



Public health and expert failure

Roger Koppl¹

Received: 18 April 2021 / Accepted: 5 September 2021 / Published online: 17 September 2021
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

In a modern democracy, a public health system includes mechanisms for the provision of expert scientific advice to elected officials. The decisions of elected officials generally will be degraded by expert failure, that is, the provision of bad advice. The theory of expert failure suggests that competition among experts generally is the best safeguard against expert failure. Monopoly power of experts increases the chance of expert failure. The risk of expert failure also is greater when scientific advice is provided by only one or a few disciplines. A national government can simulate a competitive market for expert advice by structuring the scientific advice it receives to ensure the production of multiple perspectives from multiple disciplines. I apply these general principles to the United Kingdom's Scientific Advisory Group for Emergencies (SAGE).

Keywords Pandemic · Experts · Covid · Scientific Advisory Group for Emergencies · Scientific advice

JEL Classification I18 · H4 · D8

Science is the belief in the ignorance of experts

~Richard Feynman

1 Introduction

Pandemics and other health crises demand state action. I think that the national governments of modern democracies generally should respond to health crises. In any event, they will. When a health crisis hits, it always is futile and usually inappropriate to tell modern democratic governments to stand aside and do nothing. In modern democracies, the state has become largely responsible for public health. State actors, including elected officials of a national government, are not in a good position to say, “We run the hospitals, we pay the doctors, we decide what drugs may be prescribed, but you’re on your own in this public health crisis.” It seems good and proper, then, to ask what norms, rules, organizational

✉ Roger Koppl
rkoppl@syr.edu

¹ Syracuse University, Syracuse, USA

structures, principles, processes, and procedures might best ensure that the actions the state takes in a public health crisis will do as much good and as little harm as reasonably possible. We need to set up good systems ahead of time because it is too late once crisis hits.

I raise a big question for which I have no grand comprehensive answer. I think I can answer, however, an important sub-question. How should national governments arrange for the provision of expert advice? When addressing health crises, national governments should and will seek expert advice. I have some suggestions about how to structure the provision of such advice. The gist of my idea is to simulate a freewheeling market for expert advice. My approach and my general principles apply to any modern democratic government, I think. I nevertheless will concentrate my attention on the United Kingdom (UK) and its Scientific Advisory Group for Emergencies (SAGE), in part because I am more familiar with it than comparable bodies elsewhere.

I intend my analysis to be useful and practical. To suggest impracticable “reforms” is not thinking; it is daydreaming. It would be wrong to impose an idle dream on busy readers. And it may be that SAGE is headed for reform. One report says, “The role of the Scientific Advisory Group for Emergencies (Sage) is likely to be reviewed when the Covid pandemic is over, government sources have said” (Rayner, 2021). The prospective inquiry “is expected to scrutinise Sage and consider whether such a monolithic body should hold so much power.” I intend my analysis to be useful in such an inquiry.

My suggestions are not unalterable marching orders or detailed blueprints. They are suggestive. Many stakeholders should contribute to the reform process, and they have different perspectives and know different things. Reforms that do not emerge from deliberation among those stakeholders are unlikely to achieve their stated ends. Deliberation may be guided, however, by general principles hammered out ahead of time. Should stakeholders bring inconsistent principles to the table, the conflicts must be resolved. And some general principles may emerge from the deliberative process. Some general principles, however, must be identified beforehand. It is in that spirit that I outline some general principles implied by information choice theory and sketch the broad contours of reform suggested by those principles.

Public health in the United Kingdom “is primarily a devolved matter”. Each of the four nations of the United Kingdom largely have independent authority in matters of public health. Thus, Covid lockdown policies may vary from nation to nation within the United Kingdom. And, “In the UK, the National Health Service (NHS) is the umbrella term for the four health systems of England, Scotland, Wales and Northern Ireland” (Institute for Government, 2021a). That devolution introduces some complexities to both the structure of SAGE and its role in the United Kingdom. For example, the “Welsh Government Technical Advisory Group” exists “to ensure that scientific and technical information and advice, including advice coming from the UK Scientific Advisory Group for Emergencies (SAGE) for COVID-19, is developed and interpreted in order to ensure” several outcomes, including the interpretation of “SAGE outputs and their implications for a Welsh context” (Welsh Government, n.d.). The complexities of devolution would have to be considered, of course, at some point in any inquiry scrutinizing SAGE. As far as I can tell, however, the complexities do not affect the general principles I lay down below. In this paper, therefore, I will ignore them from here on.

My proposed improvements, as I have said, would create within SAGE a simulated market for expert advice. Governmental bodies such as SAGE cannot somehow contain a free market for expert advice. An imperfect simulation is the best we can do. Improvement always is possible and always a worthy goal. But state bureaucracies necessarily are hierarchical no matter our attempts to simulate a polycentric market order within them. Thus,

they cannot always be as nimble and adaptive as we would wish. We should adjust our expectations to that inescapable reality.

2 Information choice theory

My proposal to create within SAGE a simulated market for expert advice is based on “information choice theory”, which is a theory of experts. Koppl (2018) is the most complete statement of the theory. As far as I can tell, the theory has not been fully anticipated in the past, if only because my definition of “expert” is innovative for a general theory of experts. But I have, of course, been influenced by a large body of prior work. Smith (2009), Turner (2001, 2010), and Levy and Peart (2017), are especially to be noted among works by living scholars.

2.1 Defining “expert”

It may seem natural to define experts by their expertise. But each of us occupies a different place in the division of labor and has, therefore specialized knowledge others do not. If an expert is anyone with expertise, we are all experts. But if everyone is an expert, no one is. It is better, I think, to consider what contractual role “experts” play. That role depends, of course, on how we define the term. At least one definition carves out a clear and distinct contractual role for “experts” while preserving a reasonable link to the word’s use in ordinary language. An “expert” is anyone paid for an opinion. It is, by that definition, a contractual role.

We usually are willing to pay for an opinion only when we think the person being paid has knowledge we lack. Thus, we usually hire as “experts” in my sense only persons we think possessed of “expertise” in the usual sense. But exceptions exist. A baseball umpire may have no more expertise in the usual sense than anyone else on the field of play. But we need to hire someone to say which balls were fair and which foul.

In Koppl (2018), I show that the foregoing definition of expert opens up a new class of models that is distinct, though not disjoint, from credence goods models, asymmetric information models, and standard principal-agent models. Canonical examples include Milgrom and Roberts (1986), Froeb and Kobayashi (1996), Feigenbaum and Levy (1996), and Whitman and Koppl (2010).

2.2 The theory of experts starts with the division of knowledge

Experts exist because we don’t all know the same things. Knowledge is “dispersed” as Hayek (1945, pp. 77, 79, 85, 86, 91) said. A division of labor creates a division of knowledge, which enables a more refined division of labor, which alters and refines the division of knowledge in a co-evolutionary process similar to Young’s (1928) description of endogenous growth in the extent of the market. Hayek (1937, p. 50) noted the “problem of the division of knowledge”, which is, he said, “at least as important” as “the problem of the division of labor”. But he lamented in 1937, “while the latter has been one of the main subjects of investigation ever since the beginning of our science, the former has been as completely neglected, although it seems to me to be the really central problem of economics as a social science.”

The theory of experts I develop in Koppl (2018), which I call “information choice theory”, addresses the problem of the division of knowledge in a way that Hayek seems to have overlooked. In an enduring contribution to economic theory, Hayek (1945, p. 86) explained that the price system is “a mechanism for communicating information”. But he seems to have given little attention to “knowledge markets”, which we may loosely define as particular markets in which knowledge is communicated. I think Hayek was right to see knowledge communication as a central issue in economic theory and policy. “The various ways in which the knowledge on which people base their plans is communicated to them is the crucial problem for any theory explaining the economic process, and the problem of what is the best way of utilizing knowledge initially dispersed among all the people is at least one of the main problems of economic policy—or of designing an efficient economic system” (ibid., pp. 78–79). And yet he seems to have considered closely only one way in which knowledge is communicated, namely, through prices.

Hayek seems to have neglected knowledge markets somewhat, but later thinkers have not. Hayek’s work was an important stimulus to information economics, mechanism design theory, and other treatments of knowledge and information in post-war economics. Presumably, we should consider some but not all of the large literature to be “about knowledge markets”. For this paper, it does not matter which bits count. Therefore, I will not pause to consider that literature further. I will note only that information choice theory addresses knowledge markets by considering one such market, the market for expert advice. It applies a broadly Hayekian understanding of the division of knowledge to the problem of how “the knowledge on which people base their plans is communicated to them” by experts and the “problems of economic policy” raised by the prospect of expert failure.

2.3 The nature of dispersed knowledge

Hayek did not just point out that knowledge is “dispersed”. He also discussed the nature of such dispersed knowledge. As far as I know, my treatment of dispersed knowledge adds nothing of substance to Hayek’s treatment. But I would not be troubled by the discovery that I had somehow deviated from the master. And, indeed, I have profited not only from Hayek (1937, 1945, 1952, 1967, 1978), but also Mandeville (1729), Wittgenstein (1958a, b), Smith (2009), and others. Garzarelli and Infantino (2019) suggest that I have only repackaged Smith’s (2009) “ecological rationality”. They may be right, although Smith probably gives less attention than I do to what happens beyond “skin and skull” (Clark & Chalmers, 1998). In any event, my exposition may be helpful to some readers.

I have elsewhere (Koppl, 2018) argued that knowledge is in the main “synecological, evolutionary, exosomatic, constitutive, and tacit”. Briefly, knowledge is “synecological” if the knowing unit is not an individual, but a collection of interacting individuals. It is “evolutionary” if it emerges from an undirected or largely undirected process of variation, selection, and retention. It is exosomatic if it is somehow embodied in an object or set of objects such as a book or an egg timer. It is constitutive if it constitutes a part of the phenomenon. The “knowledge” of Roman augurs studying bird flights was constitutive because it influenced events such as when or whether an enemy was to be attacked. And, finally, knowledge is tacit if it exists in our habits, skills, and practices while being simultaneously difficult or impossible to express in words. Nelson and Winter (1982, pp. 76–82) is an unusually helpful discussion of tacit knowledge. They note that incentives influence what remains tacit and what, instead, is made explicit. I have suggested the acronym SELECT

as a memory aid. The “L” in SELECT is meant to represent the L in “evolutionary”. Thus, knowledge is Synecological, EvoLutionary, Exosomatic, Constitutive, and Tacit.

The dynamic, synecological, and fractal nature of constitutive knowledge is illustrated by the “scotch egg” controversy in the United Kingdom. In late 2020, the British government promulgated its “COVID-19 Winter Plan”, which imposed a new and variegated set of restrictions on the movements and activities of persons in the United Kingdom, ostensibly to “suppress the virus” (Cabinet Office, 2020). In areas designated as “tier 2” by the plan, “pubs and bars” were ordered to close “unless they are serving substantial meals”. It became a public scandal that no one seemed able to clarify the meaning of “substantial meal”. The controversy came to focus on whether a “scotch egg” was a “substantial meal”. (That common picnic snack is a hard- or soft-boiled egg together with some sausage meat, which is breaded and either baked or deep fried.) Government officials gave the public confusing and seemingly inconsistent statements on whether a scotch egg is a substantial meal by the meaning of the government’s Winter Plan (Taylor, 2020). An editor at one newspaper spoke wryly of the “curious case of Schrodinger’s scotch egg” (Stevenson, 2020).

In many contexts, it is clear enough for practical purposes what is or is not a “substantial meal”. But such clarity is local, contextual, and variable. A substantial meal for a child may be inadequate for a lumberjack. A substantial breakfast in central Rome will not do on the farm. A substantial Christmas meal may be too much on January 1st. And so on. “Substantial meal” is a serviceable concept adapted to context. Using it as a general rule was sure to fail because any imposed rule becomes “a tool that unknown persons will use in unknowable ways for unknowable ends” (Devins et al., 2015). The term becomes a façade, a Potemkin village that hides purposes, choices, strategies, habits, and practices that cannot be known or imagined by the rule makers.

2.4 Experts are people

I have dubbed my theory of experts, “information choice theory”. The point is that experts must choose what information to convey. This label also is an allusion to public choice theory. Like public choice theorists, I assume behavioral symmetry among agents in my model. I assume behavioral symmetry between experts and non-experts. That assumption simply is the humble insight that experts are people.

The “insight” that experts are people is self-evident. Thus, the assumption of behavioral symmetry between experts and laity should be obvious and otiose. But my experience seems to suggest that it often is either denied (if only implicitly) or misconstrued. Some discussions of experts implicitly assume that experts are somehow higher and better than others. For example, Joe Biden has proudly proclaimed, “I trust scientists” (Grinalas & Sprunt, 2020). His remark expresses trust not in science, but scientists. And it seems to reflect the view that scientists are unswerving in their devotion to truth seeking. Similarly naïve views have been expressed by prominent scholars. Tom Nichols, the author of *The Death of Expertise*, has derided as “unAmerican this incredible suspicion of one another as if doctors don’t have an interest in you getting better somehow” (Amanpour & Company, 2020, pp. 6:01–6:13). It does not seem to have occurred to Professor Nichols that doctors are people and thus driven by motives no less flawed and complex than those of bus drivers, corporate executives, and movie stars. Sanjay Gupta, a physician and prominent public intellectual, described himself as “stunned” by Governor Cuomo’s loss of trust in health experts. He said, “If you start to take away some of the credence of those experts, I think that’s really, really harmful” (Budryk, 2021). Since the election of Donald Trump

as US president, it seems to have become increasingly common to view any questioning of experts or expert power as simultaneously irrational and immoral. One scholar has said that “fear distorts our thinking about the Coronavirus” and “the solution isn’t to try to think more carefully about the situation.... Rather, the solution is to trust data-informed expertise” (DeSteno, 2020).

Not everyone who proclaims that experts are people agrees with my view of human motives as “flawed and complex”. The scholar Rod Lamberts of the Australian National University seem to think that dastardly doubters dehumanize experts by viewing them as “commodities”. Professor Lamberts (2017) says, “Traditionally, an expert’s motivation for participating in public conversations as an expert will be rooted in a desire to inform, guide, advise or warn based on their specialist knowledge.” Lambert seems to think he has some nuance to add to this putatively traditional view. “But equally—and often simultaneously—they could be driven to participate because they want to engage, inspire or entertain. They themselves may also hope to learn from their participation in a public conversation.” He concludes, “At their heart, criticisms of experts often imply that they are servants, commodities or so vested in their field they can’t relate to reality” (Lamberts, 2017). Experts, apparently, are humble and self-effacing creatures who utterly disdain money, power, and prestige. Their human motives are only of the highest type and never base in any way. No wonder he could say, “To restore trust in experts, we need to remember they are, first and foremost, human beings.”

My model of human beings is less optimistic than Lambert’s. High motives surely are to be found, but low motives never are absent. Such a motivational mix applies to experts as a body and to each expert individually. Mandeville (1729, vol. I, p. 146) said, “no Body is so Savage that no Compassion can touch him, nor any Man so good-natur’d as never to be affected with any Malicious Pleasure.”

2.5 Key behavioral assumptions

Information choice theory builds, *inter alia*, on three key motivational assumptions. First, experts seek to maximize utility. The assumption that experts seek to maximize utility is parallel to the public-choice assumption that political actors seek to maximize utility. I have used the word “seek” to avoid any suggestion that experts must be modeled as Bayesian updaters or otherwise “rational” in some strong sense. Nor do I assume narrowly selfish or otherwise one-sided motives.

Second, expert cognition is limited and erring. That is a bounded rationality assumption. But I do not wish to invoke the standard model of bounded rationality. Felin et al. (2017) suggest why we should not limit our concept of “bounded rationality” to that found in the work of Herbert Simon or of Daniel Kahneman. They note that perception arises through an interplay of environment and organism and is, therefore, “organism-specific”. They quote Uexküll (1934, p. 117; quoted in Felin et al., 2017) saying, “every animal is surrounded with different things, the dog is surrounded by dog things and the dragonfly is surrounded by dragonfly things.” Similarly, in human social life, each of us occupies a different place and has, therefore, a different perspective, a different Umwelt. Though rooted in a phenomenological perspective rather than natural science, Alfred Schutz (1945) makes a similar point. Citing William James, he speaks of the “multiple realities” we inhabit. For man and beast alike, different “realities” or Umwelten have different Schutzian “relevancies”. Traditional bounded rationality models assume, instead, only one world, parts

of which may be obscured from view and only one right answer, which may be hard to compute.

Finally, expert errors are skewed by incentives. The assumption that incentives skew expert errors has two aspects. First, experts may cheat or otherwise self-consciously deviate from complete truthfulness. They *knowingly* may err to serve an external master or an internal bias. Second, experts *unknowingly* may err to serve an external master or internal bias. The first point is relatively straightforward, but not always applied consistently. A homely example illustrates the second point.

Think of what happens when an econometrician gets an unexpected or undesired outcome. The response will depend on the individual, of course. But he or she often will search diligently for ways in which the analysis has gone wrong. Should the desired or expected outcome be reached, however, the econometrician often will search less diligently for problems. The econometrician who responds asymmetrically in that way will produce biased results without realizing it or intending to. Econometricians might well be mortified by the thought that they have opened the path for bias. But such good intentions may be powerless to eliminate an asymmetry of which they are not conscious. My econometrics example is meant to be familiar and illustrative. Pichert and Anderson (1977), Anderson and Pichert (1978) and Anderson et al. (1983) provide experimental support for the claim that incentives skew even honest errors, as I discuss in Koppl (2018, pp. 176–177).

2.6 Theory of expert failure

I have provided a structural theory of expert failure. The theory of expert failure is not exhausted by the simple table I give in Koppl (2018, p. 190). For example, Murphy (2021) has expanded the theory by contributing a theory of “cascading expert failure”, whereby one expert failure may lead to another in a process similar to the regulatory dynamic described by Ikeda (1997) and others. Nevertheless, Table 1 identifies the most important institutional elements of the theory. Expert failure features two main institutional dimensions. The first is whether experts merely are advisory or whether, instead, the expert chooses for the non-expert. The second is whether experts are in competition or, instead, enjoy monopoly power. The largest chance of expert failure comes from the “rule of experts”, wherein monopoly experts choose for the non-expert. The smallest chance of expert failure comes from “self-rule or autonomy”, wherein competing experts offer advice without being able to choose for the non-expert. For my purposes in this paper, at least two other influences on expert failure matter, complexity and monopsony.

Complexity increases the risk of expert failure. The more complex the phenomena on which the expert delivers an opinion, the more likely is expert failure. In the language of Scott (1998), complex systems are not “legible”. The expert will rely on a model of the complex phenomena, whether that model be explicit, implicit, mathematical, verbal, or purely intuitive. Whatever form the model may take, it must simplify. Because the expert’s model simplifies, it will omit some causal paths, especially, perhaps, roundabout causal paths. Economists simplify when they rely on partial equilibrium models and ignore indirect general equilibrium effects. But as Hotelling (1932), Yeager (1960), Sonnenschein (1972, 1973), Mantel (1974), and Debreu (1974) illustrate, complex general equilibrium effects may overturn simple partial equilibrium effects.

Complexity may cause therapeutic advice, including policy advice, to be ineffectual or to produce undesirable unintended consequences. That consequence of complexity is most obvious in the case of “wicked problems”. The term “wicked problem” seems to have been

Table 1 A taxonomy for the theory of expert failure

	Monopoly expert	Competitive experts
Expert decides for the non-expert	<p>Rule of experts Examples include state-administered eugenics programs, central planning of economic activity, and central bank monetary policy Highest chance of expert failure</p>	<p>Quasi-rule of experts Examples include school vouchers, Tiebout competition, and representative democracy</p>
Non-expert decides based, perhaps, on expert advice	<p>Expert-dependent choice Examples include religion under a theocratic state and state-enforced religion</p>	<p>Self-rule or autonomy Examples include periodicals such as <i>Consumer Reports</i>, the market for preferences, and venture capital Lowest chance of expert failure</p>

coined by Rittel and Weber (1973) who list ten “distinguishing properties” of them. They contrast the “wicked” problems of “open societal systems” with the “tame” problems of “science and engineering”. DeFries and Nagendra (2017) have described “Ecosystem management as a wicked problem” in precisely the sense of Rittel and Weber (1973). They say, “Wicked problems are inherently resistant to clear definitions and easily identifiable, pre-defined solutions. In contrast, tame problems, such as building an engineered structure, are by definition solvable with technical solutions that apply equally in different places.” Rittel and Weber (1973), Head (2008), and DeFries and Nagendra (2017) all describe wicked problems as “intractable”.

The process of generating expert opinion likewise may be complex, raising the risk of “normal accidents of expertise” (Koppl & Cowan, 2010; Turner, 2010). Charles Perrow (1984) “used the term ‘normal accidents’ to characterize a type of catastrophic failure that resulted when complex, tightly coupled production systems encountered a certain kind of anomalous event” (Turner, 2010, p. 239). In such an event, Turner (*ibid.*) explains, “systems failures interacted with one another in a way that could not be anticipated, and could not be easily understood and corrected. These were events in which systems of the production of expert knowledge are increasingly becoming tightly coupled.”

Finally, monopsony increases the likelihood of expert failure. Monopsony is the existence of only one buyer in a market. It makes even nominally competing experts dependent on the monopsonist and correspondingly unwilling to provide opinions that might be contrary to the monopsonist’s interests or wishes. For example, the police and prosecution often are the only significant demanders of forensic science services in the United States.

As I have pointed out before (Koppl, 2018, p. 214), a kind of narrow monopsony also may sometimes encourage expert failure. An expert is hired to give an opinion to a client. The client is the only one demanding an opinion. That bilateral exclusivity is a narrow monopsony. The expert may have other customers, but none of them is paying the expert to give this particular client an opinion. In that situation, the expert may have an incentive to offer pleasing opinions to the client even if that implies saying something unreasonable or absurd. Toadies and yes men respond to such incentives. Michael Nifong, the district attorney in the Duke rape case, induced the private DNA lab he hired to withhold exculpatory evidence (Zuccino, 2006). The lab was private and thus nominally “competitive”. Presumably, it could have declined Nifong’s particular request without particular harm to its bottom line. But the lab chose to go along with Nifong’s desire to hide evidence. The District Attorney’s narrow monopsony created an incentive to do so. Only Nifong had effective control of the DNA evidence in that case. In the same narrow sense, the British government is a monopsony demander of SAGE’s advice.

If such a monopsony client seeks multiple opinions from redundant heterogeneous suppliers of expert opinion, then different suppliers may give different expert opinions. In that case, each expert has an incentive to anticipate the opinions of other experts and explain to the client why his or her opinion is best. Competition tends to make experts less like mysterious wizards and more like helpful teachers.

3 Expert failure in the pandemic

Significant examples of expert failure in SAGE seem to have materialized during the current pandemic. But examples often are contested or difficult to document unambiguously. It seems likely, for example, that the epidemiological models on which SAGE and others

relied made an inappropriate “homogeneity” assumption that, as Ioannidis et al. (2020) explain, “all people hav[e] equal chances of mixing with each other and infecting each other” (see also Murphy et al., 2021; Ridley & Davis, 2020). Models adopting the homogeneity assumption tend to overestimate the proportion of the population that must be infected or vaccinated to reach herd immunity (Britton et al., 2020; Gomes et al., 2020). Ioannidis et al. (2020) describe the homogeneity assumption as “untenable” and responsible for “markedly inflated” estimates of the herd immunity threshold. That all seems devastating. At least some researchers, however, explicitly defend simple models against Ioannidis and his co-authors because they “believe that the advantage of using less complicated pandemic models reduces the uncertainty in forecasting” (Dbouk & Drikakis, 2021, p. 021901-7). It is possible, though unlikely, that a standard and widely agreed upon list of SAGE’s errors and failures eventually will emerge. It is not possible that such a list could emerge while the pandemic is still on. At least one failure, however, seems to be important, unambiguous, and readily documented.

A report by the Treasury Committee of the House of Commons suggests that the economic costs of lockdown measures have been given inadequate weight and have not been appropriately clarified by the Treasury. The report says, “We strongly urge the Treasury to provide rigorous analysis of future policy choices which quantifies the harms and benefits of each of the plausible range of alternative policies.” The report further says, “The Treasury should be more transparent about the economic analysis which it undertakes to inform Government decisions in the fight against coronavirus and to publish any such analysis in a timely manner. The House should not be asked to take a view on proposals which have far-reaching consequences for the general population, such as those involving restrictions on social interaction, education, movement and work, without the support of appropriate and comprehensive economic analysis” (Treasury Committee, 2021, p. 4). The failure to weigh costs and benefits, including “economic” costs, is a striking failure of pandemic policy. As we shall see below, this failure seems to have been caused at least in part by the failure of SAGE to bring economists into the core SAGE group.

4 How SAGE works

SAGE seems to have committed at least one significant expert failure, and likely more. That is, the advice from SAGE seems to have been needlessly bad in at least one important way and likely others. Information choice theory and the theory of expert failure suggest that the structure of SAGE may have contributed to the problem. If so, then we may be able to say that the current structure is “bad” because it produces avoidable expert failure. Such expert failure is “avoidable” only if we reasonably can hope to find an alternative structure for SAGE likely to produce better advice. I believe that such beneficial structural changes can be found.

4.1 SAGE has no formal organization

The Scientific Advisory Group for Emergencies (SAGE) is an emergency body, not a regular, ongoing office or organization in continuous operation. The government calls each meeting separately and the list of participants varies from meeting to meeting (Government Office for Science, 2021). SAGE goes back only to 2009 (SAGE n.d.). The immediate predecessor to SAGE seems to have been the Scientific Pandemic Advisory Committee”

(Wells et al. 2011, p. 4859), which “was stood down on 5 May 2009 following the activation of SAGE” (Parliament 2021). Once “activated”, SAGE becomes the main scientific advisory body of the British government for the emergency in question.

No enabling law created SAGE or authorized its creation. Rather, the national government simply decided to “activate” it in 2009. Since then, it has been “activated” another nine times (SAGE n.d.) for a total of ten. (SAGE n.d. does not count the current activation, which brings the total to ten.) No binding formal organizational structures, procedures, or protocols exist for SAGE beyond its status as an input to the British government’s Civil Contingencies Committee, which is a part of COBR or COBRA.

The Civil Contingencies Committee “is convened to handle matters of national emergency or major disruption. Its purpose is to coordinate different departments and agencies in response to such emergencies. COBR is the acronym for Cabinet Office Briefing Rooms, a series of rooms located in the Cabinet Office in 70 Whitehall” (Institute for Government, 2021b). COBR is loose and irregular in its organization. “COBR meets during any crisis or emergency where it is warranted, but this can be ad hoc and the timing of meetings may be dependent on ministerial availability. Officials will convene a committee and use the emergency situation centre in the absence of ministers when a situation requires” (Institute for Government, 2021b).

In an emergency, one official document says, “COBR would be activated in order to facilitate rapid coordination of the central government response and effective decision-making.” And, “For a civil or non-terrorist domestic emergency, the Cabinet’s Civil Contingencies Committee (CCC) will meet bringing together Ministers and officials from the key departments and agencies involved in the response and wider impact management along with other organisations as appropriate” (Cabinet Office, 2010, pp. 11 & 23). Big national emergencies fall into the lap of COBR. And if the emergency is not related to terrorism, COBR acts primarily through CCC.

Although SAGE has no formal organizational structures, procedures, or protocols (beyond its status as subservient to COBR) a guidance document (Cabinet Office, 2012) was issued in 2012 and is still invoked by the government (SAGE n.d.). The guidance document says, “If activated, SAGE would report to, and be commissioned by, the ministerial and official groups within COBR” (p. 13). In other words, only COBR can “activate” SAGE. The guidance document contains little or nothing else that is truly binding. It says, “SAGE is designed to be both flexible and scalable. It is likely that its precise role will evolve as the emergency develops and vary by the nature of the incident” (p. 12).

4.2 Aims and objectives

The guidance document’s statement of “SAGE aims and objectives” is open ended (pp. 12–13). It says, in part, “SAGE aims to ensure that coordinated, timely scientific and/or technical advice is made available to decision makers to support UK cross-government decisions in COBR” (p. 12). Advice is “coordinated” when it is delivered to appropriate parties and based on inputs from appropriate parties. For example, if SAGE estimates that a certain number of hospital beds will be required on a certain date, the estimate should be conveyed to the NHS. And the number of beds available should be conveyed from the NHS to SAGE. Thus, again, the stated “aims and objectives” are open-ended.

The “aims and objectives” actually pursued by SAGE seem to have changed over time. Wells et al. (2011) describe how SAGE operated when it was first “activated” in 2009 for the H1N1 pandemic. “Throughout most of the pandemic, SAGE met weekly and received

papers and reports and reviewed modelling from three independent groups of mathematical modellers” (p. 4859). And, “An important part of SAGE’s role was to communicate the uncertainties, particularly from the mathematical modelling, to ministers” (p. 4860). Such emphasis on uncertainty seems to have dimmed and the idea of multiple perspectives seems to have been lost. The government’s SAGE explainer from 5 May 2020 says, “SAGE’s role is to provide unified scientific advice on all the key issues, based on the body of scientific evidence presented by its expert participants” (SAGE, 2020). While the explainer does make reference to “uncertain scientific evidence”, the call for “unified” advice seems far from the earlier use of multiple independent teams and its corresponding greater emphasis on communicating uncertainties “particularly from the mathematical modelling”. Similarly, the guidance document says, “Where there are differences in expert opinion, these should be highlighted and explained to ensure decision makers are given well-rounded, balanced advice” (p. 47). But it does not say anything to the effect that “differences in expert opinion” should be viewed as healthy and normal, let alone sought out.

4.3 The monopoly power of SAGE

SAGE is part of COBR, which is uniquely responsible for the national government’s response to an emergency. SAGE is the primary source of scientific advice to COBR in an emergency. One official document says, “The role of SAGE is to bring together scientific and technical experts to ensure co-ordinated and consistent scientific advice to underpin the central government response to an emergency” (Office, 2010/2013, p. 70). Thus, SAGE has monopoly power in the provision of scientific advice to the British government in an emergency. The monopoly power of SAGE has been noted in political discussions in the United Kingdom (Rayner, 2021). In a clear reference to SAGE, MP Mark Harper (2020) has said, “it’s time to end the monopoly on advice of government scientists.”

4.4 Narrow governance

Responsibility for selecting SAGE members is provided by the “SAGE secretariat”. When only one department is strongly involved in the emergency it is the “Lead Government Department (LGD)” and it will serve as the secretariat. Thus, the Department of Education might be the Lead Government Department if a wave violent protests was spreading across British universities, and the Department for Transport would likely be the Lead Government Department if a sequence of rail disasters occurred in the country. In other emergencies, the Cabinet Office and Government Office for Science would be the main departments acting as SAGE secretariat. “In all circumstances, Cabinet Office would be responsible for ensuring that SAGE had a UK cross-government focus whilst the Government Office for Science would be responsible for ensuring that SAGE drew upon an appropriate range of expertise and on the best advice available” (Guidance doc p. 16). In other words, the civil service is responsible for selecting the governmental members of SAGE and the Prime Minister’s office, through the Government Office for Science, is responsible for picking the extra-governmental experts.

The power to select members thus is concentrated in a small number of persons. In the case of COVID-19, SAGE’s power is concentrated in two persons, the government’s Chief Medical Officer, Chris Whitty. And the government’s Chief Scientific Advisor, Patrick Vallance. The concentration of appointment power creates the possibility of inappropriate homogeneity of SAGE membership. The guidance document articulates a set of

“principles” for membership, which include the call for “a wide-range of appropriate scientific and technical specialties” and the admonition that SAGE not “overly rely on specific experts” (p. 19). But those principles are vague and lack an enforcement mechanism. Nor does the guidance document seem to imagine the possibility that experts within a specialism might disagree. A risk of bias or partiality in the selection of SAGE members thus arises. Such bias or partiality may operate unconsciously and does not require anything like cheating, ill will, or self-serving.

Concentrated selection power also creates the risk of an *appearance* of impropriety even when no impropriety exists. For example, Patrick Vallance held a high position with the pharmaceutical company GlaxoSmithKline before being appointed Chief Scientific Advisor. Vallance had “a £600,000 shareholding in a drugs giant contracted to develop a Covid-19 vaccine for the Government, prompting claims of a potential conflict of interest” (Hymas 2020). Presumably, Vallance has not been obsequiously doing the business of GlaxoSmithKline while Chief Scientific Advisor. But the appearance of a conflict of interest may tend to diminish public trust in both SAGE and the overall Covid response of the government. It also seems possible to wonder, whether justly or not, if Vallance’s experience with a large pharmaceutical company has left him more disposed toward a policy of widespread vaccinations than he might otherwise have been.

Controversy about Chris Whitty also illustrates the problem that narrow governance creates the risk of a false sense of impropriety about SAGE. Whitty held an important position within the London School of Hygiene and Tropical Medicine when it received, in 2008, about \$46 million from the Gates Foundation (London School, 2008). About \$40 million of this sum went to a project headed by Whitty (London School, 2008). In the meantime, Gates had been something of a Cassandra, warning of the risks of pandemic, which he thought best to combat with vaccines (Gates, 2015). Moreover, “The Gates Foundation has been a key donor to the WHO over the past decade, accounting for as much as 13% of the group’s budget for the 2016–2017 period” (McPhillips, 2020). Thus, Gates is something of a big player in global health policy. Apparently, significant evidence suggests that his foundation may have profited from the current pandemic. “Recent SEC filings and the foundation’s website and most recent tax filings show more than \$250 million invested in dozens of companies working on Covid vaccines, therapeutics, diagnostics, and manufacturing. These investments put the foundation in a position to potentially financially gain from the pandemic” (Schwab, 2020). Those facts, particularly when transmitted in hushed tones, may seem to suggest that Bill Gates is some sort of sinister puppet-master manipulating Whitty and, though him, the whole of SAGE.

The notion of Bill Gates as venal puppet-master is implausible. His wealth is not proof of greed, but an incentive to generosity. The marginal value of a dollar is low for him, and the marginal psychic cost of doing wrong to get more money is correspondingly high. Besides, the gains flagged by Schwab (2020) would go to Gates’s foundation, not his bank account. We would further have to imagine that a donation made several years earlier gave Gates an enduring Svengali-like influence on Whitty, notwithstanding Whitty’s independent interest in his reputation as a researcher. Nevertheless, enough mistaken commentary in this direction has been published to induce the fact-checking website Full Fact to assure its readers that “Chris Whitty did not *personally* receive millions of pounds from the Bill & Melinda Gates foundation” (Full Fact, 2020, emphasis added). The point of my Whitty example is that narrow governance creates the risk of an *appearance* of impropriety even when the claims of impropriety are implausible or downright silly.

Having found Gates innocent of Svengalism, I had better note that he is not therefore beyond criticism. William Easterly (2013, pp. 6, 152) is the most notable of his critics

among economists. He laments Gates's "technocratic illusion". In 2013, he reports, Gates even "echoed" the "idea of a benevolent autocrat implementing expert advice to achieve great results in Ethiopia" (ibid., p. 156). Easterly excoriates the "disrespect for poor people shown by agencies such as the World Bank and the Gates Foundation, with their stereotypes of wise technocrats from the West and helpless victims from the Rest" (ibid., p. 350).

It does seem fair to suggest that Gates and the Gates Foundation have a technocratic vision. They apply that technocratic vision no less to pandemics than to economic development. And most empowered experts in the rich democracies seem to take a technocratic approach to the COVID-19 pandemic as well. It also is true that the Gates Foundation "has committed about \$1.75 billion dollars to support the global response to COVID-19" (Gates Foundation, 2021). Vaccines are important to the Gates Foundation's vision for COVID response. Whether "the world" in 2021 "gets better for everyone", Melinda Gates has said, "depends on the actions of the world's leaders and their commitment to deliver tests, treatments, and vaccines to the people who need them, no matter where they live or how much money they have" (Gates Foundation, 2020). It may be, therefore, that Bill and Melinda Gates and the Bill and Melinda Gates Foundation have contributed to a technocratic *zeitgeist* that emphasizes the potential of vaccines. If funding for research and public health policy measures tends to flow more easily to technocratic projects, a technocratic *zeitgeist* will be strengthened correspondingly. And it would not be surprising to find government experts agreeable to such a vision. If we should see such a *zeitgeist* in current events, we should not imagine that Bill or Melinda Gates somehow conjured it from thin air. They have the power only to add or subtract. And to see such a *zeitgeist* at work is quite different than imaging a conspiracy in which a sinister clique headed by a venal puppet-master is orchestrating world events.

4.5 Funding

Funding for SAGE is irregular. If an LGD is identified, the guidance document says, it "would provide the lead in ensuring that there is adequate funding for the provision of SAGE activities" (p. 29). Otherwise, "attempts should be made to reach a consensual financial solution between the key customers of the advice being provided" (p. 29). In other words, everyone involved should be reasonable and should work something out.

It is not expected that members, the SAGE experts, will be paid to be on SAGE. "Given SAGE relies largely upon the good-will of many experts; provisions should be made to cover appropriate personal expenses quickly and efficiently" (p. 29). And yet there is wiggle room. "However, given the need to ensure that SAGE remains focused on supporting UK cross-government decision-making other costs should be considered on a business case basis" (p. 29). The presumption against paying SAGE experts does not seem to be binding given the ad hoc character of the body, although I am aware of no case in which a direct payment was made.

4.6 Disciplinary narrowness and siloing

The guidance document says, "To ensure the full range of issues are considered advice needs to stem from a range of disciplines, including the scientific, technical, economic and legal." And yet the economic dimension seems to have been slighted in the COVID-19 SAGE. A figure from the guidance document, "The COBR mechanism" (p. 11), illustrates the problem. It is reproduced here as Fig. 1.

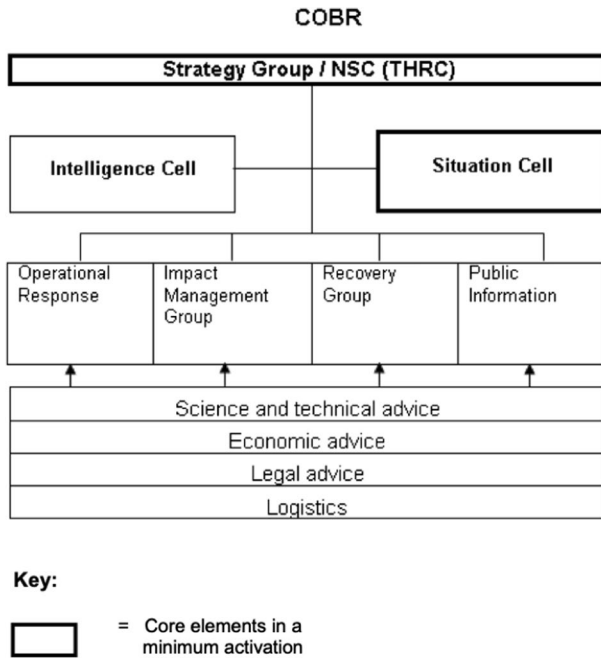


Fig. 1 The COBR mechanism

In the figure, economic and legal advice are separate from “science and technical advice”. By thus separating out economic and legal advice, the guidance document encourages the view that disciplines such as epidemiology that are thought of as “scientific” are independent of economics. “Economics” and “science” are connected in at least two ways, however. First, the policy prescriptions emergent from “science” models, including epidemiological models, should be considered from the economic point of view. Klement (2020, p. 70) says,

This involves, but is not limited to: (i) research on causal mechanisms within the healthcare system and hospital structure impacting on COVID-19 deaths; (ii) research on financial conflicts of interest of scientific advisors, research institutions and politicians and how these influence public health decisions; (iii) financial toxicity resulting from the lock down in many countries and how this impacts long-term anxiety and depression disorders and deaths such as suicides resulting thereof; (iv) modeling of the complex relationship between pandemics, economics, and financial markets.

Second, “science” models, including epidemiological models, make “economic” assumptions about human behavior. Dasaratha (2020) shows that “changes in infection risk” may have “counterintuitive effects” when people respond to the risk of infection. Dasaratha uses the term “risk compensation”, which seems to be the most common term of art for such behavioral changes. Trogen and Caplan (2021) say, “In situations that are perceived as risky, people naturally adjust their behavior, compensating to minimize that risk.” They point out that it is also known as the “Peltzman Effect” (see Peltzman, 1975). Economic theory helps to identify behavioral regularities such as the Peltzman effect, which

are generally absent from standard epidemiological models, including the famous “Report 9” (Ferguson et al., 2020) co-authored by Neil Ferguson (see Dasaratha, 2020, p. 1, n. 1).

Separating out “economic” analysis from “scientific” analysis creates both disciplinary narrowness and the closely related phenomenon of disciplinary siloing, “whereby one becomes so engrossed in one’s silo that one fails to consider, or may even be unaware of, other salient issues” (Murphy et al., 2021). It creates disciplinary narrowness by limiting inappropriately the numbers and types of scholarly disciplines brought to bear on the emergency. It creates disciplinary siloing by reducing the probability that an expert lodged in one silo will become aware of other salient issues.

The existence of a “nudge squad” within SAGE does not greatly qualify my claim that SAGE tends toward disciplinary narrowness and siloing. The guidance document notes that “it is likely to be necessary to create sub-groups” (p. 13). One of the subgroups of the COVID-19 SAGE is the “Scientific Pandemic Insights Group on Behaviours (SPI-B)”, which “provides advice aimed at anticipating and helping people adhere to interventions that are recommended by medical or epidemiological experts” (Government Office for Science, 2021). Prominent in that subgroup is David Halpern, head of the Behavioral Insights Team, also known as the Nudge Unit. The group, originally created within the British government, is now nominally a private organization whose mission is to “generate and apply behavioural insights to inform policy, improve public services and deliver results for citizens and society” (BIT n.d.). The SAGE nudge squad does not seem to be charged with bringing insights such as the Peltzman effect to the attention of the epidemiologists within SAGE. It seems rather to be charged with designing “nudges” to enhance compliance with COVID-related policy.

4.7 No formal processes of contestation

SAGE contains no formal process of challenge and contestation. Discussion occurs at SAGE meetings, and disagreements may be expected. But there is no required or customary process of “red teaming”, whereby a subgroup is tasked with challenging the emerging consensus of the larger body. Neustadt and Fineberg (1978, p. 75) warned of the danger in the context of the “swine flu affair” of 1976. “Panels tend toward ‘group think’ and overselling, tendencies nurtured by long-standing interchanges and intimacy, as in the influenza fraternity.”

4.8 PFM measures are inapplicable

Given the ad hoc character of SAGE, it seems impossible to discipline it through systems of public financial management (PFM). “Fiscal discipline, allocative efficiency, and operational efficiency have been adopted widely by the international PFM community as the explicit aims of the PFM system—the ‘holy trinity’ of standard PFM objectives” (International Working Group, 2020, p. 2). In the common view at least, a national government’s PFM system should be designed to ensure the fulfillment of these standard objectives of PFM professionals. PFM assessment tools are meant to provide monitoring of goal fulfillment. We do not need to consider competing models of PFM systems to see that the sort of monitoring and fiscal discipline imagined cannot meaningfully be applied to SAGE as it currently is constituted and governed.

It seems impossible to integrate SAGE better with the British PFM system by specifying goals and performance metrics. After the PFM reforms initiated in the late 1980s in

New Zealand, “Virtually every element of [PFM] reform has been designed to establish or strengthen contract-like relationships between the government and ministers as purchasers of goods and services and departments and other entities as suppliers” (Schick 1998, p. 124). That model is ill suited to SAGE because we cannot clearly specify goals and metrics. Sasse et al. (2020, p. 18) describe the incoherence of decision making within SAGE in Fall of 2020. They quote a SAGE member. “Ministers said: ‘What should we do?’ and scientists said: ‘Well, what do you want to achieve?’” Sasse, Haddon, and Nice comment, “Some back and forth is necessary to refine questions, but scientists said ministers’ objectives remained unclear throughout the crisis.” In an emergency, the government doesn’t know what its goals should be or how to measure progress toward them. The terms of any contract, agreement, or directive cannot be specified clearly enough to indicate whether its terms have been fulfilled. Thus, it seems impossible to reform SAGE through improved PFM.

5 Reform principles

When my description of SAGE is placed in the context of information choice theory, it suggests some basic reform principles.

As we have seen, SAGE has but little formal structure. Given the risks of expert failure created by the current loose structure (as explained below), a binding set of well-crafted protocols would tend to improve SAGE advice. It may be understandable that an emergency body largely would be devoid of binding protocols. Events are moving quickly, and decisions must be made. But an emergency can last a long time. The first activation of SAGE lasted for about 8 months. The current activation has lasted about 14 months at the time of this writing. The average across all ten activations is over four and a half months. Thus, room would seem to be available in a typical SAGE activation for adherence to protocols. Such protocols might not kick in immediately. They might apply only after the first meeting or only after, say, one week or one month. But they should kick in early and bind tightly.

5.1 Monopoly power

We have seen that SAGE has monopoly power in the provision of expert advice to the government. According to information choice theory, experts are more likely to fail when they have monopoly power. Reducing the monopoly power of SAGE’s experts would reduce the chance of expert failure in SAGE. It would tend to improve the quality of SAGE advice.

It might be possible somehow to break SAGE up into separate groups that must compete for the ear of the British government. It seems more straightforward, however, to form three competing expert teams within SAGE. Recall that when SAGE was first activated in 2009 it proceeded in much that way. Wells et al. (2011, p. 4859) described the use of “three independent groups of modellers.” Some such procedure likely would enhance the ability of SAGE to “communicate the uncertainties... to ministers” (Wells et al., 2011, p. 4860).

Koppl et al. (2008) report experimental evidence suggesting that three is the right number. In their experimental setup, redundant independent opinions tended to improve system performance, but only if at least three independent opinions were given. Expanding beyond three opinions did not improve system performance. Such redundancy will be useless if the opinions are not provided by truly independent experts or groups of experts.

Without independence, errors may be correlated across experts. In that case, the recipient of the expert opinions (COBR) mistakenly may believe that a strong scientific consensus has emerged when, perhaps, no such consensus exists. Or they may be led to think that the evidence points unambiguously in one direction when a more thorough accounting would reveal greater ambiguity. It is necessary, therefore, to construct appropriate mechanisms within SAGE to prevent collaboration across teams or any form of “information pollution” (Koppl, 2005, p. 258) that might compromise the independence of each team’s analysis from those of the other two.

5.2 Disciplinary narrowness and siloing

We have seen that SAGE tends toward disciplinary narrowness and siloing. According to information choice theory, such narrowness and siloing increases the chance of expert failure. Expanding the disciplinary base of SAGE would reduce the chance of expert failure in SAGE. It would tend to improve the quality of SAGE advice.

It might at first seem downright trivial to work up a protocol that enlarges the disciplinary base of SAGE. “Thou shalt have a lot of different disciplines represented.” But in my jocular example, “a lot” is ill-defined. The appropriate model of disciplinary breadth will depend on the emergency. Nor is it easy to define “discipline”. Is the discipline chemistry, organic chemistry, or computational chemistry? Or is it, perhaps computational organic chemistry? And so on.

No rule has been written that could identify for all possible emergencies the appropriate disciplinary breadth. Thus, the written protocols would have to be limited to vague statements about the importance of disciplinary breadth. The same goal also can be approached, however, by a very different path. Below I will discuss the value of a broad governance structure. A broad and representative governance body would help to ensure appropriate disciplinary breadth. Imagine that SAGE were activated, but with inappropriately narrow disciplinary representation. The disciplinary mix would be “inappropriate” because it made certain considerations such as causal linkages less likely to be raised and given appropriate weight. Slighting of certain issues would be bad for someone. A narrow governance body would be relatively unlikely to include persons representing those harmed by the omissions. A broad governance body would be more likely to include them. Thus, a broad governance body would leverage the interests of its multiple participants to help ensure appropriate disciplinary breadth on SAGE.

5.3 Unified advice

We have seen that SAGE has no formal processes of contestation and that its role has been construed as the provision of “unified scientific advice”. According to information choice theory, experts are more likely to fail when there is no contestation among expert opinions. Formal processes of contestation within SAGE, such as “red teaming”, would reduce the chance of expert failure in SAGE. (I agree with an anonymous referee of this paper who comments that red teaming, while helpful, “isn’t as good as actual dissent.”) Contestation would tend to improve the quality of SAGE advice.

It should be relatively straightforward to institute a protocol for “red teaming” SAGE analyses. A “red team” is a devil’s advocate. Thomas and Deemer (1957) is the earliest relevant use I have found in JSTOR. They discuss “operational gaming” as a tool of operations research. In “operational gaming”, participants play a formal non-cooperative game

meant to simulate “an actual conflict situation” (Thomas and Deemer (1957, p. 1). They say, “operational gaming derives from the old war gaming that military academies have long used to teach tactics” (ibid.). In their discussion a “red team” and a “blue team” compete. Wolf (1962) provides a broadly similar discussion of a “red team” and “blue team”, but in the context of “military assistance programs in less-developed countries”. The red team is equated with the USSR in Davis (1963, p. 596) and North Vietnam in Whiting (1972, p. 232).

Today, “red teaming” often means a process in which a team is assigned the task of challenging as strongly as possible the analytical or scientific results of the rest of the group, which may or may not be referred to as the “blue team”. In a scientific context, a “blue team” of scientists might present “the most robust evidence” in support of a given interpretation, theory, or model, while a “red team” would “seek to find flaws in the arguments”. In some versions, “the process would repeat until, in theory, consensus emerges” (unattributed, 2017).

It has been argued that the scientific process of peer review makes red teaming inappropriate in scientific contexts. In 2017, the American Association for the Advancement of Science (AAAS) objected to an EPA proposal to use red teaming to challenge the climate science consensus. “The peer review process itself is a constant means of scientists putting forth research results, getting challenged, and revising them based on evidence. Indeed, science is a multi-dimensional, competitive ‘red team/blue team’ process whereby scientists and scientific teams are constantly challenging one another’s findings for robustness” (AAAS, 2017). Whether peer review did or did not make the EPA’s 2017 proposal inappropriate, the lesson does not carry over to SAGE. Contentious issues of pure science probably matter in most or all emergencies, albeit more in some than in others. Howsoever that may be, the link from science to policy rarely or never will be obvious and uncontroversial. Red teaming can help to spot problems in that link.

5.4 Narrow governance

We have seen that SAGE has a narrow governance structure, which creates the possibility of inappropriate homogeneity of SAGE membership. According to information choice theory, experts are more likely to fail when homogeneity of opinion prevails among experts. A broader governance structure for SAGE would reduce the chance of expert failure in SAGE. It would tend to improve the quality of SAGE advice.

Broader governance means that persons with competing and inconsistent interests would share governance responsibilities. It is precisely the diversity of inconsistent interests within the governing body that ensures that multiple perspectives are represented within the organization. I noted earlier that narrow governance creates a risk of bias in the selection of members. Selected members also may be biased by their dependence on one or two parties for continued participation in SAGE. Those biases may be unconscious. Broader governance of SAGE would reduce the strengths of such biases. It also would create the possibility that one bias may be a check on another. In that sense, it would allow SAGE to “leverage” bias (see Koppl & Krane, 2016).

The Houston Crime Lab illustrates the benefits of broad governance. As many commenters have noted, Houston’s crime lab was badly run and produced spectacularly bad forensic science analyses (Garrett, 2017; Koppl, 2005). It experienced a “resurrection” (Garrett, 2017, p. 980) after 2012, however. The lab had been a part of the Houston Police Department. In 2012, it was removed from the police and “reincorporated... with an

Table 2 Basic reform principles for SAGE

Problem	Solution
Monopoly power	Three competing teams
Disciplinary narrowness and siloing	Broader set of represented disciplines
Unified advice	Red teams
Narrow governance	Broad governance

independent oversight board” (Garrett, 2017, p. 985). The lab has since “adopted important new quality controls” including a “blind quality control program, where fake test cases are included in analysts’ workloads to assess their performance” (Garrett, 2017, p. 985). Garrett properly notes the independence of the lab’s board of directors from the police. But he seems to have missed the importance of broad governance of the lab. The lab’s Certificate of Formation, its most basic founding document, requires that its nine-member Board of Directors be “appointed by the Mayor of the City (the ‘Mayor’) and confirmed by the City Council as evidenced by a resolution approved by a majority vote” (Houston, 2012).

Ultimate governance of the lab is in the hands of the elected representatives of the people of Houston. And that arrangement has been successful. Changes in the governance of the Houston lab have not, of course, spirited away human fallibility or eliminated all problems. In 2018, for example, the lab fired “a crime scene investigator who violated policy by using unapproved equipment that resulted in false negatives for biological evidence in at least two sexual assault cases” (Ketterer, 2018). But the changes, including broader governance, have improved the lab’s work greatly. A similar arrangement for SAGE likely would improve its performance as well.

6 Closing remarks

My suggestions for reforming SAGE are summarized in Table 2.

My suggestions will be useful if they can serve as the basis of further and more detailed deliberations among a broad group of heterogeneous stakeholders. If a lynchpin principle can be found among my four reform principles, it is broad governance. With broad governance, infirmities have a reasonable chance of being identified and acted upon. Without broad governance, we may get more form than substance on the remaining principles.

Democratic polities would not be better off sending their experts to hell. We need experts and we should value expertise. But “Science is the belief in the ignorance of experts” (Feynman, 1969, p. 320). We should cultivate and explore different opinions in science and policy, remembering that contestation is the lifeblood of science. Because governments act in crises, it matters what they learn from their experts. Will it be uniform orthodoxy? Or will it be a full pallet of rich, diversified, and contested scientific opinion? I cast my vote for contestation and against conformity.

Acknowledgements I thank Harry Richer for crucial guidance on the structure and organization of SAGE. I thank Rok Novak, Jon Murphy, Iliia Murtazashvili, Barkley Rosser, and Peter Boettke for help on Hayek and knowledge markets. I thank David Harper for a helpful discussion of SAGE and public financial management. I thank Brian Albrecht, Nick Cowen, Stephen Davies, John Davis Mark Koyama, Luc Marest, Tucker Omberg, and Mario Rizzo for comments given at the NYU conference on Political Economy of Public Health. William Shugart provided help with English grammar and style. Some passages of my exposition

of information choice theory are lifted from Koppl (2018, 2019) with little or no alteration. My discussion of SAGE lifts a few passages from Koppl (2020). My epigraph is from Feynman (1969).

References

- AAAS. (2017). AAAS CEO Rush Holt and 15 other science society leaders request climate science meeting with EPA Administrator Scott Pruitt (Letter to Scott Pruitt), 31 July 2017. Downloaded 13 April 2021 from <https://www.aaas.org/statements>
- Amanpour and Company. (2020). Why don't Americans trust experts anymore? *YouTube*, 15 May 2020. https://www.youtube.com/watch?v=V50HMNB4qX4&feature=share&fbclid=IwAR1jp8vWbNvVdMWGbO1VZqD-SbZtK_eibt1J_Nxq3g1fqKSIYQw47FVPGbI. Viewed 14 April 2021.
- Anderson, R. C., & Pichert, J. W. (1978). Recall of previously unrecalable information following a shift in perspective. *Journal of Verbal Learning and Verbal Behavior*, 17, 1–12.
- Anderson, R. C., Pichert, J. W., & Shirey, L. L. (1983). Effects of reader's schema at different points in time. *Journal of Educational Psychology*, 75(2), 271–279.
- BIT. (n.d.). *About us*. Downloaded 14 April 2021 from <https://www.bi.team/about-us/>
- Britton, T., Ball, F., & Trapman, P. (2020). A mathematical model reveals the influence of population heterogeneity on herd immunity to SARS-CoV-2. *Science*, 369(6505), 846–849.
- Budryk, Z. (2021). "CNN's Gupta 'stunned' Cuomo said he doesn't trust health experts. *The Hill*, 2 February 2021. Downloaded 9 February 2021 from <https://thehill.com/homenews/media/536923-cnns-gupta-stunned-cuomo-said-he-doesnt-trust-health-experts?rl=1>
- Cabinet Office. (2010). *Responding to emergencies, the UK central government response: Concept of operations*. Downloaded 12 April 2021 from https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/192425/CONOPs_incl_revised_chapter_24_Apr-13.pdf
- Cabinet Office. (2012). *Enhanced SAGE guidance: A strategic framework for the Scientific Advisory Group (SAGE)*. Civil Contingencies Secretariat.
- Cabinet Office. (2020). *COVID-19 winter plan*. Downloaded 15 April 2020 from <https://www.gov.uk/government/publications/covid-19-winter-plan/covid-19-winter-plan>
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19.
- Dasaratha, K. (2020). *Virus dynamics with behavioral responses*. ArXiv [arXiv:2004.14533](https://arxiv.org/abs/2004.14533)
- Davis, R. H. (1963). Arms control simulation: The search for an acceptable method. *The Journal of Conflict Resolution*, 7(3), 590–602.
- Dbouk, T., & Drikakis, D. (2021). Fluid dynamics and epidemiology: Seasonality and transmission dynamics. *Physics of Fluids*, 33, 021901.
- Debreu, G. (1974). Excess-demand functions. *Journal of Mathematical Economics*, 1, 15–21.
- DeFries, R., & Nagendra, H. (2017). Ecosystem management as a wicked problem. *Science*, 356(6335), 265–270.
- DeSteno, D. (2020). How fear distorts our thinking about the coronavirus. *The New York Times*, 11 February 2020. Downloaded 9 February 2021 from <https://www.nytimes.com/2020/02/11/opinion/international-world/coronavirus-fear.html>
- Devins, C., Koppl, R., Kauffman, S., & Felin, T. (2015). Against design. *Arizona State Law Journal*, 47(3), 609–681.
- Easterly, W. (2013). *The tyranny of experts: Economists, dictators, and the forgotten rights of the poor*. Basic Books.
- Feigenbaum, S., & Levy, D. M. (1996). The technical obsolescence of scientific fraud. *Rationality and Society*, 8, 261–276.
- Felin, T., Koenderink, J., & Krueger, J. I. (2017). The all-seeing eye, perception and rationality. *Psychonomic Bulletin & Review*, 24, 1040–1059.
- Feynman, R. (1969). What is science? *The Physics Teacher*, 7(6), 313–320.
- Froeb, L. M., & Kobayashi, B. H. (1996). Naive, biased, yet Bayesian: Can juries interpret selectively produced evidence? *Journal of Law, Economics, and Organization*, 12, 257–276.
- Full Fact. (2020). *Chris Whitty did not personally receive millions of pounds from the Bill & Melinda Gates Foundation*, 3 June 2020. Downloaded 8 April from <https://fullfact.org/online/chris-whitty-did-not-personally-receive-millions-pounds-bill-melinda-gates-foundation/>
- Ferguson, N., et al. (2020). *Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand*. <https://doi.org/10.25561/77482>
- Garrett, B. L. (2017). The crime lab in the age of the genetic panopticon. *Michigan Law Review*, 115(6), 979–999.

- Garzarelli, G., & Infantino, L. (2019). An introduction to expert failure: Lessons in socioeconomic epistemics from a deeply embedded method of analysis. *Cosmos and Taxis*, 7(1/2), 2–7.
- Gates, B. (2015). The Next Epidemic—Lessons from Ebola. *New England Journal of Medicine*, 372(15), 1381–1384.
- Gates Foundation. (2020). *Bill and Melinda Gates call for collaboration, continued innovation to overcome challenges of delivering COVID-19 scientific breakthroughs to the world*, 9 December 2020. Downloaded 8 April 2021 from <https://www.gatesfoundation.org/Ideas/Media-Center/Press-Releases/2020/12/Bill-and-Melinda-Gates-call-for-collaboration-innovation-to-deliver-COVID-19-breakthroughs>
- Gates Foundation. (2021). *Gates Foundation COVID-19 response FAQ*. Downloaded 8 April 2021 from <https://www.gatesfoundation.org/ideas/articles/covid19-faq>
- Gomes, M. G. M., Corder, R. M., King, J. G., Langwig, K. E., Souto-Maior, C., Carneiro, J., Gonçalves, G., Penha-Gonçalves, C., Ferreira, M. U., & Aguas, R. (2020). Individual variation in susceptibility or exposure to SARS-CoV-2 lowers the herd immunity threshold. *medRxiv*. <https://doi.org/10.1101/2020.04.27.20081893v3>
- Government Office for Science. (2021). *List of participants of SAGE a related sub-groups*, 29 January 2021. Downloaded 01 April 2021 from <https://www.gov.uk/government/publications/scientific-advisory-group-for-emergencies-sage-coronavirus-covid-19-response-membership/list-of-participants-of-sage-and-related-sub-groups>
- Grinalas, S., & Sprunt, B. (2020). “I trust vaccines. I trust scientists. But I don’t trust Donald Trump,” Biden says. NPR 16 September 2020. Downloaded 9 February from <https://www.npr.org/2020/09/16/913646110/i-trust-vaccines-i-trust-scientists-but-i-don-t-trust-donald-trump-biden-says>
- Harper, M. (2020). Lockdowns cost lives—we need a different strategy to fight Covid-19. *The Telegraph*, 10 November 2020. Downloaded 20 November 2020 from <https://www.telegraph.co.uk/news/2020/11/10/lockdowns-cost-lives-need-different-strategy-fight-covid-19/>
- Hayek, F. A. (1937/1948). Economics and knowledge. In F. A. Hayek (Ed.), *Individualism and economic order* (pp. 33–56). The University of Chicago Press.
- Hayek, F. A. (1945/1948). The use of knowledge in society. In Hayek F. A., *Individualism and economic order*. Chicago: University of Chicago Press.
- Hayek, F. A. (1952). *The sensory order*. The University of Chicago Press.
- Hayek, F. A. (1967). Rules, perception and intelligibility. In F. A. Hayek (Ed.), *Studies in philosophy, politics and economics* (pp. 43–65). The University of Chicago Press.
- Hayek, F. A. (1978). The primacy of the abstract. In *New studies in philosophy, politics, economics and the history of ideas* (pp. 35–49). The University of Chicago Press.
- Head, B. (2008). Wicked problems in public policy. *Public Policy*, 3(2), 101–118.
- Hotelling, H. (1932/1968). Edgeworth’s taxation paradox and the nature of supply and demand functions. *The Journal of Political Economy*, 40(5), 577–616.
- Houston, City of. (2012). *Certificate of Formation, Houston Forensic Science LGC, Inc.* Downloaded 13 April 2021 from [https://www.houstonforensicscience.org/resources/\\$1\\$Ceq3tCy9\\$yvywgiaibdoO0S.Sedtx.pdf](https://www.houstonforensicscience.org/resources/1Ceq3tCy9$yvywgiaibdoO0S.Sedtx.pdf)
- Hymas, C. (2020). Revealed: Sir Patrick Vallance has £600,000 shareholding in firm contracted to develop vaccines. *The Telegraph*, 23 September 2020. Downloaded 6 April 2021 from <https://www.telegraph.co.uk/news/2020/09/23/revealed-sirpatrick-vallance-has-600000-shareholding-firm-contracted/>.
- Ikeda, S. (1997). *Dynamics of the mixed economy: Toward a theory of interventionism*. Routledge.
- Ioannidis, J. P. A., Cripps, S., & Tanner, M. A. (2020). Forecasting for COVID-19 has failed. *International Journal of Forecasting*. <https://doi.org/10.1016/j.ijforecast.2020.08.004>
- Institute for Government. (2021a). *Devolution and the NHS*. Downloaded 14 April 2021 from <https://www.instituteforgovernment.org.uk/explainers/devolution-nhs#:~:text=After%20devolution%20Scotland%2C%20Wales%20and,which%20patients%20pay%20for%20prescriptions>
- Institute for Government. (2021b). *COBR (COBRA)*. Downloaded 14 April 2021 from <https://www.instituteforgovernment.org.uk/explainers/cobr-cobra>
- International Working Group on Managing Public Finance (Edward Hedger, Roger Koppl, Florence Kuteesa, Nick Manning, Barbara Nunberg, Jana Orac, Allen Schick, Paul Smoke, and Duvvuri Subbarao). 2020. *Advice, money, results: Rethinking international support for managing public finance*. New York University, Robert F. Wagner Graduate School of Public Service. <https://wagner.nyu.edu/advice-money-results/about>
- Ketterer, S. (2018). Houston crime lab fires investigator after alleged testing policy violation. *Houston Chronicle*, 26 October 2018. Downloaded 13 April 2021 from <https://www.chron.com/news/houston-on-texas/houston/article/Houston-forensic-lab-fires-investigator-after-13338820.php>
- Klement, R. J. (2020). The SARS-CoV-2 crisis: A crisis of reductionism? *Public Health*, 185, 70–71.

- Koppl, R. (2005). How to improve forensic science. *European Journal of Law and Economics*, 20(3), 255–286.
- Koppl, R. (2018). *Expert failure*. Cambridge University Press.
- Koppl, R. (2019). Response paper (in “Symposium on Roger Koppl’s *Expert failure*”). *Cosmos + Taxis*, 7(1–2), 73–84.
- Koppl, R. (2020). We need a market for expert advice, and competition among experts, 3 November 2020. *Institute of Economic Analysis blog*, <https://iea.org.uk/we-need-a-market-for-expert-advice-and-competition-among-experts/>.
- Koppl, R., & Cowan, J. E. (2010). A battle of forensic experts is not a race to the bottom. *Review of Political Economy*, 22(2), 235–262.
- Koppl, R., & Krane, D. (2016). (2016). Economic incentives and other barriers to blinding in forensic science. In C. T. Robertson & A. S. Kesselheim (Eds.), *Blinding as a solution to bias: Strengthening biomedical science, forensic science, and law*. Academic Press.
- Koppl, R., Kurzban, R., & Kobilinsky, L. (2008). Epistemics for forensics. *Episteme: Journal of Social Epistemology*, 5(2), 141–159.
- Lamberts, R. (2017). Distrust of experts happens when we forget they are human beings. *The Conversation*, 11 May 2017. Downloaded 9 February 2021 from <https://theconversation.com/distrust-of-experts-happens-when-we-forget-they-are-human-beings-76219>
- Levy, D. M., & Peart, S. J. (2017). *Escape from democracy: The role of experts and the public in economic policy*. University of Cambridge Press.
- London School of Hygiene & Tropical Medicine. (2008). *London School of Hygiene celebrates new \$59 million Gates funding*. Downloaded 14 April 2021 from <https://www.gov.uk/government/people/patri-ck-vallance>.
- Mandeville, B. (1729/1924). *The fable of the bees: Or, private vices, publick benefits*, with a Commentary Critical, Historical, and Explanatory by F. B. Kaye, in two volumes. Clarendon Press.
- Mantel, R. (1974). On the characterization of aggregate excess-demand. *Journal of Economic Theory*, 7, 348–353.
- McPhillips, D. (2020). Gates Foundation donations to WHO nearly match those from U.S. government. *U. S. News & World Report*, 29 May 2020. Downloaded 8 April 2021 from <https://www.usnews.com/news/articles/2020-05-29/gates-foundation-donations-to-who-nearly-match-those-from-us-government>
- Milgrom, P., & Roberts, J. (1986). Relying on information of interested parties. *RAND Journal of Economics*, 17, 18–32.
- Murphy, J. (2021). *Cascading expert failure*. SSRN. Downloaded 12 April 2021, from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3778836
- Murphy, J., Devereaux, A., Goodman, N., & Koppl, R. (2021). Expert failure and pandemics: On adapting to life with pandemics. *Cosmos + Taxis* (forthcoming).
- Nelson, R. R., & Winter, S. G. (1982). *An evolutionary theory of economic change*. The Belknap Press of Harvard University Press.
- Neustadt, R. E. & Fineberg, H. E. (1978/2009). *The swine flu affair*. The National Academies Press.
- Parliament. (2010). *Memorandum submitted by the Health Protection Agency* (SAGE 28). Downloaded 14 April 2021 from <https://publications.parliament.uk/pa/cm201011/cmselect/cmsstech/498/498we11.htm>
- Peltzman, S. (1975). The effects of automobile safety regulation. *Journal of Political Economy*, 83, 677–725.
- Perrow, C. (1984/1999). *Normal accidents: Living with high risk systems*. Princeton University Press.
- Pichert, J. W., & Anderson, R. C. (1977). Taking different perspectives on a story. *Journal of Educational Psychology*, 69(4), 309–315.
- Rayner, G. (2021). Role of Sage to be reviewed over fears scientists hold too much power. *The Telegraph*, 15 March 2021. Downloaded 16 March 2021 from https://www.telegraph.co.uk/news/2021/03/15/role-sage-reviewed-fears-scientists-hold-much-power/?WT.mc_id=tmgliveapp_iosshare_Aw1KjZQzGVf2
- Ridley, M., & Davis, D. (2020). Is the chilling truth that the decision to impose lockdown was based on crude mathematical guesswork? *The Telegraph*, 10 May 2020. Downloaded 7 April 2021 from <https://www.telegraph.co.uk/news/2020/05/10/chilling-truth-decision-impose-lockdown-based-crude-mathematical/>
- Rittel, H. W. J., & Weber, M. M. (1973). Dilemmas in a general theory of planning. *Policy Science*, 4(2), 155–169.

- SAGE. (2020). *The Scientific Advisory Group for Emergencies (SAGE)*. Downloaded 10 February 2021 from https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/900432/sage-explainer-5-may-2020.pdf
- SAGE. (n.d.). *About us*. Downloaded 14 April 2021 from <https://www.gov.uk/government/organisations/scientific-advisory-group-for-emergencies/about>
- Sasse, T., Haddon, C., & Nice, A. (2020). *Science advice in a crisis*. Institute for Government. Downloaded 9 April 2021 from https://www.instituteforgovernment.org.uk/sites/default/files/publications/science-advice-crisis_0.pdf
- Schick, A. (1998). Why most developing countries should not try New Zealand's reforms. *The World Bank Research Observer*, 13(1), 123–131
- Scott, J. C. (1998). *Seeing like a state*. Yale University Press.
- Schutz, A. (1945). On multiple realities. *Philosophy and Phenomenological Research*, 5(4), 533–576.
- Schwab, T. (2020). While the poor get sick, Bill Gates just gets richer. *The Nation*, 5 October 2020. Downloaded 8 April 2021 from <https://www.thenation.com/article/economy/bill-gates-investments-covid/>.
- Smith, V. (2009). *Rationality in economics: Constructivist and ecological forms*. Cambridge University Press.
- Sonnenschein, H. (1972). Market excess-demand functions. *Econometrica*, 40(3), 549–563.
- Sonnenschein, H. (1973). Do Walras' identity and continuity characterize the class of community excess-demand functions? *Journal of Economic Theory*, 6, 345–354.
- Stevenson, C. (2020). Starter or substantial meal? A curious case of Schrodinger's scotch egg. *The Independent*, 1 December 2020.
- Taylor, A. (2020). Are scotch eggs snacks—Or meals? As British pubs emerge from coronavirus lockdown, they must decide. *The Washington Post*, 1 December 2020.
- Thomas, C. J., & Deemer, W. L., Jr. (1957). The role of operational gaming in operations research. *Operations Research*, 5(1), 1–27.
- Treasury Committee. (2021). *Economic impact of coronavirus: Gaps in support and economic analysis*. House of Commons.
- Trogen, B., Caplan, A. (2021). Risk compensation and COVID-19 vaccines. *Annals of Internal Medicine*, 174(6), 858–859.
- Turner, S. (2001). What is the problem with experts? *Social Studies of Science*, 31(1), 123–149.
- Turner, S. (2010). Normal accidents of expertise. *Minerva*, 48, 239–258.
- Uexküll, J. (1934/2010). *A foray into the world of animals and humans: With a theory of meaning*. University of Minnesota Press.
- Unattributed. (2017). Scientific groups question Pruitt's 'red team' climate review. *Inside EPA's Risk Policy Report*, 24(32), 2.
- Welsh Government. (n.d.). *Technical Advisory Cell, Terms of reference*. Downloaded 14 April 2021 from <https://gov.wales/technical-advisory-cell/terms-reference>
- Wells, G., Williams, S., & Davies, S. C. (2011). The Department of Health perspective on handling uncertainties in health sciences. *Philosophical Transactions of the Royal Society a: Mathematical, Physical and Engineering Sciences*, 369(1956), 4853–4863.
- Whitman, D. G., & Koppl, R. (2010). Rational bias in forensic science. *Law, Probability and Risk*, 9(1), 69–90.
- Whiting, A. S. (1972). The scholar and the policy-maker. *World Politics*, 24(Supplement), 229–247.
- Wittgenstein, L. (1958a). *The blue and brown books*. Harper & Row, Publishers.
- Wittgenstein, L. (1958b). *Philosophical investigations* (3rd ed.). Macmillan Publishing Co., Inc.
- Wolf, C., Jr. (1962). Defense and development in less developed countries. *Operations Research*, 10(6), 828–838.
- Yeager, L. B. (1960). *Methodenstreit* over demand curves. *Journal of Political Economy*, 63, 53–64.
- Young, A. A. (1928). Increasing returns and economic progress. *Economic Journal*, 38(152), 527–542.
- Zuccino, D. (2006). Duke case worsens for prosecution. *Los Angeles Times*, December 16, 2006. Downloaded August 1, 2016 from <http://articles.latimes.com/2006/dec/16/nation/na-duke16>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.