, D. (2000), *The electric vehicle and the burden of history* (New Brunswick NJ, Rutgers University Press).

Liebowitz, S.J. & Margolis, S.E. (1990), 'The Fable of the Keys', *Journal of Law and Economics*, 33: 1–25.

Liebowitz, S.J. & Margolis, S.E. (1994), 'Network externality: An uncommon tragedy', *Journal of Economic Perspectives*, 8: 133–50.

Olmstead, A.L. (1975), 'The mechanization of reaping and mowing in American agriculture, 1833–1870', *Journal of Economic History*, 36: 399–415.

Olmstead, A.L. & Rhode, P.W. (1995), 'Beyond the threshold: An analysis of the characteristics and behavior of early reaper adopters', *Journal of Economic History*, 55: 25–57.

Olmstead, A.L. & Rhode, P.W. (2011), 'Adapting North American wheat production to climatic challenges, 1839–2009', *Proceedings of the National Academy of Sciences of the United States of America*, 480–85.

Parker, W.N. (ed.) (1960), *Trends in the American Economy in the Nineteenth Century*, Studies in Income and Wealth, National Bureau of Economic Research, Vol. 24 (Princeton, Princeton University Press).

Puffert, D. (2009), *Tracks across continents, paths through history: the economic dynamics of standardization in railway gauge* (Chicago, University of Chicago Press).

Wright, G. (1990), 'The Origins of American Industrial Success, 1879–1940', *American Economic Review*, 80: 651–68.

Note on the authors: Timothy W. Guinnane is Philip Golden Bartlett Professor of Economic History, Emeritus, Yale University. Gavin Wright is William Robertson Coe Professor of American Economic History, Emeritus, Stanford University.

Biographical Memoirs of Fellows of the British Academy (ISSN 2753–6777) are published by
The British Academy, 10–11 Carlton House Terrace, London, SW1Y 5AH
www.thebritishacademy.ac.uk