A Journal for the History and Philosophy of Science
Spontaneous Generations

Expertise, a Framework for our Most Characteristic Asset and Most Basic Inequality

Author(s):
Claire Hooker
University of Sydney

Cliff Hooker
University of Newcastle, Australia

Giles Hooker
Cornell University

Source: *Spontaneous Generations: A Journal for the History and Philosophy of Science*, 10 no. 1 (2022) 27-35.

Published by: The University of Toronto

E D I T O R I A L O F F I C E S

Institute for the History and Philosophy of Science and Technology Room 316
Victoria College, 91 Charles Street West
Toronto, Ontario, Canada M5S 1K7
hapsat.society@utoronto.ca

Published online at *jps.library.utoronto.ca/index.php/SpontaneousGenerations*

Spontaneous Generations **10:1** (2022). University of Toronto.

Expertise, a Framework for our Most Characteristic Asset and Most Basic Inequality

Claire Hooker, Cliff Hooker, Giles Hooker

Abstract This essay provides a framework of concepts and principles suitable for systematic discussion of issues surrounding expertise. Expertise creates inequality. Its multiple benefits and the creativity of technology lead to a society replete with expertises. The basic binds of expertise derive from the desire of non-experts to be able to both enjoy what expertise offers and insure that it is exercised in the social interest. This involves trusting the exercise of expertise, involuntarily or voluntarily. A healthy society provides various means to move trust from involuntary to voluntary. The social means for achieving this are laid out.

The purpose of this short essay is to briefly lay out a conceptual framework within which to construct, clarify, evaluate and apply expertises. It is not to promote some particular notion of expertise over others, or to review the vast literatures, such as that on trust in science, that make up the domain. A few notes on one work towards this essay's close may indicate what a major, and expert, process this would be.

I The Basic Framework

Expertises emerge from the non-sharable attributes of information and capacity with which life is filled. Consider Louis Armstrong's jazz trumpet and a state attorney general's case preparation. The trumpeter has a strongly non-sharable capacity to make jazz, the attorney has a non-sharable store and application of information about laws, arguments, precedents and more with which to make cases. Both expertises require such intensity and time to learn that they are in practice non-sharable, though some unusual individuals can acquire them with sufficient effort, and the attorney's information store is perhaps somewhat the more accessible of the two.[1] The roots of these non-sharabilities could be traced back to the finitude and adaptation at the root of any creature's development and thus includes the wide variety of non-human expertises; but here it suffices to recognise their ubiquity and variability.

A human expert is the possessor of a relatively unsharable attribute in respect of some domain of normed endeavour, to the degree and in the relevant domain offered. Some expertises are strongly social in character (Law, language), others more weakly social (jazz), others very weakly social (non-competitive marathon running). Because of its origin in non-sharable adaptations, expertises are inherently excluding, and this role lies at the very roots of human expert constitution.

Only an expert can properly judge the work of another expert. Even when

[1] See in general Michelene Chi, Robert Glaser, and Marshall Farr (eds), *The Nature of Expertise*, rev. ed (London: Tayler and Francis, 2014) and K. Ericsson and J. Smith, *Towards a General Theory of Expertise: Prospects and Limits* (Cambridge: Cambridge University Press, 1991).

performance is catastrophically bad or prodigiously successful, understanding the performance still requires expertise(s). Moreover, many expertises are highly valued because of the supply of goods they enable, from jet travel to X rays and the internet of things. Further, establishing an expertise also establishes a suite of further expertises that have multiple valuable applications in other domains, from car racing, for example, to wind tunnel design and precision robotic manufacture. The increasing number of, and interplay among, positively valued expertises, combines with the effort typically required to acquire and apply them and with the competition, to exploit what expertises offer. Taken together these forces lead to self-reinforcing networks of expertises, for example in building construction. In these circumstances individual socio-economic niche roles are increasingly characterised in terms of the expertise(s) used, in relation to the rest of the expertise network. Such niche roles are specialisations. Expertise and specialisation tend to increase together, and to generate increasing complexity. There is a substantial literature concerned with the psycho-neurology of expertise which is set aside here as beyond scope.[2]

But the outcome is also one where each individual is expert in a decreasing proportion of the entire expertise network while dependent on an increasing proportion of that network. Everyone relies on expert management, but only a very few specialists will even know of its existence, let alone understand the expertise involved.[3] Further, knowing about dependencies does not alter them (e.g. viral epidemics**).** This may provide a constructive way to illuminate the surge of distrust that is now occupying democracies: the problems may reflect the inheritance of an inadequate sociology of expertise rather than being regarded as narrowly political.

II The Basic Expertise Problems.

Everyone quite reasonably wants to participate fully in all the relevant expertises that affect their lives, directly or indirectly, but no individual can develop expertises for more than a small, and reducing, fraction of them. This is the fundamental individual problem of expertise. A similar problem operates socially: the wider and deeper a society's use of expertises grows, the weaker grows its capacity to adequately surveil and regulate them.

The immediacy of the need for a resolution of the basic problems is a function of the impacts that would be brought about by failures to resolve them. For example, a doctor can penetrate deep into the working of a patient's body, and is licenced to do so but, with some exceptions (e.g. in handling a plague virus), clinical error mostly impacts on relatively few people, while an error in water purifying plants may impact millions, though often not very penetratingly, and a nuclear leakage error in a power generator may subsequently impact many, each penetratingly.[4] Add to nuclear power the

---

[2] See for example, M. Bilalić, *The Neuroscience of Expertise* (Cambridge: Cambridge University Press, 2017) and H. Collins and R. Evans, *Rethinking Expertise* (Chicago: University of Chicago Press, 2009).

[3] Cf. M. Tschannen-Moran, "Collaboration and the Need for Trust," *Journal of Educational Administration* 39 no. 4 (2001): 308–331.

[4] See, among many: M. Gibbons, "Science's New Social Contract with Society," *Nature* 402 (1999): C81-C84; D. Layton, E. Jenkins, S. Macgill, A. Davey, *Inarticulate Science? Perspectives on the Public Understanding of Science and Some Implications for Science Education* (Nafferton, Driffield: Studies in Education Ltd, 1993); H. Nowotny, P. Scott, M. Gibbons, *Re-thinking Science : Knowledge and the Public in an Age of Uncertainty* (Cambridge, England: Polity, 2001); J. Solomon,

invisibility of radiation, the irreversibility of its impacts, and the relative unintelligibility of its workings, all also features of modern expert life, and you have the basis for a heightened demand for control sufficient to resolve its problems.

III Problem Amelioration and Trust

Any effective amelioration of these problems must locate some form of assessment that replaces involuntary trust with voluntary trust. Trust, in one form or another, lies at the heart of any workable sociology of expertises.

(IIIA) One prominent amelioration of expertise problems is through the provision of broadly checkable features of expert performance trustworthiness. This can take two assessable forms. (IIIA1) Design-in at the outset broadly checkable, trust-increasing features. That is, for example, the approach taken by those who aim to ameliorate the inaccessibility of learning and decision making algorithms, used increasingly widely, by requiring that the algorithms used are inherently understandable in normal agent-action terms and their rationales assessable, for example to examine for inherent bias in algorithms (see below). (IIIA2) Subject all claimed expert performances to a designed process of critical examination of these claims, an examination by other professionals as equally or more expert in the domain. This is, for example, the approach taken by those who form professional institutions with the purpose to insure that the latest standards of excellence are adopted and adhered to (see below).

(IIIA1) *Design-in of trust-promoting features*. One way to promote public acceptance of expert roles is to require that all acceptable expert applications provide some minimum degree of public assessibility designed into them from the outset. This would allow some estimation of the degree to which the system 'works like us' in approaching problems and their solutions, reducing the fear of incompetence and hidden agendas, increasing voluntary trust in the overall processes adopted. Assessibility can be complemented or substituted by *reliability,* and both can be extended using *formative participation* to maximise non-expert access.[5] In what follows, an example of this design-in approach will be briefly reviewed, using our own features above to open up a wider domain of application.

IIIA1(i) By *assessibility* we mean the degree to which the broad public understands the subject matter, including assumptions made, risk mitigation processes adopted, and so on (i.e. its accessibility), and can assess options generated, their acceptable risks, trade-offs among these, and so on. However assessibility is not always possible. Some situations are simply too complex for non-experts (e.g. nuclear reactors) and sometimes the barriers run deeper (e.g. problem wickedness: poverty, ecological management).

"Changes to Science Education: Where Next?" *Canadian Journal of Science, Mathematics and Technology Education 2* no.1 (2002): 25-30; J. Ziman, *Real Science* (Cambridge: Cambridge University Press, 2000); and more broadly, C. Hayden, "The Automated Lab," *Nature* 516 (2014): 132; R. King et al., "The Automation of Science," *Science* 324 no. 3 (2009): 85-89; and https://allea.org/truth-trust-and-expertise/.

[5] See for example B. Mittelstadt, P. Allo, M. Taddeo S. Wachter, L. Floridi, "The Ethics of Algorithms: Mapping the Debate," *Big Data & Society* (2016): 1-21; and S. Wachter, B. Mittelstadt, L. Floridi, "Why a Right to Explanation of Automated Decision-Making does not Exist in the General Data Protection Regulation," *International Data Privacy Law*, 7 no. 2 (2017): 76-99. Using instead FAccT: Fairness, Accountability, and Transparency, provides a more formal – perhaps a too formal – formulation; see facctconference.org/., https://dl.acm.org/doi/abs/10.1145/3351095.3375624)

The growing use of AI learning algorithms to replace formal analysis in 'smart' vehicles, medical diagnoses, courtroom sentencing, and the like poses a particular conundrum worth noting: the AI learning processes may prove faster and more accurate than analytic reasoning, yet to do so they use processes that have no intuitive logical equivalent, or anyway no stable translation from algorithm states to explanatory agent states.[6]

IIIA1(ii) *Reliability* makes it possible to reduce the practical effects of inassessibility by focusing on ultimate empirical performance. 'How accurately does this AI algorithm track vehicular cruise control settings?' is a question that can be posed and answered and play its role in practical drivability, even if the underlying core algorithm is inassessible. Note, however, that the strategy trades-away gaining internal assessibility of the core processes for external practical intelligibility, blind to core inassessibility. That leaves the experts to continue to grow in power and inassessibility. Moreover, unreliable expertises, which typically originate in more difficult environments - political pundits, nutrition, and the like - form a much larger class of expertises, all of which are less assessible because uncertainty about process and state blurs knowing, on top of inherent process inassessibility.[7]

IIIA1(iii) *Formative participation* is participation in the key understanding and decision processes involved, interactively shaping the process with the experts involved so as to provide for the richest non-expert roles available. This applies to algorithm design in advance of final designs for public use, while in a medical setting it might be participation in forming a medical diagnosis-treatment process constructed for a specific patient. While that does not add new structural content to assessibility and reliability, it is worth noting here because it does open up new, potentially valuable, approaches to designing-in. For instance, the process should foster respect for the user, provide for the user's well informedness, and offers the opportunity to adapt technical expert knowledge to the persons to which it is to be applied. This brings the processes closer to those of the soft institutions discussed below.

Designing-in of the three-fold sort above carries us some distance toward ameliorating the problems our expertises invoke, but not without limits and difficulties arising. To extend ameliorisation further we next consider social institutions with 'hard' institutional arrangements for promoting voluntary trust.

---

[6] A. Abid, A. Ghorbani, J. Zou, "Interpretation of Neural Networks" Department of Electrical Engineering, Stanford University, California, 2017; C. Hooker and G. "Machine Learning and the Future of Realism," in *The Future of the Scientific Realism Debate: Contemporary Issues Concerning Scientific Realism*, ed. C Forbes, *Spontaneous Generations: A Journal for the History and Philosophy of Science* 9 no. 1 (2018): 174-182; cf F. Doshi-Velez B. and Kim, "Towards A Rigorous Science of Interpretable Machine Learning," arXiv.org > stat > arXiv:1702.08608.v2; and F. Doshi-Velez and M. Kortz, "Accountability of AI under the Law: The Role of Explanation," For the Berkman Klein Center Working Group on Explanation and the Law. ArXiv: 1711.01134.v1 [cs.AI] (2017).

[7] For example, P. Allen, "Complexity and Management," in *Philosophy of Complex Systems*, *Handbook of the Philosophy of Science, Vol.10,* ed. C. Hooker (North Holland: Elsevier, Amsterdam, 783-808); T. Brinsmead and C. Hooker, "Complex Systems Dynamics and Sustainability: Conception, Method and "Policy," in *Philosophy of Complex Systems*, *Handbook of the Philosophy of Science, Vol.10,* ed. C. Hooker (North Holland: Elsevier, Amsterdam, 809-838); H. Rittel and M. Webber, "Dilemmas in a General Theory of Planning," *Policy Sciences*, 4 no. 2 (1973): 155–69; R. Farrell and C. Hooker, "Design, Science and Wicked Problems," *Design Studies*, 34 no. 6 (2013): 681-705.

IIIA2 *Formal institutional safeguards: professions*. Growing alongside expertises are professions whose members are specialists. Many, perhaps most, professions are expertise-based (engineers, medical practitioners, etc.) and many, perhaps most, experts are supported through professions, exceptions including artists, life guards. Roughly, a profession is a socially established role, to some degree involving expert specialisation, while the expert role derives from the content and application of the expertise, though often also somewhat socially professionalised, most significantly in response to the basic expertise problems.

The essence of a profession, of what makes the institutionalising of the expert role socially acceptable and stabilises its practices, is a two-way compact between the group of experts involved and the society. The experts are trusted to perform their roles effectively and ethically, in the best interests of clients and the society, including respecting client privacy and properly resolving client/expert conflicts. Experts are to promote ethical public service, including fostering their own gain of, improvement in and respect for ethics, providing pro bono service as appropriate, and undertaking to self-regulate their conduct so as to be as transparent, explicable and accountable (see IIIA1 above) as practicable. Experts are trusted to provide ethical practice via self and professional regulation, using institutional membership that provides participation in formulating content of required knowledge and skills, assesses content mastery (or licences assessors), formulates code(s) of practice, technical and ethical, and assesses code compliance. Reverse-wise society is trusted to legitimise and sustain the distinctive social privileges of that professional role, including autonomous self-regulation, the capacity to charge fees commensurate with expertise, and status and authority.[8] The essential idea is that membership of an examining professional institution, with its social rewards as well as expertise examinations, will significantly reduce the odds of encountering member corruption or incompetence, and this gives rise to voluntary trust.

The design-in approach increases direct or first order expert system information to raise voluntary trust, 'they know how the system functions', while the profession option increases indirect or second order information about process to raise voluntary trust, 'they know the system functionality has been checked', with each having clear pre-conceptions of what is involved. However, and despite their differences, these approaches share some general limitations that curb too uncritical acceptance of voluntary trust. (a) Both have difficulty in dealing with complex and/or wicked problems (see 'assessibility' above). (b) Both approaches lack an inherent capacity to test for assumptions and biases within their procedures; for instance to check for uncritical adoption of gender biases in medical data used to test drugs, and so on[9] or for AI learning trained on such data.[10] (c) Neither approach possesses an inherent capacity to test for relevance when formulating the problem classes it will recognise, for example ignoring bias in data, and the cumulative effects of neglected weak relationships in modelling. (d) Neither is equipped to mitigate structural social inequality (e.g. race, SES). (e) Both have only limited capacity to respond to socio-cultural deficiencies that undermine acceptance of trust-based approaches (e.g. the rise of popularisms). All this suggests that we look in addition for further kinds of institutions and practices that can

---

[8] A. Abbott, *The System of Professions: An Essay on the Division of Expert Labor* (Chicago: University of Chicago Press, 1988); E. Hoyle and P. John, *Knowledge and Professional Practice* (London: Cassell, 1995).

[9] C. Perez, C, *Invisible Women: Data Bias in a World Designed for Men* (Penguin, 2020).

[10] M. Gutierrez, M, "Algorithmic Gender Bias and Audiovisual Data: A Research Agenda," *International Journal of Communication (Online)* (2021); 439.

serve as voluntary, but critical, trust increasers.

IIIB *Soft problem-solving processes*. A large cluster of relevant institutions and cultural practices are all those concerned with mediation, negotiation, arbitration, counselling, and the like. These can take widely different forms. Estate processors, for example, exercise a legal expertise shaped by the written contents of a document (the will) together with current estate case law and what is fair, just and practicable under that law. In contrast, family counselling process is an expertise shaped by the broadly emotional, social and cognitive functioning of the persons involved and their histories and experiences. The solution sought is one whose performance and resilience is grounded in those particular characters (albeit also within legal and similar constraints). Though deliberative, this is informal in a way that estate settling is not.[11]

The expertises utilised by these soft institutions have two distinctive features, (i) openness to re-formulation of the initial problem, methods, data and solution status as initially proposed, as a function of what exploration of the situation throws up, and (ii) a broad and learnable conception of what the specific roles are of the parties involved. These features result in more informal problem-solving processes characterised by depth of flexibility and creativity.[12] Voluntary trust is critical to these processes because they have the power to transform every element of the situation and all participants have to be trusted to caringly but intelligibly support such a deep process. Process manuals typically set establishment of trust at the centre of their procedures for these reasons, though as often it is treated only very generally. This treatment is non-trivial because informal or soft problem solving processes have their own limitations, particularly the systematic cognitive limitations displayed by humans in evaluating information.[13] For instance if some features of human psychology promote conditions for trust (e.g. shared orienting frameworks or 'worldviews'), these same features can also severely constrain it (e.g. by excluding different frameworks). In response, it will be necessary to develop cultural practices that weaken frame entrenchment and promote rich cross-frame relationships, including through more formal professional means.

IV Applying the Framework

IV.1 Evaluating alternative approaches, an example. To see how the above framework might apply we turn briefly to consider Collins and Evans, *Rethinking Expertise*.[14] *Rethinking Expertise* conceptualises expertise as tacit knowledge,

---

[11] See for example, P. Checkland, "Soft Systems Methodology: A Thirty Year Retrospective," *Systems Research and Behavioral Science* 17 (2000): S1-S89; C. Hooker, "Understanding HPS Paradigms Through Galison's Problems," *Axiomathes* (2021), sec. II.

[12] C. Argyris, *On Organisational Learning* (Malden, Mass: Blackwell Business, 1999); C. Argyris and D. Schön, *Organizational Learning: A Theory of Action Perspective* (Reading, Mass: Addison Wesley, 1978); cf P. Checkland, "Soft Systems Methodology: A Thirty Year Retrospective," *Systems Research and Behavioral Science*, 17 (2000): S1-S89; C. Hooker, "A Proposed Universal Model of Problem Solving for Design, Science and Cognate Fields," *New Ideas in Psychology* 47 (2017): 41–48.

[13] D. Kahan, H. Jenkins-Smith, D. Braman, "Cultural Cognition of Scientific Consensus," *Journal of Risk Research* 14 (2011): 147-74; D. Kahan, "Cultural Cognition as a Conception of the Cultural Theory of Risk," in *Handbook of Risk Theory* ed. S. Roeser (New York: Springer, 2012).

[14] Collins and Evans, *Rethinking Expertise*.

knowledge had but without explanation. It focuses on the empirical interdependencies among expertises that help to explain their emergence and functioning to yield a 'periodic table' of the expertises. For instance acquiring civil engineering expertise involves acquiring selected expertises in physics, mathematics, modelling, and more, in an appropriate time sequence. From within the framework developed above, an immediate question is what kind of non-sharable knowledge is involved? Wide roles for tacit knowledge are clear. But weaker kinds of non-sharable knowledge, such as an attorney may possess, are in explicit form, the non-sharabilities arising from constraints, cognitive and/or social. And there can be expected to occur many mixed forms across a modern society, for example AI (artificial intelligence) algorithms and human skills operating together in driving and flying modern vehicles. A related issue is what holds the periodic table organisation together? How fixed is its form? What roles do social structures play in this? What norms apply, or ought to apply, to table organisation? RE does not tackle these issues. Another issue underlined by its absence is the role of expertise impacts and the construction of which suit of trust-inducing institutions are most suitable for the table circumstances? Risk analysis guided modelling is one thing for safe-for-failure design (e.g. floodways) and another for safe-from-failure design (dam walls). And what, conversely, might be their impacts on table organisation? All of these explorations are stimulated and disciplined by having a suitable framework for expertise such as that offered above.

IV.2 Ecologies of trust. A society with multiple complex processes enriching life will require trust decisions, implicit and explicit, many times a day among citizens that are otherwise strangers to various degrees (this airline pilot, safe?) and will require such judgements across many domains, familiar and not to us. Section III indicates something of what goes into it. Any modern society will find it necessary to build the web of social structures and processes characteristic of well-functioning trust relations between experts and their society. Call this patterned web plus the voluntary trust-making institutions and practises that populate it, a society's ecology of trust. The quality of that ecology is the core determiner of a society's manageability and quality of life.

References

Abbott, A. *The System of Professions: An Essay on the Division of Expert Labor.* Chicago: University of Chicago Press, 1988.

Abid, A., Ghorbani, A., and Zou, J. "Interpretation of Neural Networks." Department of Electrical Engineering, Stanford University, California, 2017

Allen, P. "Complexity and Management." In *Philosophy of Complex Systems. Handbook of the Philosophy of Science, Vol. 10,* edited by C. Hooker, 783-808. North Holland/Elsevier, Amsterdam.

Argyris, C. *On Organisational Learning.* Malden, Mass: Blackwell Business, 1999.

Argyris, C., and Schön, D. *Organizational learning: A Theory of Action Perspective.* Reading, Mass: Addison Wesley, 1978.

Bilalić, M. *The Neuroscience of Expertise*. Cambridge: Cambridge University Press, 2017.

Brinsmead, T. and Hooker, C. "Complex Systems Dynamics and Sustainability: Conception, Method and Policy." In *Philosophy of Complex Systems. Handbook of the Philosophy of Science, Vol. 10,* edited by C. Hooker, 809-838. North Holland/Elsevier, Amsterdam, 2011.

Checkland, P. "Soft Systems Methodology: A Thirty Year Retrospective." *Systems Research and Behavioral Science* 17 (2000): S1-S89.

Chi, M., Glaser, R., & Farr, M. (eds.) *The Nature of Expertise*, rev. ed. London: Tayler and Francis, 2014.

Collins, H. and Evans, R. *Rethinking Expertise*. Chicago: University of Chicago Press, 2009.

Doshi-Velez, F. and Kim, B. 2017. "Towards A Rigorous Science of Interpretable Machine Learning." arXiv.org > stat > arXiv:1702.08608.v2. (2017).

Doshi-Velez, F. and Kortz, M. "Accountability of AI Under the Law: The Role of Explanation." For the Berkman Klein Center Working Group on Explanation and the Law. ArXiv: 1711.01134.v1 [cs.AI] (2017).

Ericsson, K. and Smith, J. *Toward a General Theory of Expertise. Prospects and Limits.* Cambridge: Cambridge University Press, 1991.

Farrell, R. and Hooker, C. "Design, Science and Wicked Problems." *Design Studies*, 34 no. 6 (2013): 681-705.

Gibbons, M. Science's New Social Contract with Society. *Nature 402* (1999): C81-C84.

Gutierrez, M. Algorithmic Gender Bias and Audiovisual Data: A Research Agenda." *International Journal of Communication (Online)* (2021): 439.

Hayden, C. "The Automated Lab." *Nature* 516 (2014): 132.

Hooker, C. "A Proposed Universal Model of Problem Solving for Design, Science and Cognate Fields." *New Ideas in Psychology* 47 (2017): 41–48.

Hooker, C. "Understanding HPS Paradigms through Galison's Problems." *Axiomathes*. In Press.

Hooker, C. and Hooker, G. “Machine Learning and the Future of Realism.” In *The Future of the Scientific Realism Debate: Contemporary Issues Concerning Scientific Realism*, edited by C. Forbes. *Spontaneous Generations: A Journal for the History and Philosophy of Science.* 9 no. 1, 174-182.

Hoyle, E., and John, P. *Professional Knowledge and Professional Practice*. London: Cassell, 1995.

Kahan, D., Jenkins-Smith, H. and Braman, D. “Cultural Cognition of Scientific Consensus.” *Journal of Risk Research* 14 (2011): 147-74.

Kahan, D. “Cultural Cognition as a Conception of the Cultural Theory of Risk.” In *Handbook of Risk Theory*, edited by S. Roeser. New York: Springer, 2012.

King, R et al. “The Automation of Science.” *Science* 324 no. 3 (2009): 85-89.

Layton, D., Jenkins, E., Macgill, S., and Davey, A. *Inarticulate Science? Perspectives on the Public Understanding of Science and Some Implications for Science Education*. Nafferton, Driffield: Studies in Education Ltd, 1993.

Mittelstadt, B., Allo, P., Taddeo M., Wachter S. and Floridi L. “The Ethics of Algorithms: Mapping the Debate.” *Big Data & Society* (2016): 1–21.

Nowotny, H., Scott, P., and Gibbons, M. *Re-thinking Science: Knowledge and the Public in an Age of Uncertainty*. Cambridge, England: Polity, 2001.

Perez, C. *Invisible Women: Data Bias in a World Designed for Men*. Penguin, 2020.

Rittel, H. and Webber, M. Dilemmas in a General Theory of Planning. *Policy Sciences* 4 no. 2 (1973): 155–69.

Solomon, J. “Changes to Science Education: Where Next?” *Canadian Journal of Science, Mathematics and Technology Education* 2 no. 1 (2002): 25-30.

Tschannen-Moran, M. “Collaboration and the Need for Trust”. *Journal of Educational Administration*. 39 no. 4 (2001): 308–331.

Wachter, S., Mittelstadt, B., Floridi, L. “Why a Right to Explanation of Automated Decision-Making does not Exist in the General Data Protection Regulation.” *International Data Privacy Law* 7 no. 2 (2017), 76-99.

Ziman, J. *Real Science*. Cambridge: Cambridge University Press, 2000.