



HEALTH CARE IN THE INFORMATION SOCIETY

VOL. 1

FROM ADVENTURE OF IDEAS TO
ANARCHY OF TRANSITION

DAVID INGRAM



<https://www.openbookpublishers.com>

©2023 David Ingram



This work is licensed under an Attribution-NonCommercial 4.0 International (CC BY-NC 4.0). This license allows you to share, copy, distribute and transmit the text; to adapt the text for non-commercial purposes of the text providing attribution is made to the authors (but not in any way that suggests that they endorse you or your use of the work).

Attribution should include the following information:

David Ingram, *Health Care in the Information Society, Vol. 1: From Adventure of Ideas to Anarchy of Transition*. Cambridge, UK: Open Book Publishers, 2023,
<https://doi.org/10.11647/OBP.0335>

Copyright and permissions for the reuse of this image is provided in the caption and in the list of illustrations. Every effort has been made to identify and contact copyright holders and any omission or error will be corrected if notification is made to the publisher.

Further details about the CC BY-NC license are available at
<http://creativecommons.org/licenses/by-nc/4.0/>

All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at <https://archive.org/web>

Digital material and resources associated with this volume are available at
<https://doi.org/10.11647/OBP.0335#resources>

ISBN Paperback: 978-1-80064-952-1

ISBN Hardback: 978-1-80064-953-8

ISBN Digital (PDF): 978-1-80064-954-5

ISBN Digital ebook (EPUB): 978-1-80064-955-2

ISBN XML: 978-1-80064-957-6

ISBN HTML: 978-1-80064-958-3

DOI: 10.11647/OBP.0335

Cover image: D koi, *White H₂O Molecules* (27 June 2022),
<https://unsplash.com/photos/5nI9N2wNcBU>

Cover design: Jeevanjot Kaur Nagpal

4. Models and Simulations— The Third Arm of Science

Modelling and simulation have arisen as a third branch of science alongside theory and experiment, enabling and supporting discovery, insight, prediction and action. The Information Age gave rise to an upsurge in the use of models to represent, rationalize and reason about measured and predicted appearances of the real world. This chapter describes different kinds of model—physical, mathematical, computational—and their use in different domains and for different purposes. Solutions of mathematical model equations that defied analytical method and required huge amounts of mental and manual effort for the calculations made, before the computer, became considerably more straightforward to deal with using computational methods and tools developed and refined in the Information Age.

In the examples described, the focus is on pioneers I have been taught by, got to know or collaborated with: John Houghton (1931–2020) on weather and climate modelling, to give a perspective from a non-medical domain; Arthur Guyton (1919–2003) and John Dickinson (1927–2015) on modelling of body systems and clinical physiology; Louis Sheppard on model-based control systems for intensive care, and mathematical models applied to track and predict the course of epidemics and analyze clinical decisions. Other examples are from teams I have been privileged to see firsthand, as a reviewer and advisory board chair of largescale research projects across the European Union.

With colleagues in the UK and Canada, I previously published the Mac Series models of clinical physiology with Oxford University Press. I have established a Cloud-based emulation environment to provide access to these working models—created in the first half of my career and thus now archaic in terms of software interface—to accompany their description in one of the chapter's examples.

The most conspicuous example of truth and falsehood arises in the comparison of existences in the mode of possibility with existences in the mode of actuality.

–Alfred North Whitehead (1861–1947)¹

Science may be described as the art of systematic over-simplification—the art of discerning what we may with advantage omit.

–Karl Popper (1902–94)²

The story now moves from knowledge, observation and measurement to modelling and simulation, which are sometimes described as the third arm of scientific method, alongside theory and experiment. These in turn connect with information and engineering, which is where the story moves to in the following chapter. We build models to see how our ideas hang together and might pan out, when observed or implemented for real in everyday practice and context. Simulation is the enactment of a model, and science seeks new understanding by hypothesizing and exploring the enactment of candidate theoretical models, matched against experimental data.

There are many kinds of model in everyday use—in physical or abstract logical, mathematical or computational forms. Some are continuously grounded in experiment and measurement, and closely track observed physical reality. Some synthesize and reason with accumulated knowledge spanning different disciplines and domains. The knowledge bases introduced in Chapter Two embody models of domains of knowledge.

Alfred North Whitehead's words can be taken as a caution not to confuse the real world and the model, in this process. Models can be powerful tools, but it is beguiling easy to become a bit too fixated on them. Karl Popper's remark may be taken as another kind of caution; about how we delineate the purpose that a model is to serve and reduce to its essence the detail that we include, so that the model can be both tractable and useful. Too many particular adaptations, and the generality of the model loses its power and appeal. Too much blanket generalization and the real world loses touch. These issues have come to the fore in how models are designed, validated, communicated and used. In 2020, mathematical epidemiologists built and refined many mathematical models to guide government policies, seeking to manage and combat the future progression of the Covid-19 pandemic. In the first half of my academic career, from 1970–89, I was closely involved in

1 A. N. Whitehead, *Adventures of Ideas* (New York: Macmillan, 1933), p. 234.

2 K. Popper, *The Open Universe: An Argument for Indeterminism from the Postscript to the Logic of Scientific Discovery*, ed. by W. W. Bartley III (Abingdon: Routledge, 2012), p. 44.

the development and application of computer models of clinical physiology and pharmacology; it is a world I have known well.

The Egyptians built pyramids using a toolbox of methods, combining astronomical observation, standardized measurement, machinery and human labour. They might have built a small-scale model of the Grand Pyramid to test their ideas, or maybe they just went for broke and got lucky in attaining its summit. They had built smaller pyramids first, and if you have seen them, their summits do look a bit flatter—rather like the summit of Mont Blanc. They are not at all as sharply geometrical as the Grand Pyramid. Maybe a point summit was not aimed for or just not attained for these. Maybe they reflect erosion through time of less hardy construction materials. Or maybe they reflect gradual improvement of initially not so successful pyramid construction methods, through successive attempts to build them, and learning from the experience.

This speculation serves to introduce an important general point: that there is iterative learning involved in making and deploying useful models, and risk in extrapolating their use, precipitately or too widely, in real life. That said, innovators always push on and pursue adventures of ideas, beyond the bounds of what may, in retrospect, look to have been safer and more logical ways of proceeding into the unknown. There are always first prototypes, and these may not work—they are models used in developing and improving designs. John Archibald Wheeler (1911–2008) has a memorable quotation in his seminal ‘it from bit’ paper that I introduced in Chapter Three, where he writes that the ‘The policy of the engine inventor, John Kris, reassures us, [about the importance of testing our ideas in practice] “Start her up and see why she won’t go”’³ Of course, medical practice must be more cautious! But, medical science and clinical practice are always proceeding into the unknown, both at the level of first encounter with a presenting patient, or in exploring in further detail, with new methods, a presenting situation that goes beyond or does not quite fit with current knowledge and practice.

The terms modelling and simulation arise widely in everyday life: mathematicians, scientists, engineers, clinicians, economists, businesspeople and politicians all create, use and talk about their models. Philosophers of mind talk of it as encompassing a model of the world it inhabits. At Magdalen College (University of Oxford), I lived for a year two floors above the study of Gilbert Ryle (1900–76), who devoted his life’s work to philosophy of mind. He has been succeeded by luminary figures like Daniel Dennett. Philosophy of mind, and ideas about the nature of consciousness and intelligence,

3 J. A. Wheeler, ‘Information, Physics, Quantum: The Search for Links’, in *Feynman and Computation*, ed. by A. Hey (Boca Raton, FL: CRC Press, 2018), pp. 309–36 (p. 310), <https://doi.org/10.1201/9780429500459-19>

weave in and out of contemporary neuroscience, psychology, economics, computational science and engineering, and artificial intelligence (AI).

A seven-foot-tall machine called MONIAC (Monetary National Income Analogue Computer—the name amusingly close to the root of another word that seems suitably adjectivally-descriptive of the world it models!) was one of the first economic models, built in 1949 at the London School of Economics. I have seen it on display at the London Science Museum, consisting of a collection of tanks and tubes designed to simulate the flow of money around the economy. The human circulatory system has variously been modelled as a hydraulic mechanism, a set of biomechanical equations, and as analogue and digital computer programs. Embodied in simulations, models are used to explore structure and function of the modelled system in ways that may not be open to direct experimentation—such as to think about what will happen when a currency is devalued, or a patient's cardiac pump performance suddenly decreases by a half. Or in an imaginary enactment of what has been described as happening in 'The First Three Minutes' of the universe, remembering the graphic storytelling in the book by Steven Weinberg (1933–2021).⁴ Not much scope for experiment there, but he was a co-recipient of the 1979 Nobel Prize in Physics for his outstanding range of theoretical contributions to more testable ideas, as noted in Chapter Three.

Early simulations posed new requirements for analytical solutions of complex mathematical equations. At the start of my career, simulations of analytically intractable equations, such as those describing the propagation of the nerve action potential, were cranked out, iteratively and laboriously, on hand-operated mechanical calculators, over many months, as discussed in the previous chapter. Now, intrinsically more complex, large scale and multilevel model equations require numerical methods for their solution, which are provided in computer software.

Models of many different kinds are used in research, education and training, and in the design, implementation and operation of devices and systems. They are used in guiding policy and managing organizations, large and small, spanning from the logistics of product manufacture and distribution to the performance of national economies, and the spread of international pandemics. This wide range of usage is informative of the different purposes served, and the methods used to build, test and apply the models to simulate the domains they encompass. The examples in this chapter are mainly focused on medical science and health care.

In managing the Covid crisis, models have played a strong supportive role. It has been reported that some seven different mathematical models

4 S. Weinberg, *The First Three Minutes* (New York: Bantam Books, 1979).

have been used in the United Kingdom (UK), to predict and reason about likely patterns of infection and morbidity, and their mitigation. Probably many more than that, I would imagine; nowadays, they tend to be quite easy to conjure into existence but often remain very difficult to tune to practical ends and advantage. Different models, based on different assumptions, differently structured and differently interpreted, have predicted strikingly different courses of events. Their design and application are theoretical, experimental and applied sciences. Politics may feel a need to cover its tracks, to be seen as guided by (clear and accepted) science. It is less attractive to admit dependency and reliance on the appliance of (unclear and uncertain) research in progress. Although there may indeed be some wisdom in deferring to a crowd of models in this situation, if any model is relied upon too greatly, or used incautiously, the conclusions drawn may mislead and distort reason, and perhaps distract unduly from simpler approaches that might play out as well, or better.

Second-guessing the behavioural impact of different phases of population lockdown to contain spread of infection has been a difficult area in which to weigh model prediction against the gut instincts of political judgement. The impact of the viral mutation that set the infection in the UK on a markedly different trajectory around Christmas 2020 was a Black Swan event that could not have been predicted, although the possibility of such an event was clearly always there. Viruses mutate and this mutation turned out to be highly impactful. The wild card this morning, as I write, has been a sudden and unexpected announcement of a significant reduction in vaccine supply from a manufacturer. Again, always a possibility, but a Black Swan occurrence defying prediction, other than in general assessment of potential risks, their impact, and mitigation strategies. It is costly to allow for the many such potential contingencies, where none or several might arise and would, in combination, interact. Chance events can be a both lucky and brutal amplifier and leveller. Some we are born to, and some are cast upon us.

All this is an extremely complex domain to seek to model! Some of the mathematical modelling flown into the stormy Covid pandemic has not flown too well. Policy choices in managing the crisis, predicated on modelling studies, were impactful in the spread and impact of the infection and in the economic consequences—many lives and livelihoods were at stake. It would be interesting to discover how policy has played out in places where there was no modelling community to draw on for advice. The many imponderables about the modelling work, and how the issues it sought to predict played out in coping with and managing the pandemic, will be the subject of important and critical review, as the current crisis recedes into the past. It has been a real-life (substantially uncontrolled) experiment

that none, other than the most centrally controlled of societies, could have conducted, or wished to do so!

There is a wider problem, here. Gaining a handle on managing complex systems often looks a very good candidate for a model-based approach. In the social domain—and managing a pandemic is as much a behavioural challenge as a scientific, clinical and logistical one—models are nigh on impossible to validate experimentally, in their real-life context. In science and engineering, it is of the essence to propose and dispute alternative models as a means for gaining insight and testing ideas. It seems unlikely that the designers of simulators for training aircraft pilots would allow multiple alternative models into everyday use. They would be more specific about the purposes served and provenance of the models proposed, and test them rigorously, before deciding on use. There was no such option in reacting to the crisis of pandemic. The modelling of world economies was subjected to a not dissimilar ‘experiment’, with the financial near-collapse of the world monetary system in 2007–08.

We now have models of complex global weather systems, used to forecast the local weather that we can expect, with considerable precision. They have taken many decades to improve and scale, in balance with the feasible measurement of the weather systems they represent, the logistics of continuous collection of the data needed, and the computational capacity required to simulate and communicate the weather patterns predicted by the solution of the model equations. And these forecasts need to balance with and be tuned to the needs of different audiences—citizens, event organizers, local farmers and airline pilots.

If we only consult the weather App on our smartphone, we may yet step out into a vigorous local downpour. Somewhere along the line, individual judgement receives and processes multiple sources of information, balances risks and decides what to do—whether to stay home, what clothes to wear if going out and whether always to carry an umbrella, just in case. And it is a good idea to have a look outside the front door, first. In well-defined and monitored domains, machine intelligence may, in time, prove decisive in such judgements. If so, we must accept the loss of autonomy and self-reliant capacity to observe and think, that this dependency may engender or entail. There is an element of bargaining involved—we bargain on it not raining and leave our umbrella at home. With the Mephistophelean computer involved, it might prove a Faustian bargain.

Purpose and Method

The term model is now very widely appropriated. As a noun, it carries a connotation of ideal form, to be adhered to or followed as a way of presenting or acting. As a verb, it is a creative action, abstracting and representing something—an aeroplane, chemical structure, or body system. We talk of human role models—people we observe and follow, and thereby feel helped to shape and improve our own roles and contributions in life. Good role models are not perfect people, but they are authentic, demonstrating quality and balance in what they achieve, and how they do it. They are, in a sense, representations of who we might aim to be and become. We need to see and feel some connection between them and their lives, and us and ours, and to have belief and trust in them. The pursuit of better things needs always to be balanced by a sense of what we have, being perhaps good enough already. As the paediatrician Donald Winnicott (1896–1971), a close colleague of my parents in the 1950s, used to write, parents should not aim at perfect parenting—good enough parenting is a good enough goal to aim for. There is a wider issue, here—how do we decide what is good enough in the models we create and adopt?

Models are pervasive: in research, education and training; in practical support of design, manufacture and operation of devices; in systems and services of everyday life. They are integral to control engineering. They assist in reasoning about complexity, exploring consistency of theory and experiment, and discerning and focusing on relevant and essential detail. Purpose of use and method of design and operation connect closely. We seek good enough models.

For Science

Mathematical models and simulations have featured in theoretical physics since long before my songline. They employ differential equations to describe physical phenomena that evolve over time, solve the equations analytically, and present the results as columns of numbers and graphs. The integration of the equations is a simulation of the model they represent. Many such representations are of a scale and complexity that defies such analytical solution, and simplifications are adopted for the purpose of making analytical solutions more tractable. Non-linear relationships are linearized, for example, and large scale and distributed systems are divided into a set of smaller equivalent compartments.

Such models now range from the atomic to cosmic dimensions of physics, and living organisms and systems are modelled from molecular

to earthly (Gaia) dimensions. Discussion of the nature of consciousness and intelligence has been known to scale theoretical models, conceptually, towards the current Planck boundaries of measurement! And the quest to create new mathematical methods that better serve the modelling of complex physical and biological systems goes on. Mathematics has long sought to go beyond the integer realm of calculus, where we have the first, second and higher order derivatives of traditional calculus, and their integrals, to explore what a fractional calculus might look like, lying between those integer orders. New tools have emerged to populate models with these functions and thereby achieve a more faithful and tractable representation of complex system dynamics, for example where they cross over between different domains of behaviour and require memory of past behaviour when predicting how the future will unfold. These methods are finding increasing application in science, including the disciplines of biology and medicine, for both modelling and controlling system behaviour.

In physics, models range over field theory and particle physics of the exceedingly small, to cosmology of the exceedingly large. They simulate and match experimental data from particle colliders that accelerate matter towards the speed of light, and laboratory-, land-, ocean- and satellite-based environmental sensors, and telescopes probing to the limits of the observable physical universe. In more recent decades, models ranging from the atomic level, to the molecular, cellular, organ, body and population levels, have spread through life science, medicine and health care.

I had early encounter with the theoretical foundations of atmospheric physics that underpin today's weather forecasts, and use this as my first example, below. Coming from a field far removed from biomedicine, the story of their evolution over time provides useful counterpoint to later discussion of models in life science and health care. Rather as the story of the evolution of approaches to library classifications of knowledge, in Chapter Two, showed parallels with that of computerized medical terminologies, classifications, and knowledge bases. In like manner, stories about engineering innovation in the Industrial Revolution are invoked, in Chapter Five, when surveying the engineering of information systems in the Information Revolution of our age. I still follow with interest, but sadly no longer as much understanding, the model-related discussions, and conjectures about unification of quantum theory and general relativity. Today, as I write (18 January 2022), tidying the second draft of the book before submitting it for the publisher's peer review, I have been reading about new experimental results leading to renewed conjecture about the idea of a potential fifth fundamental force. Such glimmers stimulate hopes that they might help to clarify, and fill some of the outstanding gaps between

theory and experimental observation, in what the current Standard Model of physics can account for.

In computer science, the Turing machine arrived in the 1930s as an abstract model of computation used to represent algorithm and computer program. Models of logic, expressed by mathematical logicians and argued over by philosophers, took on computational form. Computer languages (for example, Simula (1962) and Prolog (1972)) evolved for expressing different ways of representing, simulating and reasoning with different kinds of computer-based models. Algorithms and computer programs evolved, expressed in these languages, to use these models in real life contexts. With the advance of the computer came ever more extensive computational methods and models of complex systems. These could be designed to represent the modelled system in finer detail and predict its behaviour more extensively. They used newly devised numerical methods, implemented in computer programs, to converge on accurate and stable solutions of the model equations, by iterating successive approximations in many steps.

X-ray imaging and the rod and sphere models of chemical structure of molecules, provided the methods and resources that guided Francis Crick (1916–2004) and James Watson to their imagination of the double helix of DNA, in perhaps the best-known scientific example of such physical models.⁵ What are called ‘animal models’ of disease have been extensively bred for researchers to test potential treatments, as a preliminary stage in creation of safe human therapeutics. These *in vivo* model experiments can, to an increasing extent, be replaced by *in vitro* and computer-based testing, drawing on extensive databases of the properties of already known and described macromolecules, pharmaceutical materials and products. An early example of such databases was the Swiss-Prot protein sequence databank, developed in the Geneva University Hospital, in the department of my late,

5 As I was first drafting this section, the *Times* newspaper (20 May 2020) reported the Princeton/Bergamo study of the design of Renaissance domed churches and cathedrals—at the interface of abstract geometry, architecture, engineering, and construction. The domes of Antonio da Sangallo (1484–1546), built around the 1530s, were known for their engineering feats, including foundations for building on unstable land. One such dome, built in two self-supporting layers, comprised a ‘cross-herringbone spiral pattern’—a double helix—constructed from vertical bricks and filled in with horizontal bricks. This construction avoided the need for expensive wood support framing. It felt a beautiful analogy with the chemical bonds stabilizing the flexible DNA double helix molecular structure—a double helix design of vertical ‘bricks’ (A, C, G, T), with bonding horizontal ‘brick courses’ (chemical bonds), creating a stable ‘Renaissance dome’ (DNA molecule) on ‘unstable ground’ (the living cell)! Maybe a foolish aside, but fascinating and fun!

much-respected colleague there, Jean-Raoul Scherrer (1932–2002). As early as 1969, he worked on pulling together the DIOGENE patient-centred care record for the hospital information system.

The mapping of genomic sequences to computational models of stereochemical molecular structure has evolved very rapidly in recent decades, and from this came new understanding of normal and abnormal folding of protein structures, as discussed in Chapter Three, and much else. The amazing new scientific edifice of pharmaceuticals that has resulted has, in large part, been a creation based on the work of very many scientists collaborating within the public domain. One wonders what might have happened had commercial intellectual property patenting norms prevailed over these innovations at that time. Drug and vaccine design have advanced rapidly from these beginnings, with the advent of three-dimensional computer-generated models to assist in matching their molecular design to target cellular receptors and processes.

In Education and Training

The flight simulator is a well-known example of the use of a model to assist training and assessment of a skill. The Harvey simulator, used in training doctors, nurses and paramedics in common clinical skills, is an example of the use of simulators in the clinical domain. Harvey focuses on practical handling and interpretation of clinical signs and interventions, such as intubation of breathing apparatus, intravenous injection of drugs and cardioversion. My work of two decades from the 1970s, at St Bartholomew's Hospital (Bart's), was largely devoted to the development of a range of highly experimental computer-based models, and their testing and use as simulators within educational and clinical environments. These models, known as the Mac Series, were created by a team of which I was a member, led by my then head of medicine at Bart's, John Dickinson (1927–2015), in the context of his close connection with the pioneers of the Medical School at McMaster University in Ontario, Canada. This history recurs in the examples below, of computational modelling pioneers I have known, and their related initiatives that I connected with and worked on, along my songline.

The Mac Series models were distributed widely across the world, first by me, on nine-track magnetic tapes, for incorporation within preclinical, clinical and professional educational curricula. They even cropped up and surprised my medical student daughter in her respiratory physiology course at the Nottingham Medical School, where I had previously got to know her lecturer! As I further describe in the next section, the graphical versions

that I created at the start of the transition between the minicomputer and microcomputer technology eras, were published by the IRL Press (no longer in existence), based near Oxford, as part of a trial of where computer software might fit within their future publishing business. I was a member of their Software Advisory Board, and then that of Oxford University Press, at that time. These published versions of the models received favourable review in a special issue of the *Times Higher Education Supplement*. I have resurrected them—nowadays just historic artefacts of early and now obsolete educational technology—from the only extant copies I know of. The programs are on ancient floppy discs together with their published manuals, and I have preserved them at home for now over thirty years, along with some related books written at the time.⁶ Through these years, I also collaborated closely with Leonard Saunders at the University College London (UCL) School of Pharmacy, to extend the drug pharmacokinetic modelling work into pharmacy education and research.⁷

Three Sliding Doors

This seems an appropriate point in the storyline of the book, to interleave the context of some major transitions in my career in health informatics during the years around 1990. These were marked by a switch from a career focused on the computer modelling of body systems and medical education, to one concentrated on imagining and evolving a standardized architecture for electronic health care records. They broadened the interdisciplinary and multiprofessional flavour and scope of my work, bringing unique new opportunities in the decades that followed to link health informatics within

-
- 6 These graphics versions of the Mac Series models have been reconfigured to be made available, open-access and online, with permission of the surviving authors, as part of an electronic archive of additional resources to accompany this book (available at <https://www.openbookpublishers.com/books/10.11647/obp.0335#resources>). In Chapter Five, I tell the story of how a community of programming enthusiasts, working in the public domain and dedicated to sustaining the long obsolete MS-DOS microcomputer operating system and applications that ran on it, recently enabled me to resurrect the programs in this way, and hopefully now keep them preserved in working form—something I had long-considered unachievable.
- 7 L. Saunders, D. Ingram, C. J. Dickinson and M. Sherrieff, 'A Comprehensive Computer Simulation of Drug Metabolism and Pharmacokinetics', *Computers & Education*, 6.2 (1982), 243–52; L. Saunders, D. Ingram, and S. J. Warrington, 'The Pharmacokinetics and Dynamics of Oxprenolol: A Simulation Study with Six Subjects', *Journal of Pharmacy and Pharmacology*, 37.11 (1985), 802–06; L. Saunders, D. Ingram and S. H. D. Jackson, *Human Drug Kinetics: A Course of Simulated Experiments* (Oxford: Oxford University Press, 1989).

the mainstream of health care education, research and practice. There were three sliding doors that I stepped through, along with close colleagues at that time, into new environments that we created together.⁸

Door One

On becoming a professor in 1989, one of the first innovations that I was asked to take forward for the Bart's Medical College was a proposed new joint medicine and nursing clinical skills centre supporting undergraduate and postgraduate education. The project brought together a small team drawn from the local medical and nursing academic communities, which I was asked to lead. The way in which this first sliding door arose, and the innovative health care environment that, by stepping through it, we were enabled to create over the coming years, placed me at the centre of a widening multiprofessional scope and context of the evolving field of health informatics.

This Clinical Skills Centre, as it was later characterized, was the first such initiative in the UK and arose from growing concern about rigour in the assessment of clinical skills, in formal examinations framed around the traditional apprenticeship model of bedside, in-practice teaching. The concern extended from undergraduate education into the assessment and regulation of professional practice, and the clinical team that came together around me at Bart's, then, and later at UCL—especially Jane Dacre and then Lesley Southgate, in this context—went on to lead much wider initiatives of medical royal colleges and the General Medical Council. These sought to improve accreditation and regulatory roles in the assessment of the performance of doctors, through formal examinations and continuing improvement and review of clinical professional practice. These were, and remain, complex, multifaceted and contentious concerns; highly dependent on how they are approached, the information on which they are based, and their professional leadership and acceptance. They are in some parts highly political in nature, and in others not political at all—true of wicked problems, more generally, as I discuss elsewhere in the book! Lesley and Jane both enjoyed stellar subsequent careers, elected as Presidents, respectively, of the Royal College of General Practitioners and Royal College of Physicians, and were honoured with Damehoods by the Queen.

8 Further detail of these environments is covered in Chapters Eight and Nine.



Fig. 4.1 Jane Dacre—pictured here as President of the Royal College of Physicians of London. A leading figure in undergraduate and professional medical education and a close colleague at Bart's and UCL. CC BY-NC.

Interprofessional collaborations are never easy—between medicine and nursing, professional and personal history, and some rivalry, are seldom far from the table! For our new centre to be successful, it was crucially important to work together across this divide, with a shared goal and intention to build a good environment and team culture, to deliver it. My then dean, Lesley Rees (1942–2022), had recently become the first woman dean of the venerable Bart's Medical College and had a good relationship with the head of the adjacent Nursing School, Susan Studdy. It was heart-warmingly adventurous of them, and personally motivating for me, not being clinically trained and thus inevitably an outsider in the medical school hierarchy, that I was called on and trusted in this way to lead the initiative.⁹ I reflect further

9 The project needed someone willing and able to lead from below in the hierarchy and I seemed to fit that bill—a senior professor of medicine or nursing would, likely, not have done! My sponsor professor of medicine, John Dickinson, took on the building project, liaising with the architects, and we recruited Jane Dacre as the lead medic and Maggie Nicol as the lead nurse—both then in their early careers, but clearly destined for great things—to forge a professional doctor/nurse alliance and work on the clinical skills teaching curriculum. They got on very well. I drew everyone together in running the project team, which included Diana Holroyd and Sonia Crow, senior nursing school staff who represented Sue Studdy. Together we set out to, and succeeded, in creating an effective and trusting team environment, respectful to both traditions and the sensitivities that could easily have become inflamed. That was, I think, in large part because we all approached

on this endeavour, as an example of the creation of a new environment, in Chapter Nine.

This was one of several wider leadership roles that I found myself moving into during those years of career transition: being hands-on in creating a computer infrastructure and support service, *de novo*, for the clinical departments of the Bart's Medical College; creating a novel videodisc-based educational resources to support multiprofessional training in support of cancer care in the community, for the Marie Curie Foundation; writing machine code software for a twenty-three videodisc series aiming to support the teaching of human anatomy, for an A-V industry-based project; helping to stabilize a project team that had got into team relationship difficulties when producing a computer-based resource supporting tropical medicine, for the Wellcome Trust; and then, most impactfully of all, leading the GEHR (Good European Health Record) academic, industrial and health services consortium and project of the European Union (EU), from 1990 to 1994, creating new foundations for electronic health record architecture. This latter opportunity—brought to me, persuasively, by Sam Heard and Alain Maskens, who overcame my initial nervous scepticism—led over the next twelve years to the establishment of the multiprofessional and interdisciplinary Centre for Health Informatics and Multiprofessional Education (CHIME) at UCL, in 1995, and creation of the openEHR Foundation at UCL, in 2003.

The successes and failures experienced along these separate but connected and evolving pathways of my career in health informatics have permeated widely in the narrative and framing of this book. They flowed from my presence over many years, working at the centre of health care communities and being given permission and freedom to explore. These beginnings led to conferment of more than ten full professorships in informatics and medical and health care-related fields, over the coming fifteen years. The huge amount I owe to many colleagues, in making this possible and making it work, cannot be overestimated. Their contributions feature and are acknowledged in multiple places throughout the book. Several of them have been reading and advising as its writing has progressed.

Door Two

During the time of the Skills Centre project, I was also busy thinking about creating a new environment for my academic work in health informatics. This would have had quite limited prospects within the Department of Medicine, after my sponsoring professor's approaching retirement. It was

the project in that way. Jane Dacre et al., 'The Development of a Clinical Skills Centre', *Journal of the Royal College of Physicians of London*, 30.4 (1996), 318.

an uncertain time and led to a second sliding door. Lesley Rees agreed that I could move my small team out of the Department of Medicine on the hospital site in Smithfield, and to rooms adjacent to the Medical College General Practice Department, situated at Charterhouse Square. This was at the generous invitation of its Head of Department, Lesley Southgate, and her colleagues, who included my subsequent close colleagues and team members, Marcia Jacks, Sam Heard and Dipak Kalra. In 1991, Jane Dacre and I established a small new department there, brashly called Clinical Skills and Informatics, and moved to Charterhouse Square, along with colleagues in my small team of that time. It was a staging post that positioned us for new opportunities to come, as the different stages and directions of our careers played out. It was a brave leap of faith into the unknown for us all!

The name of the department was a statement of ambition—the sort of outlandish idea that a new professor in a new field is sometimes indulged and encouraged to come up with! Such indulgence tends to wear thin if not fruitful, and unsuccessful indulged heads tend to be lopped within a couple of years or so—luckily mine survived intact in academic life for twenty more years, until my retirement!

Little could either Jane or I have imagined how far and how rapidly the interrelationship of the nature and skills of health care and professional practice, and their connections with the computer and its burgeoning aspirations and progress towards AI, were destined to advance during the following years, as unfolded in the storyline of this book! The advancing story of health informatics, and especially, now, of a hoped for benign and humanly supportive AI, has profound and increasing implications and impacts throughout the spectrum of health care education, professional practice, service delivery, research, governance, regulation and legal accountability. It has equally profound implications for every citizen in context of their access to, engagement with and expectations of health care services. It is a vector of continuous disruption of the markets and industries that support health care, in their products and services. And, thus, it is politics and business writ large and an unrelenting headache for the politicians and civil servants struggling to advance policy and strategy and control and manage the purse strings of the NHS!

Once again, the way in which this second sliding door arose—and the innovative multiprofessional and interdisciplinary health care environment that, by stepping through it, we were enabled to create over the coming years—placed my work within a still wider everyday practical context of health care services and their community and industry relationships, which were becoming increasingly central to the evolving field of health informatics.

It was in this new environment that the GEHR project, the antecedent of openEHR, was nurtured into life, as described in Chapter Eight and a Half. Alongside, we made new connections with others in the academic community that Lesley Southgate and her predecessor, Mal Salkind, had drawn together. They included Ann Bowling, a national figure in health services research, and Brian Jolly, engaged in a nationally funded higher education project exploring the teaching of medical ethics, clinical communication skills and health informatics. Another close colleague of subsequent years, Jeannette Murphy, joined us at that time, to spearhead the connecting of health informatics with the medical education curriculum. And it was from there that the third sliding door of those transitional years of my professional life presented itself, three years later.

Door Three

Our Charterhouse Square group's combined profile in health informatics and multiprofessional education, and success in leading the Bart's Skills Centre project and the GEHR project focused on a standardized health record architecture in Europe, was seen by new eyes. These matched us with ideas for new academic developments at the nearby UCL Medical School and its Whittington Hospital campus at Archway, in North London. This complementarity led to the invitation for our combined team to move there to create, and me to lead, a new department, christened CHIME (Centre for Health Informatics and Multiprofessional Education) by its facilitators and founders at UCL and the Whittington, David Patterson, John Pattison, Helene Hayman and Derek Roberts.

The unfolding story of the creation and lifetime of this new academic environment is told in Chapter Nine. The experience of creating and leading CHIME, and of the multiple new local, national and international relationships and roles that I was drawn into, there, became the central context of my evolving sense of the priorities and necessary steps to position health informatics appropriately on the changing landscape of health care, in its transition through the Information Age into the Information Society. This was also closely informed by two of my children's experiences as trainee doctors of those times, and my wife's experience, informed by her time as a doctor in a different country.

The work on the major projects I led, the changes in assessment of clinical skills and performance that Jane and Lesley led, and my switch in those years to leading teams envisaging and creating new foundations for digital care records, internationally, confirmed for me the oncoming inevitable migration of informatics, as a discipline, to the heart of medical science and health care. This migration occurs alongside transformational

change in the nature and ways of working of health professions, in their skills and services. I was probably one of only a few people working in academic medicine and health care services of those times who was seeing things from that blue-skies and transitional perspective. But fortunately for me, I was trusted and supported to engage in that spirit, imagining and working to help create new environments supportive of transition towards future Information Society health care.

I had also come to believe that no amount of talking or writing about these still nebulous ideas was likely to enable much of significance to be learned about them, created and sustained—rather the opposite, in fact. A head-down focus on the practicalities of implementation, and thereby learning by doing—as Jo Milan brilliantly exemplified at the Royal Marsden, as described in Chapters Five, Seven, and Eight—was the order of the day! I draw on this experience in Part Three of the book, where I consider the challenges faced in creating and sustaining a future care information utility, designed, operated, led and governed across the disciplines, professions and supporting industries of health care services, in partnership with the citizens and communities they serve.

Reversing back through these three sliding doors, I continue now with the overview of the purposes and methods of modelling and simulation.

As Tools of Design and Engineering

In this section, I move into the world of design and engineering, where models are used to observe, predict, make decisions about and control behaviour of diverse systems, from small-scale devices to industrial plants and national power grids. They may embody algorithms for sampling, filtering and analyzing measurements of the system under consideration, fitting parameters that characterize the model to match as closely as possible to the observed experimental data collected, and estimating how precisely the model predictions can be known and relied upon. The validation and usefulness of the model rests on how well performance aligns with purpose.

As a young boy, I spent many happy hours making a quite large model aeroplane, from balsa wood struts, parchment-based bodywork hardened with cellulose dope, and a miniature diesel engine to drive the propeller. It was a thing of beauty and the maiden flight on a nearby hill a memorable event. I filled the small fuel tank, flicked the propeller to start the motor, and gave it a gentle launch. It took off and climbed with immaculate trim, soared higher, disappeared into cloud—to crash somewhere unseen and be found, broken-winged, with me broken-hearted, an hour or so later! It was a model aeroplane, a real thing but not the real thing. It represented many

of the features of the real thing, enough to give me a sense of the real thing. I had fun making it and may have learned a bit about design, construction and flying of aeroplanes, but clearly nothing about controlling them! No harm done!

Real plane crashes cause loss of life, however. The de Havilland Comet passenger jet aircraft crashes of my 1950s childhood, arose for lack of prediction of the impact of metal fatigue in propagating the collapse of its structure under pressure in flight, and failure to design accordingly, to minimize that risk. To reduce fuel consumption in flight, holes were made in struts to reduce their weight. Windows had rectangular shape—creating seeds of dislocation and fracture at the corners of the surrounding metal structure. I have seen them, preserved by aircraft manufacturers and shown to school children attending their aeronautics open days. These early passenger flights were essentially experiments with what proved fatally flawed aircraft designs. They were learned from in those years at the expense of lives. The crashes derailed the nascent jetliner business in the UK—the site of the de Havilland airfield at Hatfield is just a mile from our house and now houses commercial warehouses, elite-car showrooms and open common land for walkers. The old runway can still be traced and a plaque there commemorates the Comet.

In the early 1960s, I saw the prototype engineering systems for the Concorde supersonic plane laid out within a huge hanger at the Filton aerodrome, in Bristol, to test the evolving design concepts in play. A small, odd-looking paper-dart-like plane was built and flown, to simulate the delta wing design. A Vulcan bomber aircraft was fitted, and flight tested, with the prototype Olympus jet engine. The cramped cabin interior was mocked up in wood—I walked inside and wondered how those rich enough to fly in it would fare in such cramped conditions. The prototype flying models were star attractions at the site, for avid visiting school children, me among them, and at the annual Farnborough Air Show, where businesses and spectators merged for this national showcase of aeronautics.

Computer models of aeroplanes and flight simulators are now extraordinarily realistic. Engineers iterate aircraft design concepts using them, and pilots learn to fly using them. Of course, there is always an inaugural first flight, but the Comet-like disasters of former times are now avoidable, albeit that it seems, just recently, that commercially driven short cuts led to unsafe, unstable designs that persisted into large scale production and early flights of the Boeing 737 Max aircraft.

Flight testing of a new aircraft design and construction is now closely linked with tests performed with simulations. The models are used to explore intermediate ranges of behaviour of the modelled system, where there is limited or unduly expensive opportunity to measure and monitor

the full range that would be encountered in the real world. Multiple simulated test flights replace real ones in the programme of trials employed. A test that forces the trim upwards until the plane stalls may also be easier to envisage conducting in a simulation, provided, of course, that one already has sufficient confidence in the results obtained, when extending the model set up that far!

Models of buildings are made to show off architectural designs and Lego model villages are built as tourist attractions. The dressmaker's dummy for fashioning and making clothes is a model still widely used. The catwalk human model shows off the finished product to potential retailers and purchasers, promoting a world of glamour and make believe. Today, body profile can be scanned, and used to build and calibrate a three-dimensional computational model, embodying both form and movement, using Computer-Aided Design (CAD) software. The clothes designer has access to database archives of clothing materials and their properties, from which to select when creating a new design. A real prototype can then be constructed, and the modelled design iteratively improved, until the product is approved. The computational model then integrates with cutters and sewers, setting out patterns of materials to be cut out and stitched together in its largescale production. In the clothes shop or online, the model of the tailored garment can be calibrated and displayed, based on the size and preferences of a specific potential purchaser.

A final contemporary example, exhibiting the power of bioinformatics, has been the speed of design of candidate Covid-19 virus vaccines. The development cycle, from the identification of viral genomic sequence, modelling and picturing of the viral surface, and selection of candidate binding targets, through to detailed design and testing of the vaccine, is a process drawing on multiple scientific and computational methods and resources. And the shrinking of the timescale, from conception to approved use in at risk populations, an outstanding achievement.

Within Products and Services

As just illustrated in the case of design and production of clothes, nowadays the design and visualization of manufactured goods, and plans for their production, are often supported by computer models built with CAD software. Prototypes of new products are made and tested at key points in their production cycle; the underlying computer models connect and integrate with design, production and lifetime maintenance and repair, and with quality control and monitoring at all stages of the product life cycle. Costs throughout are estimated, based on historic accountancy data.

The model is thereby integral with related production facility, organization and personnel plans. A similar integrative approach is available for much larger-scale engineering and infrastructure projects. It provides a rigorous and consistent framework of workflow and cost projections required to underpin management and oversight of the process.

In the area of services, models of weather and climate are used to forecast outcomes in detail over coming hours and days and, in broader-brush trends, over coming months. Models of traffic flow simulate how the networks of transportation will be affected in different operating contexts, such as bad weather, scheduled repair and breakdown.

Models created to simulate behaviour of a system may take a quite different form than models used to control that same behaviour. In one case, they are used to predict outcomes of choices under consideration in the design of a new device or system. In the other, the purpose concerns control of its behaviour in practice. In control theory, a class of methods is the model-based control system. In this, a model of the behaviour of the system is built into a system controller integral with the system itself. When the system diverges from the desired operating state, or transition to a new state is required, potential changes to the system's adjustable parameters are tried out first in the onboard model and tuned iteratively to a level that would be expected to achieve the desired change, as predicted by the model, before being applied to the live system.

Models beyond Experimental Validation

In the early 1970s, Jay Forrester (1918–2016) and the Club of Rome, sought to model the global economy. It was a widely ridiculed but nonetheless worthwhile attempt, if perhaps not a realistic expectation. It seemed unimaginable that such a model could achieve more than the broadest of broad-brush approximations to reality or reach more than largely common-sense conclusions. As Walter Sellar and Robert Yeatman might have said, as in their 1930 spoof history of England, *1066 and All That*, 'wrong but wromantic'¹⁰ The quest fizzled out. At this level, only one experiment is conceivable—the uncontrolled experiment of how events play out in real life. In the modelling of complex and interconnected systems, rigorously controlled experiment is a rarity and may be impossible. Clinical practice uses controlled trials in its efforts to tame variability in experimental

10 W. C. Sellar and R. J. Yeatman, *1066 and All That. A Memorable History of England Comprising, All the Parts You Can Remember Including One Hundred and One Good Things, Five Bad Kings, and Two Genuine Dates* (London: Methuen, 1930).

observation and measurement of human subjects, to home in on and quantify outcomes reliably ascribable to interventions made.

Models and simulations have found their way into the parlance and practice of economic and social systems—domains where human thought and behaviour are closely embedded, and experimental method not of the essence. My Daniel Kahneman inukbook, *Thinking, Fast and Slow*, describes how patterns of human thought influence economic choices—for example, weighing risk of gain and loss disproportionately.¹¹ In another quite recent inukbook, *The End of Alchemy*, Mervyn King cautions that mathematical and computational models and simulations addressing such matters, should be supped with a long spoon.¹² He argues that narrative and storytelling should feature more strongly in what he describes as a current crisis of ideas, more than of institutions and methods employed.

Governments have the unenviable task of setting fiscal and monetary policies in the context of economies which are extremely difficult to understand and predict. They review policy options using computer models that describe the underlying principles and behaviours that mathematicians, scientists and economists believe to be in play, and chose policies guided by these model-based predications. And in the world of commerce, computer models of financial markets are used to drive trader advantage—for example, by anticipating market movements and being first to the start gate for buying and selling of shares, thereby achieving first mover advantage, or by gambling to nudge their market value in a preferred direction that would benefit the trader.

Models that Can Lead Astray

Richard Feynman (1918–88) notably once remarked that you can prove anything by analogy. That amounts to the same as saying that you can prove nothing by analogy and mirrors Whitehead’s cautionary advice quoted at the head of this chapter. Ability to help clarify and guide understanding are aspects of the usefulness of a model but the quest for a perfect model is an illusory goal.

One can make a simulation fit with experimental data in many ways, providing that one has enough requisite adjustable input parameters of the model at one’s disposal, to shape the predictions it makes to achieve that fit. In mathematical language, these might be called degrees of freedom.

11 D. Kahneman, *Thinking, Fast and Slow* (New York: Macmillan, 2011).

12 M. King, *The End of Alchemy: Money, Banking and the Future of the Global Economy* (New York: W. W. Norton and Company, 2016).

The process may work for the data at hand, but what about data yet unseen, where a continued good fit may require further embellishment of the model, introducing more degrees of freedom. Where are Occam's Razor and Popper's advantageous simplification, here? This can become disadvantageous complexification. The validation of a model must combine assessment of its purpose, performance and usefulness, and its feasibility depends on ability to test and evaluate these under suitably managed and controlled conditions. The more complex and extensive the system modelled, and the context in which it operates, the harder this is to achieve. An unvalidated model is essentially a loose analogy.

In science, experimental validation is the arbiter of theory, however enthralling the theoretical abstraction. The modelling process is a creative one, and especially so if it stimulates new ideas and leads to tractable new experiment to test and confirm them. Many models will prove unsuccessful or misleading, but failure is there to be learned from. Problems and disputes easily arise where belief rather than experiment becomes the arbiter of the validity of a model.

With the generally more tightly bounded models used for design and engineering purposes, there is less excuse for, or tolerance of, getting the models wrong or misusing them. Unforeseen behaviours—such as the lock-step lateral vibrations that built up in the Millenium Bridge across the Thames, between Tate Modern and St Paul's Cathedral, when first opened for public use—can quickly derail the designs that the modelling had led to, quite apart from tripping up both designers and users in that case! The Tacoma Bridge must presumably have had a test outing *in silico*, too, but rigorous perturbation testing, to simulate the traffic and wind-induced vertical resonance that destroyed it, must have been lacking in the test schedule.

Where human understanding of the modelled domain is complex and uncertain, the model may sometimes be taken as offering beguiling and yet spurious certainty. As Voltaire (1694–1778) wrote, 'doubt is uncomfortable, certainty is absurd'.¹³ All models rest on assumptions, simplifications and approximations. As with the wobbly bridge over the Thames, where these are wobbly, the predictions and their consequences may prove wobbly, too! Returning to Whitehead's caution, we must be cautious about tendency to allow the model to supplant the reality in our thinking and reasoning about the domain it represents. We must always keep in mind the purpose that the model is intended to fulfil, and the evidence that it can and does fulfil that purpose for the situation at hand. This can only be achieved if there is

13 Letter to Prince Frederick William of Prussia (28 November 1770).

clear and rigorous connection of the concepts, structures and behaviours embodied in the model and the real-world evidence and experience it relates to.

For this reason, we should be cautious of predictions based on opaque models. Models that should best be open to critical review and inspection are sometimes kept under wraps, for a mixture of academic, commercial or political reasons. We should be cautious, too, when modelling seeks to accommodate the uncertainties of human behaviours, however relevant these might be for the purposes the model serves, because they may also, likely, prove very difficult to handle confidently.

Given the intrinsic limitations of their model formulations, modellers often perform a sensitivity analysis to explore how the model predictions change in response to perturbation of the model parameters and initial conditions, and thereby to place confidence limits on predictions made when simulating future events. This has become a talking point in the context of the wide-ranging predictions of different teams simulating the expected impact of different lockdown strategies for limiting the spread of Covid-19 infection. The proof of the pudding is in the eating, but here it must also take account of what was cooked, how and why, and who was eating, when and where! Bland assertion of reliance on science, engenders public scepticism and suspicion that it is the books that are being cooked!

Notwithstanding these caveats, computational models are now central to the study of complex systems and the mathematics and science of complexity has moved forward alongside efforts to build such models. There is now greater understanding of the practical difficulties that modelling, and simulation methods pose, and their intrinsic limitations. Clarity of purpose is essential. Models should correctly utilize the methods and materials they employ but cannot be relied on as correct representations. They can perform well or badly for the purposes they serve, and the struggle to make them better and thereby more useful, is a strong motivation for exploring, using and learning from them. This may necessitate long-term endeavours and require sustained capacity and resource.

The next section visits physics, life science, clinical science, and health care, for examples of models used to explore and represent systems—what they aimed for, how and why they prospered and how and why they failed. These are stories I have personally encountered, firsthand, and some in which I also participated. They illustrate how models have become central to the creation of new knowledge and the taming of its related and exploding data sources, over the timeline of my career.

Pioneers of Computational Modelling

Physics–John Houghton–Modelling Weather and Climate

We share weather and climate. They connect our lives locally and our welfare, globally. This is not to say that the local weather is ever the same at different times and places—it clearly is not, anywhere: minute by minute, a rainstorm moves. This is not to say that the climate of Malaysia, central Europe or sub-Saharan Africa are ever the same—they clearly are not, either. But to model weather and climate in any of these contexts, one must understand and be able to compute with the same science and know how to use that knowledge in the different contexts in which one seeks to understand and make forecasts. One must know how to allow for the differences of local environment that play out in determining the weather there, from locality to locality, and time to time, and how accurately these can be forecast.

Why start with physics and weather? Well, illness is like bad weather and good health is like a sunny climate. Illness and health are local and global. Okay, a bit sophistic, but pandemic and health can learn from weather and climate. Discipline from around the circle of knowledge is now in play in these models, and mathematics and physics is where the weather and climate models started from. Arguably, though, machine intelligence may now sometimes prove a better bet for making accurate short-term forecasts, exploiting the measurements available. *Zobaczmy [we will see]!*¹⁴

Weather forecasting has a basis in measurement, science and engineering that have all advanced steadily through the Information Age. It is now largely based on mathematical models and simulations, extending over large-scale patterns and local variations of weather. Today's forecasts are based on hundreds of millions of data points, from land, sea and air radiosondes, covering the earth's surface on a one kilometre square grid. Forecasts have improved—in recent decades they have been shown capable of predicting accurately one day further ahead in time, each decade.

In my childhood, the village doctor diagnosing and treating illness in a child and the local weather forecaster predicting weather over the coming period, shared some common traits. There was science and craft in play—skilled observation of what could be seen and felt, combined with limited measurement, put together with a wealth of knowledge and experience. A rough and ready seaside hydrometer was a strand of seaweed hung on an outside door of the house. Its feel reflected moist or dry air, and this in turn reflected and helped in anticipating changes in the weather.

¹⁴ On this Polish expression, see Preface.

There was good reason to want answers about what was going on and about to happen. Is it going to rain? Is this child's fever a serious concern? Of course, the doctor could, or was expected, to do something about the illness. Weather forecasters were not expected to do anything about the weather, but villagers could put on raincoats and the farmer could act in anticipation of a storm. Similarly, the village general practitioner (GP) might take one look at a sick child and know what was wrong and what to do (or not to do)!

Life has moved beyond invocation and sacrifice to unseen gods, or employment of mystical rainmakers to dance and bring on rain for the crops. But predicting the weather remains a part of culture and folklore. In England, weather arrives mainly from the west and departs mainly to the east. In Poland, the unpleasant damp clouds come from the west and the unpleasant dry and dusty winds come from the East—a metaphor of how Polish people feel, and with good reason, about the currents that have buffeted the history of their lovely country! According to Somerset folklore: 'Lundy high, fine and dry; Lundy low, it's going to snow', reflects how this small island appears on the horizon, from land several miles away. Cynical doubters had a (very English!) variant: 'If you can see Lundy in the morning, it's going to rain; if you can't see it, it's already raining!'

Weather stations at home were quite common in my childhood—anemometer, hydrometer, thermometer and barometer used in creating personalized weather forecasts. Early morning and late-night radio intoned the shipping forecast, based on measurements and observations from a network of lightships and weather stations, at sea and around the coast. These forecasts were published as contour maps of atmospheric pressure and temperature, and patterns of air movement reflecting the rotation of the earth. They were correlated with winds circulating air across warmer and cooler areas of the earth's surface, and seasonal variations according to position of the earth in its orbit, with rotational axis inclined to the ecliptic.

Now, a worldwide network of sensors and satellites collects data for creation of images and calibration of models that span continents. These depend on a functioning global digital network of communications, as do the telescope terrestrial networks and solar system voyagers through which astronomy and astrophysics advance. These networks were being invented and piloted at the time I attended Peter Kirstein's (1933–2020) lectures on telecommunication in the late 1960s, at the London Institute of Computer Science, later incorporated into UCL. He set up the first transatlantic connections with the US Arpanet and is recognized as a founding father of the Internet. This was five years before Conway Berners-Lee (1921–2019) and Ted Coles were selling computers into hospitals, for International Computers and Tabulators (ICT), and Conway's son, Tim Berners-Lee,

arrived at Oxford to study physics, before going on to the Conseil Européen pour la Recherche Nucléaire (CERN) in Geneva, where he conceived and implemented the software network protocols which formed the basis of the World Wide Web.

One of my Oxford physics lecturers in the mid-1960s was John Houghton (1931–2020). His course on atmospheric physics was memorable—clear, concise and well presented. Olympian erudition, in welcome contrast to the bafflingly unexplained, but clearly brilliant and enthusiastic, genius of some. I occasionally attend physics alumni events in the department, today, and the standard of presentation all round is now hugely better. Houghton started his course by considering only a small volume of moist air, for which understanding at the level of the gas laws of Robert Boyle (1627–91) combined with the later understanding of the solid, liquid and gas phases of water, enable modelling and prediction of its behaviour in terms of pressures, volumes and temperatures. He was fluent in thinking on his feet, showing how the vector methods and calculus that we were coming to grips with in our first, highly mathematical year of study, could translate into a set of equations that modelled this small volume of air. These, in turn, became the core of generalization to more extensive and complex models of the atmosphere.

Imagine a weather forecast predicting that, on the one hand, the country may be covered in cloud and snow, at -25 degrees Celsius, or, on the other hand, it may enjoy a clear sky and heatwave at +25 degrees Celsius. Those sorts of swings occur naturally over several months from winter into summer—the frozen Niagara Falls giving way to a steam bath, between January and July, as I have seen with colleagues and friends, during visits to the nearby McMaster University. Such swings get wider and more erratic as climate becomes more chaotic. Earthly air temperatures now range over approaching 150 degrees Celsius. Life is tolerant of extreme cold but less so of extreme heat.

And hurricanes do happen in Hampshire, although hardly ‘hever’, as intoned in George Bernard Shaw’s (1856–1950) play, *Pygmalion*.¹⁵ Our children’s home summer camp on a farm near Beaulieu in Hampshire, opposite the Isle of Wight, was hit by one in the mid-1950s—trees everywhere were uprooted. My dad’s ten carefully erected bell-tents, sheltering twenty-five young campers, withstood the storm. He always brought one-metre wooden tent pegs to the camps and was teased for his caution. He sledge-hammered them into the soft ground at the height of the night-time torrential storm, and ours was the only camp still standing by the next morning, that had not been blown away in the night, like all the others! Drenched scouts and guides had been hurriedly decamped to shelter in the farmer’s huge

15 G. B. Shaw, *Pygmalion and Major Barbara* (New York: Bantam Classics, 2008).

hay barn, where we saw them as we walked to the milking parlour next morning, to collect our day's supply of delicious fresh, cooled milk!

The physics of the atmosphere from sea level upwards, is extremely complex. It is in a dynamic balance with circulating ocean currents and land masses, small islands and major continents—some moistly forested and some aridly dry regions. Gyroscopic forces from the earth's rotation, solar energy incident, scattered, absorbed, and radiated, from air, cloud, icecap, desert and ocean, all interact. The modelling of ocean depths and flows, such as the El Niño, entraining with atmospheric weather conditions, is still rudimentary. The underwater geography channels a complex system of ocean currents and tides, such as the Gulf Stream, which mixes, ebbs and flows, through tides and seasons.¹⁶ Given these still intractable unknowns, how has weather forecasting bootstrapped from mystical prognosis to everyday utility?

Weather forecasting evolved along complementary axes, of measurement, modelling and simulation. Limited intermediate goals and methods could be framed and explored, with a clear measure of success of the exercise—how well did the forecasts based on these simulations perform in relation to measurement and observation? There was a well-understood and articulated purpose and value in achieving increasingly accurate local forecasts—guiding farmers in planting crops and reaping harvests, and all of us in what to wear.

Chaos theory came into the discussion, through the butterfly wings metaphor describing the widescale magnification and propagation of minute influences. Analysis of the mathematical properties of the underlying equations of the models and their measured parameters enabled probabilities to be attached to forecasts—for example, a model predicts that the probability that it will rain in St Albans between 10am and 11am today, is nine percent.

Paradoxically, places where the weather patterns are predominantly very stable and predictable are difficult places about which to model the disturbances that do arise. Daily alternating patterns of heavy rain and fierce sunshine in Singapore and Malaysia, were highly consistent, but disturbances to this stable pattern cropped up quickly and unpredictably. Snow in desert-encircled Riyadh is not unknown and is a somewhat mystical event. For me, venturing outside from the air-conditioned, refrigerator-like hotel lobby, to travel to the University Hospital to run courses there, was like walking into an oven. Travelling in the desert wadis nearby on a weekend

16 The General Bathymetric Chart of the Oceans (GEBCO) is being pieced together from robotic surface vessels, commercial and public data, and crowd sourcing of sonar measurements from large ships and private yachts. Six percent of the area was charted by 2017, twenty percent is expected to be accomplished by the end of 2020, and the survey completed by 2030.

trip, one felt extreme vulnerability when walking even a short distance away from the air-conditioned vehicle and supply of cooled water. It was extraordinary to pick up coral from the sand, showing how much landscape and climate can change. One understood how Saudi colleagues who I once accompanied on a winter visit to cold, windy, rainy and gloomy, Wuthering Heights-like Lancashire and Cumbria (to visit a shipyard there, during my time working in the heavy engineering industry) found it an ecstatic experience of beauty and wonder!

Uncertainty in the prediction that it will rain, reflects both the quality of measurements on which the forecast is based, the appropriateness and reliability of the model itself, and the assumptions it rests on. It may reflect that the system modelled is intrinsically variable, perhaps in some statistically well-characterized or assumed manner. This may favour a stochastic modelling approach rather than a deterministic model of the relationships. It may reflect pragmatic choices made to model the system as an assumed average representation, fitted to averages of measurements made. The uncertainty may relate to the quality and coverage of measurements employed—how many sensors, where situated, how often sampled and how finely calibrated.

As previously discussed, all models are simplifications of the reality they represent. All are created with some purpose in mind, and this focus colours the simplifications made. Statisticians, quite rightly, caution against extrapolating models of data beyond the range of measurements on which they are based. The key test is one of utility and there may be qualitatively different considerations in play. A highly accurate forecast computed slowly may come too late to be useful. If fitting model to measurement requires highly granular measurement, the limited availability or high cost of making those measurements, may render such use infeasible.

In weather forecasting, new challenges have emerged, as capability, need and ambition have grown. This ambition now extends to how well we can model climate, in a way that accommodates both normal patterns and extremes and provides a useful and trusted handle on predictions made. Houghton went on to lead the meteorological service in the UK and to serve as co-chair and chair of the scientific advisory group of the UN Intergovernmental Panel on Climate Change (IPCC), from its inception in 1988 until 2002. He is quoted as having become pessimistic about the capacity of institutions at that level to act cohesively, saying that ‘If we want a good environmental policy in the future, we’ll have to have a disaster’.¹⁷ He sadly died from complications of Covid-19 infection, in April 2020, as I started to write about him in this book. His commentary, and this coincidence, adds poignancy to the connection of global weather and pandemic.

17 Interview of Houghton, ‘Me and My God’, *Sunday Telegraph* (10 September 1995).

Health care faces arguably even more complex challenges in the modelling and simulation of its personal, local and global contexts. But whereas tackling climate change requires action at a global level, medicine can (the need for global collaboration on pandemic mitigation, notwithstanding) make progress in battling the inequalities of health, through local and regional, private and public, institutions and initiatives. There is a lot to do.

Biology and Medicine—Modelling the Human Body

Having started by opining that illness and weather have much in common, and affect us all, I must now row in another direction. They are also very different. There are marked differences in what is possible: in observation and measurement, experiment and modelling, and connecting these together to deliver useful outcomes.

It would be a daunting challenge, and not very useful, to attempt a full account of the many ways in which the modelling of living systems has evolved and contributed to understanding of human biology, medicine and health care services. Writing and reading about such models can only scratch the surface and provide only partial knowledge of their purpose, what they embody, and how well they contribute. They must be used and understood in context, to be appreciated. Most published models and simulations appear rarely to have been used or further developed beyond their place and time of origin and have thus provided little of sustained value for the domains that they represent. The same ideas and work are often repurposed into new publications, over and again, as I discovered when researching my 1991 Royal Society of Medicine (RSM) talk.

Through these decades of my songline, the evolving story of computational biomedical science was, in significant measure, one of human connections made across disciplines, professions and organizations. I focus here, once more, on examples that I know best, and why and how they stood out. They are of teams and innovations I observed or participated in, as researcher, colleague or reviewer. They differ greatly in their goals and methods. Some illustrate innovation that increased understanding and moved a field forward, while perhaps of little practical application or ambition to extend beyond research into practice. Some have been transformative of ways of thinking and working, more widely. Some have accomplished their immediate goals well, only to be soon rendered obsolete; remembered as historical artefacts of no further consequence.

When considering potential topics for a PhD, in the early 1970s, I explored two broad options. One was to focus on improving medical records, with the example of devices and algorithms for monitoring patients in intensive care. Modular instrumentation systems were becoming widely

available and the use of the computer to aggregate, process and summarize the increasing amounts of data, and relate them to clinical decisions, in an accessible manner, was an interesting possibility. Instruments signalling false alarms had become an increasing problem for busy clinical staff.¹⁸ The other option arose by chance when I met John Dickinson, who became my sponsor, head of department and close colleague for twenty-five years. I have my diary of 1970/1971 to refresh my memory of early meetings with him. John introduced me to his interest in clinical physiology and proposed that we work together on modelling the human circulatory system. It is interesting to reflect that having spent twenty years pursuing option two, I switched, when John came towards retirement, to spending the next twenty largely centred on option one!

John, with Moran Campbell (1925–2004) and Jeremy Slater ('Willie', 1928–90) wrote the authoritative textbook of clinical physiology of that era. His career-long interest was in the aetiology of essential hypertension, on which he was a world authority. My parachute into medical informatics was thus an unusual one—from the modelling of physics of the atmosphere into the modelling of human circulation. From an environment of excellence in physics into an environment of world leaders in physiology and medicine. It was a scary jump into a foreign land. It was a road less travelled and, as expressed in Robert Frost's (1874–1963) poem 'The Road Not Taken', it did make all the difference.

In Europe at that time, Jan Beneken (1934–2021) was also modelling the circulatory system, working at the pioneering Netherlands Organization for Applied Scientific Research (TNO) in Utrecht. He worked there, alongside another physicist, Jan van Bemmelen, a founding father of medical informatics and its worldwide organization. Van Bemmelen originally investigated methods of signal analysis in the electrocardiogram of foetal monitoring. In Scotland at that time, Peter MacFarlane was a notable pioneer of the computerized electrocardiogram, and we touched base many times as I followed his work in analyzing and modelling the electrocardiogram (ECG) signal. Van Bemmelen became Professor of Medical Informatics at the Vrije Universiteit Amsterdam in 1973 and then, in 1987, at the Erasmus University Rotterdam, where he became Rector of the University between 2000 and 2003. He led the International Medical Informatics Association (IMIA) in its formative years. I visited TNO, and met the two Jans, when travelling to present a paper at a meeting of the European Society of Clinical Investigation, of which Dickinson was a prominent member.

18 This has not changed—the incessant din from unattended multiple alarms in nursing homes was a sad and worrying feature of my parents' care towards the end of their lives.

One of my other early encounters nearby in UCL was in life science. Working in the Biochemistry Department nearby to me in 1971 was Ted Chance, who was interested in mathematical modelling of enzyme kinetics, using newly available timesharing computer services of the early 1970s, to express and solve the equations and create a computer simulation of the reactions. Ted Chance's father was Britton Chance (1913–2010), a luminary biochemist and biophysicist of the era. As so often, the personalities connect. I was introduced to Ted Chance by Peter Sheppard, a junior doctor working with John Dickinson in the UCH Department of Medicine, who subsequently became joint supervisor of my PhD programme, with the bioengineer Stephen Montgomery. Through Ted Chance, I was introduced to the computational problems arising in the solution of non-linear and so-called stiff differential equations, which embodied markedly different rates of change and were costly to integrate, in computer processor time.

This was a problem I encountered later, when studying the behaviour of the Guyton-Coleman model and the Mac Series simulations of circulatory system physiology, which I worked on with John, as discussed in the following examples that interleaved with one another along my songline. The first describes my extended visit to the laboratory of a colossus of modelling of medical physiology of the era, Arthur Guyton (1919–2003), at his home base in Jackson, Mississippi.

Arthur Guyton—Modelling the Human Circulatory System

To consolidate my knowledge of the wider context of my PhD research with him, the run up to which I describe in the next section, John Dickinson wanted me to gain a grounding in the physiology of circulatory system function. He thought the best way to approach this would be to arrange for me to visit the team and study the work of his colleague and friend Arthur Guyton. Arthur had pioneered animal models for studying the circulatory system—it was the then predominant experimental paradigm, pursued also by John in his own research. These labs of yesterday would rightly cause a shudder, today, and be much more tightly constrained and regulated. The experimental findings were used to elucidate the building of computer models of the system.

John and Arthur had become close colleagues because of their shared interest in the aetiology of essential hypertension. In John's case, cerebrovascular hypertension was the focus; in Arthur's, it was his characterization of the 'infinite gain' bestowed on the system by the kidneys that was central. It was interesting to me that two such towering intellects, so pre-eminently aware of the multiple complexities of the control of

blood pressure, were, nonetheless, so focused on the idea of one primary mechanism. Would thinking of it as an emergent property, arising in some deeply embedded way from the very complexity of the system, be equally realistic, I wondered? Of course, I had no ideas to offer by way of traditional cause and effect-focused scientific experiment, that might be applicable in that kind of hypothetical arena, and, in any case, the money was directed elsewhere.

Both John and Arthur were clinically trained, but Arthur was no longer active in the profession. They were blessed with similar polymath and practical engineering talents and highly complementary focus and expertise in physiology and clinical medicine. They were also quite different personalities and an example of opposites attracting. Arthur was patriarchal and utterly focused on his work; he told me he had no time for the arts—preferring to relax at home with Disney cartoons. John was libertarian and loved music and opera and the good things of life. He smoked considerable numbers of cigars, sometimes just to annoy what he saw as the unduly censorious attitudes of colleagues, I suspect! It was a long time ago!

Funded and supported by long-term research grants and local patronage, Arthur and his extremely bright and motivated team of colleagues took the physiology of the circulatory system to a new level. An initially threatened and defensive scientific establishment widely disparaged them, and a wiser, unblinkered future world acclaimed them. They studied blood pressure and flow, cardiac output, interaction between vascular, interstitial and cellular body compartments and control mechanisms providing regulation, from the immediate to the long term. Arthur, himself, worked on the interrelationship of resistance and capacitance of blood vessels, blood flow and cardiac pump performance, interstitial fluid and gel dynamics and renal function, with the multifocal eyes of life scientist, clinician and engineer. It was a grand challenge fitting to the grand person.

Arthur was indeed, as I discovered, an amazing person. Initially training to be a surgeon, he contracted polio and became wheelchair-dependent for the rest of his life. He switched to the study of medical physiology and was the last person, John told me, to write a textbook covering this huge domain, single handed. He attained a chair at the University of Mississippi in his twenties—maybe his father, as the Dean of the times, played a part in that early promotion, as some said to me, but it was surely destined to be.

He designed his own crutches, to assist him to get to and fro between car and wheelchair. He adapted the car to enable him to drive. He designed and supervised construction of his family home—he and his wife, Ruth, had a large family and they were harmonious patriarch and matriarch there. Their children, his colleagues told me, were all extremely clever, too. Arthur also designed and supervised the construction of a boat to enjoy sailing on a

local lake. He was polymath and poly-competent, as thinker, researcher and engineer, a power in the historic land of Mississippi.

Recognized later in his career, and by dint of his amazing energy and focus, Arthur became President of the American Physiological Society. He was quite conservative in the Mississippi of his times, born into a leading family of the State. But also quite radical, confident and unafraid to be his own person and challenge orthodoxy, as well. Talking about science funding, he argued in one conversation with me, when meeting him and his wife for a meal they invited me to at their home, that overspending on research and focusing too much on large grants, led not just to less research productivity but sometimes to negative outcomes, taking science backwards! He showed me his most recent application for a multi-million-dollar grant to support his laboratory, saying he had already largely accomplished the goals set out there, to be pursued. He said, the grants system coerced researchers to hold back their results and use each grant to publish and promote what was already in the bag and set the scene, experimentally, for success in the next one. He quizzed me about nuclear fusion and the safety of fusion reactors—not something I knew much more about than he did!

I had raised a grant to pay for the air fare to get to Jackson, and John and Arthur paid the living costs. My diary tells me I travelled there on 8 November 1971. I even have the flight numbers, such is the longevity of the paper record! I rather doubt that electronic diaries will persist that long! Arthur put me up in the University Medical Centre Alumni House, an extremely comfortable small hotel on the campus. This was, I discovered, also in use at the time by the Governor of Mississippi, while his official residence was being renovated. There were darkened glass limousines and armed guards in constant attendance. Breakfast in the drug store across the road was eggs sunny side up, grits, orange juice and coffee. It was my introduction to the languid south of the era of Martin Luther King. Over the Thanksgiving holiday later that month, I was invited to stay on a cotton plantation owned by relatives of the family I had stayed with in Kentucky, in school days. They were hunting, shooting and fishing family folk. Half the year was American football and half was dawn to dusk cotton. I made my excuses for not getting up at three in the morning to climb into trees on the plantation, ready to shoot unsuspecting deer, grazing at dawn.

Arthur was the most welcoming and generous of hosts during my month with him. He looked after me very generously, no doubt at John's request as I was a young newcomer and scarcely worthy of so much of his time. He presided over one of the smoothest running and focused teams and departments I have ever experienced. Everyone there (Aubrey Taylor (1933–2015), Thomas Coleman (1941–2021), John Hall, Alan Cowley and Harris Granger (1945–2018)) was either already a tenured professor

or destined to become one. Physiologist and biomedical engineering colleagues of theirs of those times, whose work I also followed, were Fred Grodins (1915–89), known for his work on biological control systems, at the University of Columbia; Howard Milhorn, a mathematician and physicist, turned physiologist, doctor and author; and Vincent Rideout (1914–2003), an electrical engineer at the University of Wisconsin.

I was allocated to the different teams in the department for several days each, to see their work and hear their ideas. Arthur himself spent a lot of time with me, describing his work on an integrative model, that he and Tom Coleman collaborated on, and which Tom revised and updated throughout his career, providing it as a resource for teaching, worldwide. The Guyton-Coleman physiological models have proved of widespread educational interest and value and provided important insight guiding world-leading research. They put together all that was known about circulatory system dynamics, to create a model representation. The focus was on synthesis of knowledge about the whole system. On how different mechanisms identified through animal experiment and clinical investigation, came together within the body to regulate blood flow and blood pressure.




Fig. 4.2 Left: Arthur and Ruth Guyton at their home in Mississippi. Photographer and date unknown. Right: the extraordinarily complex circuit diagram of the Guyton-Coleman model of blood pressure regulation (1974).¹⁹ CC BY-NC.



¹⁹ YouTube hosts a fulsome tribute to the Guytons (UMMCnews, 'Remembering the Guytons: The Story of Dr. Arthur and Ruth Guyton', online video recording, *YouTube* (3 May 2019), <https://www.youtube.com/watch?v=pWKMjYd8748>).

Arthur and Tom published these models in successive papers and books. The most notable was that published in the *Annual Review of Physiology*, in 1972.²⁰ This drew together experimental quantification of functional relationships between different variables throughout the system, operating on different timescales and impacting on blood pressure and body fluid and electrolyte distribution and regulation.

This was John's special area of clinical knowledge, being relevant to measurement and observation in the management of conditions such as essential hypertension, haemorrhage, heart attack, Addison's disease and Conn's syndrome. John's rubric for testing these computer models was his immense knowledge and experience of clinical physiology and practice. He would let the models run freely and observe how changes to different parameters played out—taking blood from the circulation, dropping the pumping power of the left ventricle, disabling hormonal control of aldosterone on sodium metabolism, for example. A bit like iterative modelling of the evolution of the early universe from the Big Bang, conditioned only by current observation and theory of physical law.

Tom and Robert Hester continued the Guyton and Coleman work of the 1960s and 1970s, in the form of the Human and HumMod models. You can see Hester's talk about this venture on TED.²¹ Many models now assert their prowess by enumerating how big they are—numbers of variables and their interconnections. Modellers aiming to see what can usefully be omitted sometimes look to have given up in the context of complex biological systems such as this. The goal seems now directed more towards keeping many variables tractably within computational scope!

Over the coming years, I used our Mac Series models (which are described in the next case study, here) to analyze clinical data collected prospectively in clinical care. Could they be used to gain useful insight into the condition of individual patients, and provide useful guidance about whether, and how, to intervene to support and treat them over time? As I demonstrated in my PhD thesis, they were able to represent persuasively the sequence of changes in the circulatory system over time—after a heart attack, for example. But they were much less successful in diagnosing and predicting a pattern of disease within the circulatory system, when matched to clinical measurements made, in the way that a weather forecasting model

20 A. C. Guyton, T. G. Coleman and H. J. Granger, 'Circulation: Overall Regulation', *Annual Review of Physiology*, 34.1 (1972), 13–44, <https://doi.org/10.1146/annurev.ph.34.030172.000305>

21 Hester's TED Talk on this venture is available at TEDx Talks, 'The Most Complete Computer Simulation of Human Physiology | Robert Hester | TEDxJackson', online video recording, *YouTube* (16 July 2019), <https://www.youtube.com/watch?v=HP6wA-H1R7M>

can be used to predict the weather, when matched to measurements from weather station sensors.

The reasons for this are illuminating, more generally. The Guyton-Coleman model, though an intellectual *tour de force*, is hugely overdetermined in relation to the measurements that can be made in the real system. Insufficient amount and detail of data is available to be sampled and there are too many adjustable parameters of the model. And thus, as Feynman observed in another context—that of the modelling of elliptical orbits of the planets—the model is not useful as a theory to explain the measurements. Biological variation, subject by subject, is considerable in both healthy and diseased states. There is also considerable overlap and redundancy of feedback mechanisms within physiological systems—as we can see in the Guyton-Coleman model, where there are very many components of the model contributing to its description of blood pressure regulation. This makes the goal of identifying a unique configuration of the model, matched to an observed patient state, impossible to achieve with any confidence or usefulness. To have a chance of success, a much simpler model would be required, and its usefulness assessed in relation to a more closely defined purpose—in terms of scope of model and context in which applied. All this, I learned through trial and error, although many looking back in time, might now suggest it should have been obvious.

I am not knowledgeable about econometric models but imagine that similar parallels must exist there—usefulness in gaining understanding and interpreting ‘what if’ scenarios, in general terms, but of less value in predicting and deciding on action in response to events unfolding, in highly variable context, day by day. Here, there is also an intrinsically complex and interrelated set of dynamic relationships in play. There is, perhaps, greater ability to measure and observe data collected in the modelled, real-world system, but more limited scope for controlled experiment to inform building or tuning of the model.

Living systems have evolved in a labile manner, accumulating much redundancy of their mechanisms in the interests of maintaining homeostasis and ability to survive in diverse and challenging contexts. For the Guyton-Coleman model, this means that very many parameters of the model can be plausibly adjusted to mirror an observed behaviour of the modelled system. This is because of limitations in the measurements possible, but also a reflection of their intrinsic biological variability. Such intrinsic variability also challenges statistical modelling of data collected, in the ways that statistical methods characterize and model the distributions of interacting variables observed.

As discussed above, one way of ameliorating this problem is by sensitivity analysis of the model, revealing the extent to which its simulated

outputs vary with changes made in its defining parameters. Another is to model variables and processes as stochastic phenomena, characterized by defined probability distributions. One way or another, a level of confidence must be stated for the results reported.

I had experience in another area of research, seeking to characterize confidence in correctness of diagnostic decisions. Confidence limits on the probability of a correct diagnosis, based on the measured data and the assumptions and simplifications made in framing the model, yielded a range from zero to one hundred percent. Not a particularly earth-shattering or helpful result! Neither would be the prediction that numbers of deaths in a pandemic will lie somewhere between 5,000 and 5,000,000. Might we better accept, in some such situations, that many estimates about admittedly potentially devastating outcomes are likely to prove very unreliable, signifying that, essentially, we do not know. We should not dwell too long in modelling the unknown but recognize and communicate that we are coping as best we can, and couch our statements and decisions, accordingly.

In contrast to the science of weather forecasting, it is apparent that in modelling human physiology, both purpose and related model must be framed precisely, and tied down experimentally, to a considerable and sometimes still unachievable extent, before the range of predictions of future behaviour of the system modelled can be usefully narrowed. Such models of narrower scope and greater simplicity have proved useful in achieving practical clinical goals. This was the approach adopted by James Kirklin (1917–2004) and Louis Sheppard, whose centre I visited after leaving Jackson, and whose work I describe in another of my examples, in the section on exploratory clinical applications of models, below.

A personal note, here, about the late and great Arthur Guyton. When he came to London, many years later, to deliver the prestigious annual Harveian Lecture at the Royal College of Physicians, in the year of its five hundredth anniversary, no doubt arranged by John Dickinson in his time as a censor of the College, he asked me to sit with him, to help with set up and delivery, when needed. It was a great honour to be beside him, there, and hear his masterly presentation of his life's work. Prince Philip and some of his friends were in attendance, further along the front row. It is hard to know what they might have made of it! I stood aside outside the building, afterwards, while the Royal Party departed. Guyton heaved himself past me on his crutches, pausing to enquire why I chose to wait on royalty! He probably did not think of himself as southern states royalty, but he was!

In his eyes, his work had made understanding of the circulatory system simpler. And yet to more traditional, less practical and differently educated eyes, he made it very much more complex. It is more complex than they saw, and their perspective was not of a kind that could piece together all

that was known. Charles Sherrington (1857–1992) wrote a hundred years ago about the integrative role of the human nervous system and its central role in neuroscience. Guyton did the same for the circulatory system in physiological science, seventy years later. His Harveian Lecture was a fitting commemoration of William Harvey (1578–1657), described as the leading medical scientist of the seventeenth century for his discovery of the human circulatory system, as recorded in *de Motu Cordis* (1628), and as the founder of modern physiology. It felt, and still feels, such a privilege, for me, to have been there with him.

John Dickinson—Modelling Clinical Physiology—The Mac Series Models



Fig. 4.3 John Dickinson—one photo chosen by his family; one with Khursheed Ahmed and the author, at McMaster University (1970s); and one with the author in his office at Bart's (c. 1985), CC BY-NC.

My story of John Dickinson interleaves at many points along my songline, such was his personal importance to and for me. It has already featured in the previous example in this section, based on his connection with the pioneering work of Guyton. Here, I track back to the beginning, to place my story of him in wider context. There is thus some overlap of the narrative of these two sections. Having described, above, three sliding door moments that marked my career around the time of this picture, my Bart's office sliding door that it shows is symbolic (Figure 4.3). Indeed, my first meeting

with John proved a life-changing sliding door moment. The picture hangs on the wall of my office at home, today.

After leaving Oxford, I worked in London for two years in the medical engineering company which had funded my industrial scholarship to study physics at Magdalen. In those years, I saw the early stages of development of a new technology for automated chemical pathology laboratory testing in hospitals and the computer system being developed to control it. I also saw the operation of commercial consortia bidding for hospital development contracts in different countries.

A new beginning was then forced on to me after the company lost its way and I had to leave. A deferred national Science and Industry Scholarship award for postgraduate research, secured as a safety net on leaving Oxford, provided me with a working salary for a three-year PhD programme, which I decided to use to venture into the nascent field of computation in medicine. It was a testing and anxious time in our family life. Such a shift in career direction was full of risk, of a kind that would be hard to contemplate other than through necessity. Things could easily have gone badly wrong. In the times ahead, I stepped, nervously, through many sliding doors.

Stephen Montgomery an engineering faculty academic at UCL, who had a consulting connection with the company, generously sheltered me there while I found my feet, offering to become my academic supervisor. He kindly allowed me to work in his office in the Engineering Building in Malet Place and arranged for me to start by attending courses on the Master of Science (MSc) Computer Science programme at the London Institute of Computer Science, nearby to UCL. Keith Wolfenden (–2003), who taught the database module on the programme, had a growing interest in medical databases. He took me under his wing, there, and we kept in close touch for many years.

I had for some time been looking for a good clinical environment in which to pursue a computational medicine research project. Assisted by Bernard Lucas (a consultant anaesthetist who was also a consultant for the company in which I had worked, advising them on equipment design), I started to meet clinicians in different academic departments and attend ward rounds. For some reason, he decided that I should attend during cardiac bypass surgery in theatre, which was an eye opener to the sophistication of innovative technology deployed, as well as nearly an eye closer emotional shock for me!

Very generously, I was allowed to set up a second base for myself at the University College Hospital (UCH) Medical School Rockefeller Building on Grafton Way, working in the computer room on the fifth floor of the building, adjacent to the surgery and anaesthetics academic unit. I also got to know the hospital physics department led by John Clifton (1930–2023)

and formed a close link with him as he, too, had a developing interest in computing in medical physics. Through him, I joined the Hospital Physicists' Association and Institute of Physics, to widen my network. John Clifton became its President a year or two later.²²

The computer room housed the then very modern Digital Equipment Corporation (DEC) PDP-8 minicomputer. It doubled as the radioisotope lab and was not used and looked after as safely as it should have been, as I soon realized. A situation reminiscent of the old aircraft hangar used for testing linear accelerators, that I described in Chapter Three! Health and Safety procedures are, fortunately, very much more protective of such exposure, today.

At this point, it was by amazing chance and good fortune that I met and got to know John Dickinson, who became my luminary mentor and academic sponsor for the next twenty years. We were introduced by the biochemistry laboratory director in the Metabolic Unit downstairs, David Cusworth, who worked with Charles Dent (1911–76), a world authority on calcium metabolism. Realizing that I was struggling to find my way to a viable PhD topic, David very kindly introduced me to John, one day in December 1970. He had recently returned from an extended sabbatical break at the ground-breaking new Medical School at McMaster University in Canada, as I describe further below. I have a diary record of my early meetings with John, from that time.

It was from this fleeting encounter that our close working relationship and friendship developed in the following years, until his death in 2018—the most important and consequential of my professional life. John was one of a kind—an extraordinary mix of humane clinician and experimental physiologist. He had polymath skills and abilities (doctor, physiologist, organist, squash player, engineer—maintaining his ancient electric typewriter, car and Honda mopeds; designing and installing a central heating system and fairground organ into his family's Hampstead home, and more!). His father had been an engineer at the North London Polytechnic, later City University. Above all, he was a selfless, quite shy, and not personally ambitious person, blessed, though, with a strong drive and sense of self that marked his charisma and personality. He was truthful to a fault and able

22 Medical Physics is now an academic department within the UCL Engineering Faculty, and the now very much larger UCL Biomedicine estate extends across the entire previous UCH hospital building: to newly merged and constructed research institutes and teaching hospitals across central London, out to the Royal Free Hospital, in Hampstead, the Whittington Hospital in Archway and the Royal National Orthopaedic Hospital in Stanmore. Biomedicine has thus expanded to comprise some fifty percent of the academic constituency of the University.

to ignore irrelevant nastiness and vanity in life, to focus on what he was interested in and found fun.

At the time we met, John was a clinical senior lecturer in medicine and rising star in the academic medical unit, then led by Max Rosenheim (1908–72). Rosenheim was one of the first professors of medicine in London, subsequently President of the Royal College of Physicians and elevated through multiple civil honours to a baronetcy in 1970. He was originally appointed as a professor in the days when the University was sceptical that medicine deserved such academic status. Somewhat like professors of computer science in the 1960s. How times change!

John and Stephen co-supervised me through my extended PhD years at UCL up until 1975. John, in turn, introduced me to a world where my interest and engagement with the nearby early flowering of computer science at the Institute of Computer Science, and my mathematics and physics background, made me a good fit with his own growing research interest in medical computing, as the field was then called. He had just returned from a sabbatical at McMaster University in Hamilton, Ontario, invited there with his lovely and sparky wife, Elizabeth, by their long-term friend and colleague Moran Campbell, a doyen of respiratory medicine and physiology of the era. Moran had gone there from the academic powerhouse of the Hammersmith Hospital in London, invited by the then Dean of the new McMaster University Medical Centre, John Evans (1929–2015), to become its new Dean and plan an innovative curriculum of medicine for a new, graduate-entry medical school. Evans went on to become the ninth President of the University of Toronto from 1972–78. MUMC, as it was affectionately known, was an inspiring new building and environment, and a humming academic community. John and Moran shared a career-long interest in clinical physiology, having published the first edition of their key book in the field.²³

John determinedly stuck to his guns with his hypothesis that the primary cause of essential hypertension was resistance to blood flow in the cerebral arteries. The mechanisms regulating blood pressure in the circulatory system were a multifaceted research conundrum, challenging any human brain to fit together the breadth and variety of data from animal and clinical research, into a plausible integrative hypothesis. Over his career, and well into his active retirement years, John charted the hundreds of mechanisms that impacted on and were impacted by blood pressure in the circulatory system. Later in his retirement, he continued to review and summarize the literature, working with Julian Paton, now a Professor in New Zealand

23 E. J. M. Campbell, C. J. Dickinson and J. D. H. Slater, ed., *Clinical Physiology* (Oxford: Blackwell Scientific, 1961).

At McMaster, John got to know the other newly appointed faculty members there. He discovered the computer unit and its Hewlett-Packard (HP) computer system, installed there by David Sackett (1934–2015), the Professor of Epidemiology. It was he who nurtured, and was the founding spirit, of evidence-based medicine, a notably important and influential new way of thinking in the context of burgeoning numbers of new methods of clinical intervention and treatment, coming to the fore in the Information Age.

Computer Centres were not common features of any academic environment in those times—least of all in medicine. Sackett's department's primary use of the facility was for the development of the SPSS statistical software. He wrote about how new online resources of medical knowledge and practice could be integrated with ward-based teaching of clinical medicine. In later years, I was the international appointee in a Canadian Government team reviewing his successors' work, for the Canadian Research Council.

John became engrossed in thinking up potential uses he might make of this Hewlett-Packard 3000 computer, then state of the art. It was installed in a carefully protected, air-conditioned computer room. McMaster was intent on breaking the established mould of separate modules of life science and clinical practice in the curriculum, preferring to mix and interrelate the two from the start. It also focused on recruiting students who were already graduates, and some from non-scientific disciplines, who wished to move on into medical studies. It had a magnificent building but no laboratory facilities for students. John became interested in how computer simulation might usefully augment the medical curriculum.

He set to work to write a quite simple computer programme to simulate blood pressure and flow in the human circulatory system, which he christened MacMan. This was the start of the work he and I developed in London and at McMaster, with Khursheed Ahmed, George Sweeney, Ralph Bloch, Moran Campbell, Norman Jones (1931–2021) and others, over the coming fifteen years. It became known as the Mac Series of physiological models and, in addition to MacMan, covered body fluid and electrolyte distribution and renal function, respiration and pharmacokinetics—the latter the brainchild of Ralph Bloch. These were christened MacPee, MacPuf and MacDope and my subsequent graphics-based implementations were published in 1984.²⁴ These versions were favourably reviewed in detail in

24 As mentioned above, over the past year, I have resurrected the last published editions of the Mac Series, in the graphical form that I wrote them in the 1980s, sucking them from still extant three-and-a-half-inch floppy discs and implementing them, with the help of a very obliging hobbyist in the Netherlands

the *Times Higher Education Supplement*. They became of interest to groups eager to explore ways of minimizing animal experimentation and details were published also in this movement's journal, *ATLA Abstracts*.²⁵

I visited McMaster with John several times over those fifteen years, funded in part by money I had raised by packaging up and distributing the Mac Series programs around the world, first on huge magnetic tapes and then on floppy discs and in manuals published by IRL Press. Khursheed and his family became wonderful and very hospitable friends. I used to stay with them, while John and Elizabeth were hosted by the MUMC Dean. Over the final two years of my PhD programme, I had created versions of MacMan for the PDP-8 and on the University mainframe, and devised methods to optimize parameters for this model and its extension to cover body fluid and electrolyte mechanisms and renal function—Guyton's main interest.²⁶ What John had constructed in his effective (but difficult to disentangle) working code, I distilled into sets of functional relationships with analytical solutions and differential equations amenable to numerical integration methods. This involved substantial rewriting of the code, replacing some of the iterative solutions of short-term adaptations, such as the baroreceptor reflex adjustments in vascular tone, by analytical solutions. It also involved introducing a computational framework for numerical optimization, fitting parameters of the model to clinical measurements in myocardial infarction, both from published papers and a study based on data collected in intensive

on the DOSBox platform. This is a cloud-based resource that simulates the Microsoft PC operating system used at that time, MS-DOS. DOSBox is the only surviving platform for many computer games of the era. It is lovingly preserved, open-source, by gaming aficionados, who delight in wielding soldering irons and hoarding electronic components, to keep the hardware and software of their beloved games, alive, as I further describe in Chapter Five. It was a heart-dropping moment for me when, after some struggles with indecipherable error codes, a tiny code patch arrived from the Netherlands and all four programs were reincarnated, in a flash!

- 25 C. J. Dickinson, D. Ingram and K. Ahmed, 'The Mac Family of Physiological Models', *Alternatives to Laboratory Animals*, 13.2 (1985), 107–16, <https://doi.org/10.1177/026119298501300204>.
- 26 I vividly recall the first outing of these early versions of the programs, at a 1971 meeting of the Physiological Society, at UCL. John showed the MacPee programme, connecting with it through a slow, ten characters per second teleprinter, and I used the PDP-8's much more dynamic oscilloscope display, transferred by a novel scan converter device to a television screen in the nearby lecture theatre, to demonstrate MacMan. Its simulations of the haemodynamic consequences of blood loss, cardiac insufficiency or vascular hypertension attracted a large audience. John was discussing much more detailed topics involving renal and interstitial fluid dynamics, on the other side of the room, and came over with Lord Rosenheim, as he then was, to see what the excitement was all about!

care unit (ICU) patients at UCH. Numerical optimization methods of this kind were in an early stage of evolution, and I experimented with the main exploratory approaches. At the time, there was much research into computational solutions for stiff systems of differential equations, and I implemented these methods, too, to replace the quite simple iterative program loops whereby many, including Dickinson and Guyton, were simulating these dynamics.

My changes speeded computation in some parts of the simulation, but the non-linear differential equations were too unwieldy to be integrable with available numerical integration packages of the time. John, unaware of these issues, had ploughed his way through to the endpoint he had sought, which served his purpose well, by using numerical smoothing functions to ease the 'stiffness' of the modelled changes as the simulation evolved forward in time, ensuring the calculations did not become unstable, and the solutions oscillate. Thankfully, the body, as a distributed system, does not normally face unstable oscillations of this kind, although instability in its biological control systems does arise in some contexts—such as in periodic breathing of the immature lungs of premature babies.

I mined the literature to unearth published sources of clinical haemodynamic data for patients treated for myocardial infarction—these were very few and far between—and wrote numerical optimization procedures to match the model to these, adjusting model parameters according to patient height and weight and optimizing the cardiac pump performance parameter to match the published blood pressure, heart rate and cardiac output values. With considerable effort, I collaborated with John's clinical house officer of the time, and ICU nursing staff, to collect data from their patients, recorded over ten days following heart attack. I used further numerical optimizations to match the model to haemodynamic and body fluid and electrolyte measurements, to estimate the cardiac pump performance, day-by-day. As an independent check, I correlated these estimates with enzyme studies, used routinely to assess extent of cardiac tissue damage. In this way, the time course of damage and recovery of heart pump performance was charted, and the results checked against the enzyme picture.

All this was very laborious, a considerable burden on routine clinical care, and very approximate—and completely useless in any practical sense! But I did learn a lot that I drew on in subsequent years, for example as a reviewer of major EU modelling research projects, such as the Oncosimulator developed by Georgios Stamatakis and Norbert Graf, in the EU Advancing Clinico-Genomics Trials on Cancer (ACGT), p-Medicine and Computational Horizons in Cancer (CHIC) projects, as further discussed, below, in the final example of this section.

*Denis Noble and Peter Hunter—Virtual Physiological Human and
In Silico Medicine*

In the early 1970s, physiologists were exploring the mechanical performance of human muscle fibres and the stimulation of their contraction that gives rise to the rhythmic pumping action of the heart ventricles. From the time of Ernest Starling (1866–1927), the heart pump performance had been characterized as a function curve, relating the pressure from venous blood entering through the right atrium to the rate of flow of blood into the systemic circulation via the aorta (cardiac output). One of the first books John Dickinson gave me to read on clinical physiology was Guyton's monograph, *Cardiac Output and Its Regulation*.²⁷ This was a topic I also encountered at conferences addressed by the UCL and Oxford physiologist, Denis Noble, whose interest was at a more granular level, recording muscle tension and length in laboratory experiments, twitching muscle fibres electrically and observing the effect of perfusion with catecholamine hormones, mirroring the manner in which the body stimulates and regulates cardiac performance. Other researchers that I got to know, such as Derek Gibson (–2021) at the Brompton Hospital in London, a founding father of echocardiography, later extended this knowledge into a three-dimensional model of the left ventricle, integrating individual muscle fibre mechanics into the muscle wall dynamics, shaping the contraction phase that ejected blood from the ventricle and the following relaxation phase that allowed incoming blood flow from the pulmonary circulation to re-expand it.

An increasing number of research teams experimented with physiological models. Noble's work extended into the modelling of cardiac cell metabolism, to study metabolic aspects of cardiac disease. His work was ground-breaking in the methods developed for modelling across domains of physics, chemistry, biology and physiology of the cell and cellular transport. The Guyton-Coleman model was essentially an integrative assembly and simulation based on experimentally derived function curves, from clinical and physiological studies of the circulation. Noble dug deeper into the modelling of function within the cell, developing and generalizing modelling methods to assist integrative understanding of more complex systems of biological and clinical science, bridging from genomics to immunology. In Oxford, he teamed up with Peter Hunter in the Oxford Engineering Department, which had long pioneered bioengineering, from the time of Brian Bellhouse (1936–2017) in the 1960s, developing artificial heart valves. I remember Bellhouse as a fellow of Magdalen College in my

²⁷ A. C. Guyton, C. E. Jones and T. G. Coleman, *Circulatory Physiology: Cardiac Output and Its Regulation* (Philadelphia, PA: Saunders, 1973).

time there. Lionel Tarassenko, a subsequent head of the same department, has pioneered signal analysis methods applied to clinical data analysis and decision making.

In the USA, through the 1980s and 1990s, the Visible Human Project was established in support of the study of anatomy. This interdisciplinary field of endeavour extended into bioscience and medical science, more widely, under umbrella terms such as systems biology and systems medicine. Hunting for a unifying initiative, in 1997 the International Union of Physiological Sciences (IUPS) launched plans for its Physiome Project. The scope of this was well exemplified by the Oxford–Auckland Cardiac Physiome Project. In short order, genome led to epigenome, transcriptome, metabolome and biome, propelled by the explosion of bioinformatics data, both experimentally and in clinical contexts. The quest for integration was championed by one of Guyton’s key team members that I had spent time with in Jackson, Alan Cowley, who had subsequently moved to a chair at Milwaukee. By then President of the IUPS, he wrote in 2004: ‘now is the time to begin building the scientific infrastructures that will enable an integrated understanding of the function of complex organisms and chronic diseases’.²⁸

The European Union has championed brave attempts to implement this mission and it is likely on a very long runway; possibly even longer than that of controlled nuclear fusion! Building on ideas developed in the Physiome Project, EU research funding was invested in the Virtual Physiological Human (VPH) Network: ‘to enable collaborative investigation of the human body as a single complex system’.²⁹ This brought together academia, clinical practice and industry, to explore how deeply and widely this synthesis might run. The present day International Virtual Physiological Human Institute was established to take this work forward. Noble’s colleague, Hunter, then working in New Zealand, became a leading light.

The term, *in silico* medicine, arrived, championed in the work of another group that I became close to in the early 2000s, the Advancing Clinico-Genomics Trials (ACGT) on Cancer initiative of the European Union. One question has recurred throughout VPH-style research: How well can the advancing science of *in silico* medicine connect with the practicalities of improving clinical care? For example, in confronting the panoply of experimental and clinical data from a patient with nephroblastoma, and using these to match a model of chemotherapy, Graf and Stamatakis and their colleagues and teams tracked the impact of treatment on a

28 IUPS Newsletter, 7 (September 2004).

29 STEP Consortium, *Seeding the EuroPhysiome: A Roadmap to the Virtual Physiological Human* (n.p.: STEP Consortium, 2007), p. 2, https://www.vph-institute.org/upload/step-vph-roadmap-printed-3_5192459539f3c.pdf

tumour's growth and its sought for reduction and elimination. Their work in the ACGT project and its successors is described in my next and final example of pioneers of innovation in modelling the human. In the choice of treatment regimens and their outcomes, there are harder challenges, where the uncertainties that I had come up against in my comparatively tiny PhD study, now five decades ago, still persist. Intrinsic biological variability, the uniqueness of each clinical problem and of the people and teams facing and addressing it, and the wider context in which the treatment applied plays out all interact and may confound feasibility of transition from the science to its effective and useful application in clinical care.

In silico medicine is a natural science and it is natural to be curious about and study it. In a clinical context, it is another specialism and, as with all such specialization, it has the potential to integrate and inform, and to fragment and confound. As has unfolded since the beginning of my songline—from mathematics and physics into biophysics, physiology, medicine and health care in everyday contexts—even the most sophisticated and well-endowed clinical environments have struggled to keep abreast of all that can now be measured and analyzed. In health care, information technology has enabled the advance of the best of the best. It has not achieved comparable impact at the other end of the spectrum of excellence—the worst of the worst. In some respects, through wasteful expenditure on pursuit of unrealistic goals, and commandeering of resource for elite priorities, it has arguably caused and allowed the gap, relatively, to widen.

Interest in what now comes under the banner of VPH research has persisted through six decades, since the 1960s,³⁰ seeking to help establish and consolidate progress in what is still a rapidly evolving field. A notable early team, whose work I collected, was that of Ed DeLand. He obtained a mathematics PhD at UCLA, at age thirty-four, and worked in the innovative environment of the RAND Corporation in Santa Monica, from the 1960s. When I came across his work, he was modelling red blood cell membrane transport equilibria, based on computation of the Gibbs function of statistical thermodynamics, with a view to applying the model to the interpretation of clinical laboratory measurements. The model was impressive in its capacity to predict these equilibria very precisely, in changing conditions. DeLand was a colleague of Thomas Lincoln (1929–2016), the clinician I mentioned in the Introduction as someone I met in London in my early PhD days with John Dickinson and Stephen Montgomery.

30 W. Ware, *RAND and the Information Evolution: A History in Essays and Vignettes* (Santa Monica, CA: RAND, 2008), https://www.rand.org/content/dam/rand/pubs/corporate_pubs/2008/RAND_CP537.pdf

The breadth of ambition for the Virtual Physiological Human has brought new levels of synthesis and insight into play, just as Sherrington did, in 1906, when he first envisaged the integrative character of the nervous system. When researching this connection, I alighted on Sherrington's reply, when asked about the purpose of his university, which still echoes a hundred years on:

After some hundreds of years of experience, we think that we have learned [...] how to teach what is known. But now with the undeniable upsurge of scientific research, we cannot continue to rely on the mere fact that we have learned how to teach what is known. We must learn to teach the best attitude to what is not yet known. This also may take centuries to acquire but we cannot escape this new challenge, nor do we want to.³¹

A University discipline teaching about the unknown—Dick Cheney might have been pitching for a faculty position! The Information Age is expanding the domain of what is known and creating new dimensions of unknowing. The physicist, Max Born (1882–1970), once poetically described scientific discovery as a process of opening windows onto the stars, that simultaneously increases our vision of the unknown—that has stuck with me, not having read his works since my college days. Perhaps the common ground being sought for the Information Society is as much about coping with the unknown as it is about sharing of the known. One wonders what Sherrington might have made of a Novacene concept of computers that can know, where humans do not and cannot.

Georgios Stamatakos and Norbert Graf—The Oncology Simulator

In the early 2000s, I was appointed to several UK medical and engineering and physical sciences research council boards, and, as with my work for the NHS and EU, was involved in oversight of 'e-science', as it came to be termed. The EU boards were especially interesting and satisfying because they provided continuity over many years between reviewers and research teams, drawn from across academia, health care and industry. The EU Commission worked us hard and for little financial reward, but the relationships both ways became deep and enduring. One such incredibly hard-working and creative initiative, that I especially enjoyed working with, was the Advancing Clinico-Genomics Trials on Cancer (ACGT) project.

31 J. Eccles and W. Gibson, *Sherrington: His Life and Thought* (Berlin: Springer International), p. 24.

This set out to build a master ontology of data used in management of clinical trials, linking with advances in genomics science. The cancers studied were nephroblastoma, breast cancer and leukaemia. In parallel, a computer model and simulation of chemotherapy, christened *Oncosimulator*, was pioneered and pushed forward with amazing energy and commitment by Georgios Stamatakos and Norbert Graf. Through dedication and hard work over many years, they became much-admired pioneers in the advance of *in silico* medicine.

Norbert was the inspirational and hard-working clinical leader of the ACGT team. Biomedical scientists from leading cancer centres across Europe were also members. Norbert encouraged the initially sceptical clinical academics, himself included, to build a very fruitful and data rich environment for Georgios and his team. Manolis Tsiknakis and Mario Cortelezzi from the Hellenic Mediterranean University, led and held together the technical teams, including researchers from the prestigious Fraunhofer Institute in Germany and strong industry and health care institution partners. These people, and their enthusiastic colleagues of all ages became a memorably motivated community. Olle Björk, an oncologist from the Karolinska Institute in Stockholm and head of the Barncancerfonden [Swedish Childhood Cancer Foundation], and Elena Tsiporkova, an incisive biomedical scientist and mathematician, were colleagues with me in the review and advisory boards appointed for the project and its p-Medicine and CHIC successors, over the following ten years.

The master ontology was an ACGT project workstream pursued by a subgroup linked closely with Barry Smith's international biomedical ontology movement, mentioned in connection with the world of medical knowledge bases, in Chapter Two. The project loyally followed this lead as it was the vision on which the grant had originally been made. It nearly became the project's undoing further downstream, as the messiness of clinical reality came up against the philosophical drivers of formal ontology. The focus on the master ontology initially pulled activities together and, as it creaked against the realities of harmonizing real data and real database implementations, nearly pulled them apart. This goal had gradually to be deemed to have failed and was downgraded in priority.

Norbert was the powerhouse for development of a tool to formalize clinical trials design and consistent collection of data across sites. The engineers loyally tackled data aggregation and descriptive metadata, across the different institutions and sites. The ethico-legal framework for the sharing of data was a notable success story—ably and professionally handled by Nikolaus Forgó, a professor of law and subsequently Dean, at the University of Hannover.

ACGT and its successor projects were an outstanding conjunction and collaboration of many teams and aspirations. It received many tens of millions of euros of European Commission investment, over three rounds of five-year funding, to create and sustain its teams and environments. As with all such projects, it produced many volumes of written reports, keeping their reviewers awake to the small hours, preparing for the regular three-day review meetings!

The clinical goals were well-expressed, and the clinicians and life scientists involved were already world authorities, joining research across many centres, in Europe, America and Japan, including my twin alma maters, UCL and the University of Oxford, in the UK. There was great success in pulling together clinical and genomics science data and building informatics infrastructure for research on different cancers. These wider connections allowed the research outputs to find a place in many international conference proceedings and journals. The range of expertise and the age and gender balance and culture of the teams was excellent. It proved the fundamental importance of good environment—not just as something nice to have, but as essential.

Georgios emerged as a world figure in *in silico* multiscale modelling—scaling from modelling chemotherapy at a molecular level, to its impact on tumour angiogenesis and growth, and tumour destruction by chemotherapy. The model utilized finite element methods from engineering science to model at multiple levels and scales of cellular and organ function. The Oncosimulator work brought different groups of modellers into a productive conjunction—mathematicians at Oxford working on analytical models of angiogenesis and the finite element approach of Georgios's team in Athens working on cellular mechanisms. Norbert, already a world figure in paediatric nephroblastoma research, and running a wonderful paediatric cancer service and caring community in Saarland, was the father figure across all domains of the project—an energetic, humane and infectiously enthusiastic clinician who also held together the project's links in Japan and the USA.³²

32 As I write today, the Pfizer/BioNTech Covid-19 vaccine is announced to be approved for use in the UK. The earlier press releases had connected its inventors, two Turkish doctors working in Germany, to the University of Saarland, Norbert Graf's University, where they had met as students. Norbert became a dean of this medical school. He told me in an email last week, that he knew them well from the time they collaborated on research with him. Another inspiring story of international connection between people and teams, and the environments that enable them and their work to prosper.

Exploratory Clinical Applications

The examples now move on further, to models of clinical decision making and intervention in the everyday practice of health care. In the early 1980s, I worked with my Mac Series colleague at McMaster, Ralph Bloch, to draw together and edit a two-volume collection, *Mathematical Methods in Medicine*, to accompany the Wiley six-volume *Handbook of Applicable Mathematics*. The first volume, focused on statistical and analytical techniques, included a chapter on clinical decision analysis.³³ The second volume focused on clinical applications.³⁴ In this section, the examples are drawn from a range of clinical applications of modelling that I connected with in the first half of my career in medical and health informatics.

Statistical Modelling of Diagnosis—The Royal College of Physicians of London Computer Group

Classification and statistical analysis of clinical observations and measurements has long been used to segment patient populations within diagnostic and therapeutic groupings. These have ranged from simple scoring to intricate mathematical methods. Mathematical methods for analyzing and guiding clinical decision making took root in the 1970s, bringing together academic departments of medicine, statistics, and psychology, in the context of everyday clinical practice. Bayesian statistical methods were championed by Dennis Lindley (1923–2013), head of statistics at UCL, and later by Adrian Smith, whose later career was as Vice-Chancellor of Queen Mary University of London, UK Government Chief Scientist and Director of the Alan Turing Institute, established close to UCL as a national initiative in data science. I first came across Adrian Smith's work in the early 1980s, when he published a novel method for forecasting renal allograft rejection.³⁵ This was based on mathematical analysis of serial measurements of creatinine clearance in urine, using a Kalman filter technique, and showed that the event could be predicted up to several days before clinically manifested in

33 D. Ingram and R. F. Bloch, ed., *Mathematical Methods in Medicine, Part I: Statistical and Analytic Technique* (Chichester, NY: John Wiley and Sons, 1984).

34 D. Ingram and R. F. Bloch, ed., *Mathematical Methods in Medicine, Part II: Applications in Clinical Specialities* (Chichester, NY: John Wiley and Sons, 1986).

35 A. F. M. Smith, 'Change-Point Problems: Approaches and Applications', *Trabajos de Estadística Y de Investigación Operativa*, 31.1 (1980), 83, <https://doi.org/10.1007/BF02888348>; I. M. Trimble, M. West, M. S. Knapp, R. Pownall and A. F. Smith, 'Detection of Renal Allograft Rejection by Computer', *BMJ*, 286.6379 (1983), 1695–99, <https://doi.org/10.1136/bmj.286.6379.1695>

patients. This work has an interesting connection with the contemporary unfolding story of artificial intelligence, as discussed in Chapter Two.

Wilfrid Ingram Card (1908–85) was head of gastroenterology in Glasgow and a leading figure in the profession in the 1960s, with a special interest in medical education. I remember him visiting UCL and speaking at the Royal College of Physicians in London in the 1970s. The Glasgow University Record describes his distinctive contribution: '[...] in October 1966 he was appointed to a personal Professorship at the University of Glasgow in the Department of Medicine by Mathematical and Statistical Methods [...] In 1967 he published a classic article entitled "Towards a Calculus of Medicine" (*Medical Annual* 1967, 85, 9-21)'. Henrik Wulff credits him with 'a new paradigm of clinical thinking'.³⁶

Card had collaborated with Lindley in further developing his ideas for mathematical modelling of medical decisions. In this they coined the term 'indicant' to cover clinical measurements and observations used to confirm a diagnosis, and linked these indicants with candidate diagnoses, using a Bayesian model of probability. There were several competing, but not necessarily conflicting perspectives on the nature of clinical diagnosis at that time, leading to different approaches to its formal study. Was it based on the weighing of statistical probabilities connecting what was observed and measured with potential underlying causes? Was it a human acquired skill based on recognition of patterns in these indicants? Was it a Popperian hypothetico-deductive method, acquiring evidence and gradually homing in on a conclusion about the causes of a presenting complaint, by ruling out alternative possibilities? What kind, specificity and sensitivity of evidence and model was needed? How did clinicians weigh multiple kinds and amounts of such evidence, in practical context, often when working under extreme pressure in seeking to save lives? Moran Campbell, on a sabbatical break from McMaster, took an interest in the topic, advocating for the framing of diagnosis of disease as a hypothetico-deductive method.³⁷

The experiments conducted with these different approaches, exploring their technical and clinical contexts, and the accompanying debates within associated scientific and professional communities were the focus of the Royal College of Physicians of London (RCP) Computer Group. This was established through connections of Card with clinical professional luminaries of the era, including the RCP President from 1977–83, Douglas

36 'Measuring Gut Feelings: The Scientific Basis for Clinical Medicine', talk given at a Memorial Festschrift, RCP London, 17 June 1986.

37 E. J. M. Campbell, J. G. Scadding, and R. S. Roberts, 'The Concept of Disease', *BMJ*, 2.6193 (1979), 757–62; E. J. M. Campbell, 'The Diagnosing Mind', *The Lancet*, 329.8537 (1987), 849–51, [https://doi.org/10.1016/S0140-6736\(87\)91620-5](https://doi.org/10.1016/S0140-6736(87)91620-5)

Black (1913–2002). Membership of the group was drawn widely from across UK academic medicine and statistics. David Spiegelhalter, (an alumnus of Oxford and UCL, and student of Adrian Smith, who went on to become President of The Royal Statistical Society and Professor of Public Understanding of Risk at Cambridge, and a nationally prominent advisor and commentator on statistical aspects of public policy) teamed up with Robin Knill-Jones (a Glasgow gastroenterologist working alongside Card) to develop a Bayesian model for differential diagnosis of acute abdominal pain. A system they christened Gladys, standing for Glasgow dyspepsia! Abdominal pain was readily recognized, but what character of pain, how exemplified and with what underlying aetiology? And in context of clinical management, what presenting situations required urgent action and what manifestations were deemed non-specific of underlying disease, and best to be watched over, to see how they evolved.

Another luminary figure of that time, Timothy (Tim) de Dombal (1937–95), also shone there. He was a surgeon at Leeds, who, with a mixture of persuasive charm and iron determination, master-minded large-scale collections of data relating to problems of differential diagnosis in gastroenterology. His field trials were conducted first in the hospitals in which he worked, collecting cohort sets of a standardized group of indicants for the most common diagnoses of acute abdominal pain, and then using these to predict the most likely diagnoses for newly presenting patients. The work extended to many countries under the auspices of the World Gastroenterology Organization and the International Federation for Information Processing, where Tim played leading roles. He was less well supported near to home—seen as neither a proper surgeon nor a proper statistician, perhaps. It seemed that he rather relished that notoriety, and he was better respected and anchored in the wider world community in which he worked. But it cannot have been an easy mix of roles and reputations to sustain.

Tim's method for modelling the diagnosis of acute abdominal pain combined standardized methods of clinical data collection and a simple Bayesian method for estimating the probabilities to be assigned in linking an observed patient profile with the six diagnoses he chose to work with as explanatory of the pain. For each of these six diagnoses, he collected learning sets of a hundred cases—in some rare conditions only fifty—with which to calibrate his method. The approach was based on assumption of conditional independence of these data indicants, when calculating the probabilities assigned. This assumed property of the data was recognized not to hold, but the limitation was glossed over. It functioned more as a scoring system, as the highest estimated probability was chosen as indicative of the correct

diagnosis and the actual probabilities assigned were not assessed for their accuracy.

Nonetheless, the method (sometimes referred to as 'Idiot Bayes'!) produced some persuasive evidence that it could assist clinical management, for example in avoidance of negative laparotomy through misdiagnosis of non-specific abdominal pain. The method was straightforward to compute, in comparison with the more complex Bayesian models of these probabilities that were used by others. Implementation of these mathematically more rigorous Bayesian models proved difficult and lacked sufficient range and granularity of available clinical data, on which to calibrate the model and compute its predictions. The computational resource required increased steeply alongside the increasing complexity and variety of clinical measurements arising in the modelled domain.

Tim faced professional opposition from some senior clinical colleagues who disparaged the importance of his work and obstructed its adoption, meaning that his progress had to be very hard-won—a not uncommon scene for such pioneers! The message that a program algorithm might, to some quantifiable degree, parallel human diagnostic skill and performance, was an uncomfortable idea for many, if not most. It challenged the self-belief of practitioners and brought the nature of their roles and contributions under a new spotlight. Tim was feisty, as well as testy, at times, but created a loyal team around himself and battled on, with notable charisma and *sangfroid*. Sadly, he died very young from unsuccessfully managed complications of cardiac surgery.

I got to know Tim and tracked his work for several years. He shared his data with me, as I tried to support a colleague clinical lecturer in the Department of Medicine at Bart's, Huw Llewellyn, with his MD (Doctor of Medicine) research project. Huw had envisaged a novel method of reasoning about diagnosis, based on manipulation of mathematical sets, expressing linkage of indicants with diagnoses. To help him, I wrote software to implement his approach, as a way of testing and presenting his ideas in his MD thesis. His reasoning was not wrong, it seemed to me, but the method did not provide useful clinical results in the work we did together. It estimated a range of probabilities whereby the observed indicants could be linked with candidate diagnoses and suggested a rationale for step-by-step investigation in reaching a diagnostic decision. The confidence limits it could assert were quickly too wide-ranging to have practical significance in guiding decision. Huw faced difficulty in defending the ideas presented in his MD thesis. He battled on and won through to the award of the degree, and still pursues those ideas, today. He was a strong and goodhearted professional doctor, hugely proud and loyal to the Welsh valley origins of his family. I noted, as I made checks in writing this section, that he is now a co-author of the

prestigious *Oxford Handbook of Clinical Diagnosis*. Clearly, he was talent, grit and staying power personified! I watched a short video connected with the promotion of his book. He sounded exactly the Huw I knew.

In the discipline of psychology, there was continuing focus on how humans could structure, refine and improve decision making. Was there a structured tree of decision points, traversed systematically according to questions asked and findings reported? Did humans learn their skills through experience whereby a decision could be made in a 'blink', based on a recognized pattern of observation and measurement? How many variables could the brain hold in mind, when interpreting and deciding the diagnostic significance of all this information? Some studies had suggested the magic number of seven variables, after which human capacity, having initially steadily improved, started to decline, and quality of decision making likewise declined. The phenomenon of information overload became of interest in the study of clinical risk, pioneered at UCL by my colleagues, Charles Vincent and Pippa Bark, in the Psychology department. Pippa subsequently joined my department and Charles moved to Imperial College London, working under Ara Darzi, who also pioneered innovations in surgical robotics and clinical informatics.

Many people contributed to the RCP Computer Group throughout its life, with Jeremy Wyatt becoming a notable participant, writer and advocate of the work. The group fulfilled a valuable coordinating focus but its role as a national forum declined as key personalities moved on and momentum was lost. It was replaced by a more generically focused activity, concerned with the impact of information technology on medical records and terminology. This new unit was led by colleagues and friends of those days, Martin Severs, John Williams and Iain Carpenter, who put tremendous energy into seeking national consensus on the structured content of clinical records. They were doughty pioneers at the coalface of informatics in clinical practice, health policy and research, for several decades.

Martin maintained his clinical professional work in care of the elderly, and as Dean of faculty at his local University of Portsmouth. He led the formative years of the International Health Terminology Standards Development Organization (IHTSDO), as described in Chapter Two and went on to become Medical Director of the NHS Health and Care Information Centre, subsequently NHS Digital and NHSX. These organizations metamorphosed and changed their name many times over these years, as described in Chapter Seven. John and his wife, Jane, pioneered a gastroenterology record-keeping system in Swansea, and John went on to lead Research and Development for the NHS in Wales. He created and led the RCP Informatics Unit that grew from the Computer Group. This works collaboratively with the NHS, still, in standardization of the clinical content of care records.

Operational Research–Clinical Trials and Epidemiology

In the mid-1980s, UCL established a Clinical Operational Research Unit (CORU). It was led by Ray Jackson, a mathematician who came to the fore in battling the orthodoxy of clinical trials methodology, proposing that clinical outcomes from treatments could be modelled mathematically, drawing on operational data from care records of populations of patients, thereby lessening the requirement for costly, independently established and extensive randomized trials. This proposition did not persuade minds, but Ray was a wily politician as well as a skilled mathematician, and he succeeded in broadening the academic focus of CORU, extending mathematical methods across health services research and biomedical informatics. Ray's contributions to government operational research are recorded on a civil service website.³⁸

It was an era when mathematics and engineering were both pitching for ownership of the fledgling discipline of computer science, not unreasonably given its founding fathers. Alan Turing (1912–54), Alonzo Church (1903–95) and John von Neumann (1903–57) were mathematicians, and the field was lifting off, powered by advances in semiconductor physics and electrical engineering. Operational research had a clear pedigree in mathematics and CORU remains in that UCL faculty to this day. With Steve Gallivan and Mark Leaning, colleagues of mine from those times, Ray initiated a programme of mathematical modelling, widely across biology and medicine. He was a battling individualist, and a very able and successful one.

I first met the world of mathematical models in medicine, in reading work of the mathematician, John Maynard Smith (1920–2004). He became a professor of mathematics, much interested in its applications in evolutionary biology and genetics, having started as an aeronautical engineer. Pure mathematics is a pursuit that treasures its isolation and sometimes almost scorns its wider application, relevance and importance! It does not do 'trade' and exists on a level of abstraction and within community that can defend an ivory tower-like existence and perspective. That aside, where would we all be without it? Maynard Smith brought mathematical insight to the centre of biological and medical discipline, as Ian Stewart has, today.

My closer contact with mathematical epidemiology arose through my work for the Wellcome Foundation in the late 1980s, on computer-assisted learning—my main interest in those years. A project had been established to build an educational resource based on the Henry Wellcome (1853–1936)

38 M. Hudson, *A History of the Government Operational Research Service 1968–1980* (n.p.: GORS, 2018), p. 20, http://www.operational-research.gov.uk/public_docs/history-of-gors.pdf

collection of historical items, known as the Wellcome Museum of Tropical Medicine, later incorporated into the Wellcome Collection Museum of Medical Science, and now to be closed. The Wellcome Foundation had traditionally had strong worldwide roots in research in this domain. The Foundation's original major shareholding in the Burroughs-Wellcome drug company was sold, and the proceeds reinvested more widely into the recast Wellcome Trust. The Trust has enjoyed stellar subsequent growth of its investments. This enabled its rapid development as a major international funder of research and public awareness of biomedical science. It was led through this transitional period from 1991–98 by its then Director, the parasitologist Bridget Ogilvie, who I worked with closely at that time.

Among its luminary Trustees of the time was the mathematical biologist Roy Anderson, author, with Robert May (1936–2020), of the highly cited book, *Infectious Diseases of Humans: Dynamics and Control*.³⁹ May went from mathematical epidemiology to become Government Chief Scientist and President of the Royal Society. Anderson led departments at both Imperial College and Oxford and succeeded the formidable former Chairman of GlaxoSmithKlein, Richard Sykes, when becoming Rector of Imperial College. He reviewed the museum project that Bridget had asked me to look after when it had got into difficulties, after conflict within its senior team.

Anderson was one of an extremely talented grouping of mathematical biologists and tropical disease researchers, working closely with Wellcome. Neil Ferguson, who studied physics at Oxford and migrated into mathematical epidemiology, became a senior member of Anderson's team, at Oxford and Imperial. He was prominent in the modelling and simulation controversies that populated academic and public discussion of policy options for containing the Covid-19 pandemic. My UCL colleague head of academic medicine of earlier years, Patrick Vallance, presided over the furore, as the Government Chief Scientist of the time. Chris Whitty, a University College London Hospitals (UCLH) physician and epidemiologist, acted, also remarkably calmly, as Chief Medical Officer (CMO) alongside him. The preceding Chief Scientist, Mark Walport, also a former chief of the Wellcome Trust, and formerly also a chief of medicine, at Imperial College, acted in support to help carry the public load. The preceding CMO, Sally Davies, was also a physician from Imperial College. They all needed to know and trust one another well, and they gave sterling service.⁴⁰

39 R. M. Anderson and R. M. May, *Infectious Diseases of Humans: Dynamics and Control* (Oxford: Oxford University Press, 1992).

40 Sparkling academic and medical careers often lead to high office in universities and leadership of government and international agencies. Some migrate into academic leadership from careers in industry and some in the reverse direction,

In the Covid-19 context, a wide range of assumptions, simplifications and models has been in play among proposing and competing groups of modellers: about susceptibility to and degree of infection, transfer within populations, and the course of illness arising—unnoticed, mild, serious and fatal. Experimental data that would help to pin these down more precisely is only slowly possible, dependent on clinical interventions, reliable testing methods and population level sampling of virus and viral response. Predictions about the degree of infection, interactions with containment measures, vaccination and treatment regimens, and expected clinical outcomes have been wide-ranging. Such predictions must, necessarily, be hedged with caveats that render them difficult to interpret as a basis for what are major policy decisions, balancing the safeguarding of personal safety and livelihood alongside capacity of the caring services.

My Polish nephew, who is working at a high level in investment banking in New York, believes the world has committed catastrophic error of overreaction, in response to what was, he admits, a terrible pandemic. Different assumptions used in designing and calibrating models, and fitting them to contemporary measurement and observation, led to radically different predictions and strategies. Countries that took early and more drastic action on containment seem to have succeeded better at first, than those where policy was gambled and nuanced more in favour of ‘play it by ear’ and ‘wait and see’. Events in such as New Zealand and China, have more recently been reversing this trend. In whatever way a pandemic is modelled, the policy choices made in seeking to contain it involve choices about values. Information has been assembled, communicated and weighed, worldwide and in new ways, in this pandemic. Being the first of the modern scientific era and Information Age, it seems destined to become a set piece of future analysis by epidemiologists and historians.

and to and fro, vitalizing connections. Those who are London-based or close to London appear to have a stronger hold on levers of political power and influence in these echoing hallows. I owe it to my NHS-employed family members, further north, to make that observation, which they would, no doubt, have me say, more loudly! Bridget Ogilvie once remarked to me about the three-fold talent of leadership she observed in academia—brain, political nous and human insight, I seem to recall. This categorization reminds me now of the philosopher Immanuel Kant (1724–1804), who wrote about the threefold resources—intelligence, power and money. He also wrote about the three degrees of evil: frailty, impurity and depravity or perversity! Everything in the world maps to three dimensions, it seems, but time always tells—especially so with simulations! *Zobaczmy!*

Electrophysiology

From the 1970s, electrophysiology was expanding into the computer domain with methods for signal analysis of the electrocardiogram (ECG), accumulation of large data banks, and automatic characterization of disorders of cardiac rhythm. This led to new methods for their detection, mitigation and correction in patients. Jan van Bommel, in Utrecht, Peter Macfarlane, in Glasgow, and Bruce Sayers (1928–2008) and Richard Kitney at Imperial College, were pioneering colleagues of those times. There was parallel work in other areas of signal analysis, notably study of the electroencephalogram (EEG) and its importance in brain science and clinical neurology.

Physicists and electrophysiologists developed what were termed inverse methods for the modelling of cardiac function. These worked backwards from electrical signals propagated through the chest wall and detected by an array of sensors fixed to the skin around the chest surface. The signals were combined and analyzed with software, to infer the structure of a postulated model of the beating heart as a transmitter of electrical signals. This was akin to the methods of axial tomography image reconstruction, used in imaging devices, such as X-ray, MRI and PET scanners.

Modelling Ventilation Management

Among the Mac Series models, MacPuf achieved the greatest success. It was envisaged by John Dickinson, in close collaboration with Moran Campbell and Norman Jones at McMaster University, both titans in the field of respiratory physiology. Other teams developed simpler programs that modelled limited aspects of lung function and achieved success, for example in matching these to measurements in gaseous anaesthesia and as tools for practical classes exploring lung mechanics and gas exchange. The idea of MacPuf was more ambitious—its goal was to model the cardiorespiratory system of the body, in clinical context.

MacPuf found its way into an unusual book, explaining line by line the clinical and physiological rationale of the program. It was an adventurous idea and won appreciative reviews—a polymath clinician writing about how he wrote a program to simulate what he saw as the essence of the system he was describing. My signed copy from John is one of my most valued books, a superinukbook! The program also found its way into preclinical courses of physiology, notably through the graphics versions that I created for the International Business Machines Corporation (IBM) PC microcomputer, when it made its appearance in the 1980s. I took it into undergraduate and postgraduate clinical teaching of anaesthetics and intensive care medicine,

in combination with structured learning materials that I developed with colleagues at Bart's and McMaster.

The question of how the Mac models might assist clinical insight and guide treatment had been the basis of my early 1970s PhD research at UCL. The example developed there had used two of the other models—of circulation and fluid and electrolyte balance—to investigate recovery from acute myocardial infarction. In the following years at Bart's, I worked with an anaesthetics and intensive care doctor, Charles Hinds. Charles went on to become Professor and Head of Department at Bart's, author of a very successful textbook and President of the Intensive Care Society. From the experience of the earlier PhD project, it seemed that a situation where extensive measurement of the body system was required for traditional management, combined with a clinical scenario where there was a small set of possible choices and adjustments of treatment, would be a good candidate area in which to explore the applicability of a model-based approach to patient management.

We set out to investigate how the MacPuf respiratory model might be used to interpret and guide ventilation management of patients in the ICU. This involved extensive collaboration with the anaesthetics research laboratory team at Bart's, in connection with new modalities of measurement and monitoring of patients. Mass spectrometry was being tried as an *in vivo* respiratory measurement device, along with routine measurements of blood gases, gas exchange and body metabolism. This involved a considerable investment of time and effort in liaison with industry partners involved in the project. Little of this came to fruition but it did provide useful context for the modelling work, when considering how improved measurement might enable the model to be matched more accurately with the clinical interventions being simulated. Rather, in the way that the expanded network of atmospheric sensors has underpinned advances in modelling of the weather.

It was a long haul over some five years, to bring these developments into alignment and devise a numerical optimization of just four key model parameters to match them to the measured respiratory variables. The high quality and dependable software productivity tools of today for mathematical analysis and software development—MATLAB, Mathematica, Eclipse and many more—were a distant vision, and those rudimentary ones that were already available, stuttered, changed and most of them quickly became obsolete. These were the temperamental equivalent of early motor cars, in need of constant, competent and time-consuming adjustment and maintenance.

The project succeeded in its primary goals and the model predictions were used in two ways. First, to explore options for managing patients

with severe respiratory failure, such as in viral pneumonia, and correlate them with the clinical pathway that unfolded.⁴¹ Second, within interactive computer-assisted learning courses for anaesthetics and ICU trainees, published in collaboration with Charles's opposite number at UCL, Rod Armstrong.⁴²

There were several reasons why this line of research came to an end. First, it proved of little interest as a research topic for the research council funding schemes of the times. Second, it would have required a larger team, working across the clinical and technical domains, which was not available to me and beyond my personal capacity to create at the time. And perhaps most significantly, the pilot project was likely to prove difficult, if not impossible, to improve and generalize more widely, beyond our setting at Bart's and the protective sponsorship provided there. I had been fortunate beyond words in the trust John Dickinson had shown in me, but as he approached retirement and I achieved a personal chair, it was important to look for new opportunity to extend the range of my academic activities.

The MacPuf model of human respiration was the most generic and widely used of its kind of its time. John was not a respiratory physician. He was interested in the challenge of representing clinical physiology and its system behaviours with computer models. Principally, he did this to explore their use as educational resources. He was not very interested or engaged in their validation and application in clinical practice, although he encouraged me to pursue that line. He saw his own experience, published research and expert colleague practitioners as providing the best available, realistic and reliable guides to their improvement. That was his nature and reflected the roles he had to balance in his professional work.

-
- 41 C. J. Hinds, D. Ingram, L. Adams, P. V. Cole, C. J. Dickinson, J. Kay, J. R. Krapez and J. Williams, 'An Evaluation of the Clinical Potential of a Comprehensive Model of Human Respiration in Artificially Ventilated Patients', *Clinical Science*, 58.1 (1980), 83–91, <https://doi.org/10.1042/cs0580083>; C. J. Hinds and C. J. Dickinson, 'The Potential of Computer Modelling Techniques in Intensive Care Medicine', in *Computing in Anesthesia and Intensive Care*, ed. by O. Prakash, Developments in Critical Care Medicine and Anesthesiology (Dordrecht: Springer Netherlands, 1983), pp. 153–69, https://doi.org/10.1007/978-94-009-6747-2_13; C. J. Hinds, M. J. Roberts, D. Ingram and C. J. Dickinson, 'Computer Simulation to Predict Patient Responses to Alterations in the Ventilation Regime', *Intensive Care Medicine*, 10.1 (1984), 13–22, <https://doi.org/10.1007/BF00258063>
- 42 C. J. Hinds, D. Ingram and C. J. Dickinson, 'Self-Instruction and Assessment in Techniques of Intensive Care Using a Computer Model of the Respiratory System', *Intensive Care Medicine*, 8.3 (1982), 115–23, <https://doi.org/10.1007/BF01693430>; J. B. Skinner, G. Knowles, R. F. Armstrong and D. Ingram, 'The Use of Computerized Learning in Intensive Care: An Evaluation of a New Teaching Program', *Medical Education*, 17.1 (1983), 49–53.

My work with Hinds was specific to the topic of respiratory system management in the ICU. It would have been possible to iterate further, in new clinical series of patients and with new kinds of data and optimization methods, but it appeared that the benefits this might bring to the everyday management of patients would prove very limited. Some modelling research groups that combined clinical and technical membership iterated their contributions in this way, within a single narrow domain of application. They succeeded for a while in publishing and republishing their work, maintaining a high profile, thereby. Extremely few of such outputs established and sustained useful clinical application. I would guesstimate that, twenty years on from publication, at least ninety-nine percent of them had already disappeared beyond the event horizon of research endeavour. Such is the anarchy of transition into the Information Age and its information explosion!

Modelling and Controlling Cardiovascular Dynamics

In the early 1970s there was much interest in what were termed model-based control systems. The design of feedback control systems was already a well-established engineering discipline and practice, enabling continuous adjustment of the system's controls to achieve a desired level of its performance. As represented, for example, by the control of the operating pressure of a steam engine by a steam governor, a purely mechanical device, or of the temperature of a water bath by a thermostat, typically in the form of an electromechanical device. How could a computer model of the controlled system be used to achieve this control, by predicting the effect of possible adjustments and using a numerical method to compute an optimum adjustment of the settings?

The PID (Proportional-Integral-Derivative) Controller was the basis of one engineering design that applied a corrective change in system settings, according to the magnitude, integral over time and rate of change, of the difference between actual and desired operating level, or performance, of the system. Another approach experimented with was to model the system to be controlled in purely mathematical terms—as a generically structured black box connecting inputs and outputs, with no attempt made to represent what was known about the actual structure and function of the system concerned. With this method, experimental perturbations of the inputs and measurements of consequent changes in outputs, were used to infer and update detail of the structure of the model. This model was then used to predict and control behaviour of the real system being modelled. It is a method akin to machine learning algorithms of today.

In order to characterize such models, various methods of signal analysis, pioneered in electrical engineering, were employed. These techniques involved decomposing measured signals into a linear combination of different frequency components (Fourier series), and experimental perturbations of the system settings, based on a pseudorandom binary sequence of inputs. This was the approach adopted in a leading clinical cardiology research centre of the 1960s and 1970s in Birmingham, Alabama. I flew there after my visit to Guyton and his team in Mississippi, in November 1971. It was soon after the era of loud gubernatorial politics of George Wallace, and the protests and peace movement of Martin Luther King. It was a place I approached with some trepidation. After the visit, I stood anxiously for an hour, alone outside on a humid, gloomy, stormy evening, waiting for an expected, but late arriving taxi, to take me back to my hotel. I can still see that alarming scene, in my mind, now.

At the hospital, I had arranged to meet a clinical cardiologist whose work I had read about and who was investigating cardiac muscle biomechanics in the context of cardiovascular disease. The more lasting and impactful event, by chance, was a visit he arranged for me, as we spoke, to the hospital's postsurgical intensive care unit, run by an already luminary, but still young, cardiac surgeon and medical computing pioneer of the era, James Kirklin. He had worked earlier at the Mayo Clinic and in collaboration with IBM at its Yorktown Heights research centre.

I was not able to meet Kirklin himself. I gathered that he worked through the day in theatre, came to the academic department in the evenings and worked there through to the small hours, before going home to sleep and arriving back to repeat the cycle, early next day! The informatics focus of his department was on the opportunity provided by real-time sensors to monitor and manage post-operative recovery of patients in the ICU. He had for some time been working to improve clinical outcomes for these patients, by introducing frequent measurement and a related set of 'house rules', as he termed them, for determining clinical management. The unit conducted extensive studies of the outcomes achieved through this close attention to detail of management, based on regular measurement of key variables. The computer system used to capture and process the data was developed and run by an engineer, Louis Sheppard, who welcomed me to the unit. The air-conditioned computer suite, located immediately above the ICU, was almost as large as the ICU itself, I recall.

Lou Sheppard later came on leave from his employment, and then on regular subsequent visits, to Bruce Sayers's (1928–2008) (subsequently Richard Kitney's) electrical engineering department at Imperial College, to complete a PhD there. He developed a model-based approach to patient management, based on data collected in the practical clinical setting

in Alabama, with Kirklin. We continued to meet from time to time, over several years. In his PhD project, he developed a linear frequency domain model of the response to infusion of sodium nitroprusside, for control of blood pressure. This used pseudo-random binary pulsed administration of the drug dose prescribed, under Kirklin's clinical supervision, to identify the defining parameters of the model when configured to represent an individual patient, and from this to predict and optimize the time-course of actual administration of the drug, to achieve and maintain stable blood pressure at a desired target level.

Lou's methods showed impressive results in controlling blood pressure and was extraordinarily successful in adjusting continuous infusion to cope with all manner of changes in clinical situations, acutely and over time, smoothly and effectively. A subsequent PhD student at Birmingham, John Slate, built on this to design and commercialize a model-based controller for infusion devices. I have a personal copy of his excellent doctoral thesis. I recall the remark of the then Professor of Medicine at The London Hospital Medical College, Robert Cohen (1933–2014),⁴³ when I shared these results at a meeting. He said that an interesting question arising was whether these patients would have done okay in their recovery, anyway, cared for with the prevailing human clinical skills of the time, even if perhaps a bit more chaotically.

In other words, another kind of control was required: a controlled clinical trial involving a suitably large number of cases, to demonstrate and convince that the new approach was clinically viable as part of everyday treatment and good practice, beyond the special environment in which it had been developed and brought to fruition. Such investigation and regulatory process is central to the approval of new pharmaceuticals but has proven harder to organize in gaining support for new devices—especially one such as this, which substitutes for human decision and uses a closed loop controller, impacting directly on treatment of the patient. This new possibility was inevitably unnerving for all concerned, including the regulators. Problems that such machines address, and the scope of the methods adopted for solving them, must be closely pinned down, to meet legal requirements for governance and accountability of clinical care.

43 Robert (Bob) Cohen was a close colleague of John Dickinson, then his opposite number at Bart's. They transcended politics through the turbulent years of rivalry leading to merger between the two hospitals, and incorporation of their medical schools within Queen Mary College (now QMUL). He was active in a leading initiative at the London Hospital, computerizing its patient administration, as I describe in Chapter Seven.

Where did the legal responsibilities and accountabilities traditionally carried by a clinician lie, when the actions taken were being determined, closed loop, by an algorithm. These questions posed both ethical and philosophical dilemmas, and they have been further highlighted as informatics has moved into the era of artificial intelligence. It is not dissimilar to the ethical issues surfaced in some debates about how model-based reasoning about epidemics has been used to decide policy for managing the Covid-19 epidemic. That said, at least the latter do not embody closed loop control. The decisions are human decisions.

Parenthesis—Purpose

Philosophers have debated noumenon and phenomenon over many centuries, in seeking clarity of reasoning about the world—the reality of the world underpinning what is experienced and observed. The golden rule when embarking on creation of a model of observed appearances, as Whitehead characterized the ways in which we experience the world, is to start with purpose. What are the appearances of the world that we seek to represent in a model, and why are we doing it? Purpose is a human quality; it embraces human values and goals communicated in the language of stories. There are machine goals, but not machine values—at least, not yet!

The 1972 mathematical model of the global economy, *Limits to Growth*, championed by the group called The Club of Rome, led to vocal controversy about assumptions, methods and predictions of models. The model became the focus of attention and argument, incapacitating as much as enabling debate and leadership on the issues. The exercise might perhaps have been better described as dealing with limits to our capacity to think about and act on the limits to growth! In such situations, the tails of special interest wag the dog of common purpose and goal. Conference of the Parties (COP) conferences seeking political traction on climate change have continued to navigate this familiar obstacle course.

Seeking illusory perfection, and losing sight of usefulness, models may become overly elaborate and intractable for the purposes they serve. They may equally be framed too simply, also limiting their usefulness. They may be overly restrictive or permissive of customization within different contexts of use, in both cases making their use more complicated. In *The End of Alchemy*, Mervyn King discussed the past twenty years of crisis in finance and banking and how ‘[In] the space of little more than a year, what had been seen as the age of wisdom was viewed as the age of foolishness.

Almost overnight, belief turned into incredulity'.⁴⁴ He attributed this not to failure of banking or policy but rather to a crisis of failed ideas. He cautioned wariness of over-dependence on rational models of economics—this was the alchemy he dramatized in his title. He emphasized, by contrast, the importance of narrative and storytelling.

When thinking about purpose, and goal of care information utility for the Information Society, we must also confront values. Machines that we create may come to embody and lead us to act according to values we do not hold—unwitting and unrecognized, but implicitly the case. A story illustrates how, unbeknown to us, such a machine may take us somewhere we would not wish to be, open to massive criticism about our values and governance.

Some decades ago, a medical school introduced an algorithm to assist the admissions team in selecting among students applying to study there. Courses in medicine are typically manyfold oversubscribed in relation to the numbers of places available. The algorithm was based on analysis of the school exam grades achieved by previously selected students, prior to entry, and their final exam results on completing the medical course. This was combined with the data provided by prospective students on their application forms. The goal of the algorithm was to predict which applicants were most likely to succeed in the course and guide the admissions team in their selections, accordingly. The discriminant analysis came up with a system to score applicants, based on these data. This was duly put into use, aiming to lighten the workload of admissions tutors, pressed for time in reviewing thousands of applications and deciding which of the applicant to invite for interview. A while later, the school discovered to its horror that it had been assigning points according to ethnicity!

44 M. King, *The End of Alchemy: Money, Banking and the Future of the Global Economy* (New York: W. W. Norton and Company, 2016), cover note.