Steve Fuller*

Expertise as a Form of Knowledge: A Response to Quast

https://doi.org/10.1515/auk-2020-0017

Abstract: Christian Quast has presented what he describes as a 'role-functional' account of expertise as a form of knowledge that purports to take into account prior discussions within recent analytic social epistemology and allied fields. I argue that his scrupulousness results in a confused version of the role-functional account, which I try to remedy by presenting a 'clean' account that clearly distinguishes such an account from what Quast calls a 'competence-driven' one. The key point of my account is that 'competence' pertains to knowledge in closed systems and 'expertise' in open systems. I observe that the invocation of 'reliability' as an epistemic standard simply serves to confuse the difference between the competence-driven and role-functional accounts.

Keywords: expertise, competence, reliability, social epistemology, client, lay

Christian Quast's 'How expertise ascriptions work' unwittingly illustrates the ease with one can lose sight of the phenomenon that one wishes to explain, in spite—or because—of the amount of debate that has taken place over it. In that spirit, let us begin by recalling Quast's core thesis, which constitutes his 'new and improved' definition of expertise:

> More precisely, this paper introduces two different kinds of contextuality by advancing and advocating the thesis that expertise ascriptions are true if and only if their content within their context of use is true against standards in the context of assessment. This means that expertise ascriptions have indexical content and are also assessment sensitive. (Quast 2020, 399)

If you read this statement without being told that it was about 'expertise', what would you guess is being defined by it? You might guess Aristotle's *phronesis* or, in a more modern vein, 'information', as in 'the difference that makes the difference' that resolves an open-ended situation to one's satisfaction. Both guesses respect the sense of epistemic relativity implied in the definition. But would you guess

***Corresponding author: Steve Fuller,** Department of Sociology, University of Warwick, Warwick, UK, e-mail: profstevefuller.net

'expertise'? Probably not. This is not because what the definition says is irrelevant to the nature of expertise. Indeed, it is quite relevant, but it fails to distinguish expertise from any other situated form of knowledge. To be sure, Quast's article conjures up most of the missing ingredients for a proper account of expertise, but they are raised *ad hoc* as Quast attempts to distinguish his own position from that of others who have tried to define expertise in recent times. In contrast, I shall begin in the *first section* of this paper by providing a relatively 'clean' account of expertise that does justice to it as both a sociological and an epistemological phenomenon, without detouring into the recent debates with which Quast is preoccupied. Afterwards, in the *second section*, I shall assess his position in light of my account.

1 A Social Epistemology of Expertise

First, while expertise certainly possesses a socially situated character, the 'situation' is largely defined by 'expert', whose expertise is grounded not in the situation itself but in the expert's professional accreditation, which typically involves specialised training. However, this training may or may not have consisted of prior experience with the situation. In this respect, the etymological root of 'expert' in 'experienced' sends an equivocal signal. The 'prior experience' may consist in the expert's education and/or direct acquaintance. Thus, the knowledge on which experts draw is alternatively cast as based on 'templates' or 'precedents', allowing for both a rationalist and an empiricist basis for the epistemology of expertise, respectively. Moreover, the ambiguity strengthens the expert's hand in justifying his or her own practice. Even if the expert has never directly encountered the current situation, s/he may claim that it belongs to a kind of situation that others in the expert's field have previously encountered. In this way, the unfamiliar is rendered familiar to the expert in a way that secures the confidence of the client.

What I have presented above is no mere social constructivist nicety. Anyone who contracts the services of an expert effectively licenses the expert to exercise discretion over the matter that concerns the client, which prior to the expert's intervention remains vaguely defined yet no less urgent to the client (Fuller 1988, ch. 12). The client's expectation is that the expert will define the situation in a way that addresses the client's concern. This will typically involve approaching the situation in ways that differ from the client's default modes of epistemic access. This shift in perspective aims to convert the situation of concern into a 'soluble problem'. The 'success' of the expert-client transaction is judged primarily in terms of the client's acceptance that the expert has made a 'good faith' attempt to solve

the client's problem. Whether the client's problem is actually solved—or the client simply comes to understand the nature of his or her situation better—is a secondary concern. Indeed, if the client wishes to contest the expert's handling of the client's problem, then other experts of the same kind need to be engaged in the ensuing litigation to determine the occurrence of any 'malpractice'. Malpractice is not something that clients can judge for themselves without additional expert input. However, if the expert is found guilty of malpractice, then the expert may be formally expelled from the professional peer group. In that case, the knowledge possessed by this defrocked expert would no longer count as expertise, even though the content of the defrocked expert's knowledge would not have changed.

This account of expertise does not fit comfortably within the competing stereotypes of knowledge in modern epistemology: *knowing how* and *knowing that*, the former associated with practices and the latter with propositions. The difference lies in the ontology of knowledge and hence the mode of access properly called 'epistemic' with regard to the relevant objects. Expertise involves a rather different approach to epistemology—a different ontology of knowledge, if you will—one that is perhaps most familiar from the religious sphere. Here we need to recall that when 'expert' started to be used regularly in legal proceedings (i.e. 'expert witnesses') in Third Republic France, the non-expert was called a 'lay' person, a word whose implications philosophers appear to be curiously oblivious to, even though it is still regularly used in public discussions of expertise. The original suggestion was that experts constituted a secular clergy, perhaps even a replacement for the Roman Catholic version in terms of their authority. To be sure, this characterisation has not always been made to the advantage of the experts over the clergy, but it is worth dwelling on why the comparison has stuck over the years to understand the distinctive epistemology of expertise.

Perhaps the most intuitively straightforward way to get to the heart of the matter is to compare a priest and, say, a physician or psychiatrist, focusing on their epistemically relevant similarities. Much of what the secular experts say and do could be—and have been done—by priests in the past and, in some countries, even to this day. Moreover, in terms of 'knowing how' and 'knowing that', there is no reason to think that the reliability of one group has been better than the other, with regard to the efficacy or truth of what they have done or said. We simply lack a consistent track record of what might be called 'client satisfaction', which takes seriously the client's own judgements of their transactions with experts. After all, alongside the many homoeopathy patients who end up accepting their death from cancer are the many chemotherapy patients who similarly accept their fate. Both groups may be satisfied with the choices they made, even if third party social epistemologists are not. Moreover, it is not clear whether there is a specifically

'epistemological' problem here—or a more straightforward 'cultural' problem about how people should conduct their lives.

In any case, we live in a world in which a wide variety of drugs and treatments can be administered *only* by qualified medical practitioners, even though others not so qualified may display the same level of competence in 'knowing how' and 'knowing that' with regard to such matters. Here I mean people who act just as the experts would in the relevant conditions but they lack expert authorization. In a sense, it is the complement of the 'placebo effect', whereby people claim to be cured with 'fake drugs' because someone they regard as an expert has prescribed them. In contrast, I mean non-experts who prescribe what the relevant experts would, but are not trusted (as much) because they are not licensed as experts. To be sure, sometimes these people successfully masquerade as experts, at least until they overplay their hand by making a serious practical error or seeking to leverage their pseudo-expertise. This 'overplaying' amounts to the pseudo-expert provoking an investigation into credentials that s/he would otherwise have been presumed to possess.

At this point, recall our earlier observation that even when an expert makes errors that are so serious that the client files a claim of malpractice, the expert's peer community plays a significant role in determining the liability of the charged expert. In effect, the possession of expert credentials may serve to shield the practitioner from forms of punishment to which the pseudo-expert is automatically liable, which may then be amplified by any deception involved. Yet, the cognitive error and the material harm may be the same. What accounts for the difference—and should there be a difference? Here social epistemologists have been inclined to appeal to something called 'trust', which functions mainly as a euphemism for a kind of risk-taking, whereby a portion of one's own sphere of judgement is forfeited to someone else who is presumed to be more capable of acting on that person's behalf. 'Forfeited' is used deliberately to convey the fact that the person taking the risk recognizes their own incapacity to address a matter of concern to them. Political scientists and economists, closely following the conceptual framework of the law, characterize the client-expert relationship as one of 'principal-agent', which captures well the voluntary subordination of the will that is involved in clients' 'trust' of experts (Ross 1973).

In religious times, people trusted priests because of their faith in 'God', however defined. Nowadays people trust medical practitioners because of their faith in 'Science', however defined. In both cases, 'truth' is the philosophical term of art used to cover the object of the faith shared by the principal and the agent, with the proviso that the relationship between the principal's trusted agent and that larger object of faith is bound to be imperfect, yet epistemically superior to the principal's own relationship to the object, especially with regard to the matter of immediate

concern to the principal. The epistemically distinctive feature of expertise, then, is the *distributed* nature of the process of knowing, whereby the principal knows *through* the agent, in both of the main senses of 'knowing' in modern epistemology: the 'knowing that' something is true and 'knowing how' to apply the truth to reach a desirable practical outcome.

The original context for 'principal-agent' theory is relevant for understanding what exactly is 'distributed' here. Clearly agency is distributed, which may invite thoughts of an 'extended mind'. But materially speaking, *risk* is distributed, such that the principal aims to minimize the cost of personal misjudgement by placing a bet on the agent's chance of making a better judgement on the principal's behalf. In effect, what social epistemologists substantively mean by 'trust in experts' is a version of 'risk pooling' in the insurance and financial trades. Both Descartes and Pascal, in their different ways, were arguing against engaging in such activities: One should bear the risks for oneself, with Descartes being somewhat more bullish than Pascal about the likely outcome. Hence, while both were avowed Christians, they were widely seen in their day as anti-clerical: They would rather place their faith (or, in Descartes' case, 'thinking') in God directly than in a priest who then exercises that faith on their behalf.

In this sense, the expert is in a double-sided relationship of 'representation': on the one hand, to the client for whom the expert is a trusted surrogate and, on the other hand, the object of knowledge to which the expert must remain loyal. I do not wish to argue conclusively here whether 'representation' is used univocally or equivocally in these two contexts. Suffice it to say, social constructivists (myself included) follow in Hobbes' footsteps in treating the two uses *univocally*, which in turn reflects the long historic drive toward increasing literalism in language, starting with Augustine and the Franciscan scholastics (Duns Scotus, Occam, etc.), running through the Protestants, Bacon, Descartes, Pascal and their secular progeny, not least Rousseau. In practice, contrary to its contemporary connotations, this 'literalism' has fostered the view that one's language should reflect what one thinks because speaking is an act of self-authorization: You should not say it unless you mean it. Here words like 'belief' and 'representation' function as euphemisms for this more potent idea.

All of the above philosophers and theologians are haunted by the creativity of God's Word (*logos*), which 'literally' applies to humans as having been created *imago dei* (in Augustine-speak), notwithstanding the Fall recounted in Genesis. The modern preoccupation with the use of language to authorize control in both the legal and scientific spheres is its secular descendant, 'positivism' being its most self-conscious expression (Turner 2010, ch. 3; cf. Passmore 1961). The arc of Anglo-German thought from Mill and Russell to Wittgenstein and Kripke made this concern central to what became 'analytic philosophy' in the twentieth century. It

also helps to make sense of the persistent controversies surrounding expertise as a form of knowledge.

The expert's discretion to define the client's situation for purposes of ministering to it amounts to a *power to name*, understood as a pretext for actions of a certain sort that the expert advises or takes on behalf of the client in an otherwise indeterminate situation. More to the point, the expert purports to resolve this indeterminacy by replacing what the client has identified through a 'proper name' (i.e. something unique) with a 'definite description' (i.e. a specific complex of repeatable properties). The question then is whether that definite description—or even some succession of such descriptions—is likely to be exhaustive of the situation originally identified by the client. Practically speaking, the answer is bound to vary from case to case. Theoretically speaking, resistance to expert advice may be seen as akin to the refusal to reduce proper names to definite descriptions.

In the philosophy of science, this sensibility is reflected in historically informed doubts that the truth about some 'rigidly designated' (in Kripke-speak) part of reality is likely to be reached by successively improved versions of the descriptors used in current authoritative accounts of that reality—Hilary Putnam's 'pessimistic meta-induction' (Putnam 1978). Philosophers routinely call this position 'realism', but it is more Platonic than Aristotelian in inspiration: It presumes a conception of reality to which we may gain access but not necessarily by our currently established means, which include those whom we now deem 'expert'. Truth be told, it is also the implied metaphysics of 'the end justifies the means', which licenses even substantial deviations from established expertise if one—and perhaps more importantly, one's 'society', however defined—is willing to absorb the risks involved.

2 Quast's Social Epistemology of Expertise in Light of My Account

Quast's account of expertise contains several features of the above account. First, he rightly observes that the nature of expertise is bound up with expertise-ascriptions, and that it is more 'role-functional' than 'competence-driven', in his terminology. He is also right that the 'role-functional' approach to expertise stresses the expert's authority and responsibility. However, Quast doesn't quite wish to renounce the competence-driven account altogether, since he believes that expertise involves some sort of 'reliability', the Holy Grail of naturalistic approaches to epistemology. It turns out that on his account, an expert reliably makes judgements and/or performs actions within a specific domain that meets

the client's standards. So, the expert is competent, after all, but what is at issue is the nature of the knowledge that constitutes this hybrid sense of 'competence'. It is here that Quast's account becomes confused. He would have done better to avoid talk of reliability and competence altogether, which tends to go against what he calls the 'role-functional' account, as we shall see below.

The very idea that expertise should be seen as something above and possibly beyond competence arose in the 1980s and was associated with the drive to automate complex decision-making in so-called 'expert systems', which remain a staple in the 'knowledge management' field in business schools (Fuller 2002, ch. 3). It was a response to the long proven failure of human experts—starting with Paul Meehl's research in the 1950s—to perform as well as mechanical counterparts in diagnosing various medical and psychiatric disorders, at least as judged by professional handbook standards. In epistemological terms, it had been demonstrated that human expert judgement is 'unreliable' in real world settings (Faust 1984). This led ethnographers to interview human experts to understand how they would process cases as they made their decisions under a variety of hypothetical conditions. Implied in this strategy was that the experts' basic approach was sound but that it ran into difficulties when they reached their 'natural limits'. 'Natural limits' should be understood to mean limits to both one's professional training and processing capacities, including hot and cold cognitive biases. The intended result of this research was an 'expert system', which consisted of a user-friendly computer interface informed by a decision-tree-styled algorithm, the design of which was based on the expert interviews. It inspired many cognitive scientists and philosophers of science, myself included, to countenance that a sufficiently advanced form of artificial intelligence may be the true reference class of the various qualitative and quantitative accounts of 'rationality' that philosophers have historically proposed (Fuller 1993).

The backlash against this entire strategy, philosophically inspired by Hubert Dreyfus and still informing Harry Collins' approach to expertise, was to argue that any such 'expert system' would always be insufficient to replace the human expert (Dreyfus/ Dreyfus 1986; Collins 1990). The backlashers claimed that human performance in the relevant domain minus the above 'natural limits' would always be better than the 'debugged' human represented in the computer algorithm programming the expert system. Their stance reflected a larger background concern—namely, that an advanced form of artificial intelligence might significantly supersede human performance in the sorts of complex cognitive tasks traditionally seen as the exclusive preserve of humans. Unsurprisingly perhaps, it resulted in a sharper distinction being drawn between 'competence' and 'expertise', almost as a proxy for 'machine' and 'human'. The intended contrast was between relatively routine domain-specific judgements, which a well-programmed machine might

deliver, and more 'creative' judgements that may suspend some of the problem-solving constraints governing the domain, but without completely undermining the domain's epistemic framework.

Nevertheless, this did not deter artificial intelligence enthusiasts following the lead of Herbert Simon (myself included), who believed that the competent-expert distinction could be deployed to capture the difference between, say, on the one hand, Kuhn-style 'normal scientists' in classical physics who were incapable of thinking outside of their paradigm to solve the long-standing problems facing Newtonian mechanics, and on the other hand, Einstein, Heisenberg and the other early twentieth century revolutionaries who managed to radically transform physics without destroying it altogether (Nickles 1980; Langley et al. 1987). Subsequent research drew the implied distinction in less world-historic terms, but the intuition guiding it is clear enough. It is one thing to acknowledge that Einstein and his comrades dealt with outstanding physics problems in a different frame of mind from their more classical colleagues, and another to say that the success of their approach amounted to their 'knowing something more' than their classical colleagues, in some univocal sense of 'know'. It would be reasonable to grant the former but deny the latter. Indeed, a classical physicist such as Henri Poincaré was probably *more* competent than Einstein by the academic standards of the day, yet that did not prevent Einstein from proving more expert in accounting for relative motion. These are good *prima facie* grounds for concluding that competence and expertise bear some sort of *orthogonal* relationship to each other. Let me briefly try to tease out the nature of this 'orthogonality'.

What unites this world-historic case of 'expertise' with more ordinary cases involving, say, doctor-patient is an acceptance that the decision-making context is *open*. In other words, the 'normal' ways of making sense of the situation are suspended—to an extent that remains to be specified—so that certain 'abnormal' ways of dealing with it are licensed. Due to the normalization of expertise in contemporary society, it is easy to forget the alien—perhaps even 'incommensurable'—character of how a doctor typically approaches a patient. Nevertheless, patients license that alien treatment because they have come not to trust their own judgement on matters of health, notwithstanding their personal nature. Similarly, what Thomas Kuhn (1970) called 'revolutionary science' is made possible once normal scientists take an estranged stance toward their 'normal science' practices because of their failure to solve long-standing puzzles of central concern to them. Kuhn calls this shift in orientation a 'crisis'—and it reflects a recognition of the limits of what heretofore had passed for 'competence'. To be sure, the intuitiveness of Einstein's superior expertise depends on a *retrospective* evaluation of the different problem-solving approaches taken by him and Poincaré—that is, after their uptake

by the relevant peer community in physics. By that standard, Einstein succeeded where Poincaré had failed. I shall return to this point shortly.

What the above suggests is that expertise, rather than being an incremental or even a step-change advance on competence, operates in a radically different cognitive space. Competence is about knowledge in a closed system, and expertise in an open system. Quast implicitly makes the point in his time traveller thought experiment, the conclusion of which is that the traveller retains her competence while her expertise depends on whether an appropriate context of use figures in the world she visits. Instead of Quast's time traveller, one could imagine Einstein himself travelling back in time, even just to the start of the nineteenth century. Would his approach to problem-solving be seen as expert or crazy? The answer would depend on the likelihood that the corresponding peer community would consolidate around Einstein's treatment of light as a constant in the understanding of motion. Alternatively, fast forward into the future that Simon and other AI enthusiasts have envisaged—one in which an Einstein-like computer could not only adjust the parameters of the variables in its programme but change variables into constants, and vice versa, resulting in a substantially different programme. Would such a 'superintelligent' machine capable of projecting paradigm shifts whole cloth be regarded as a salutary revolutionary agent or a threat to the entire scientific enterprise, if not the human condition generally? Much will depend on both first-order views about the state of science and second-order views about the conduct of science at the time such a machine becomes available.

The temporal character of expertise evaluations points to the inappropriateness of 'reliability' as a standard for judging experts. As a methodological concept associated with what is sometimes called 'internal validity', reliability is about the regularity with which the same conditions bring about the same effects. The term 'mechanism' is often used both literally and metaphorically to describe something that encompasses the 'reliable' relationship. However, if the conditions are not fully specified, then it is not possible to establish that relationship. Yet that is precisely the sort of situation in which a client would engage an expert—and part of that engagement would involve granting a license to the expert to complete the specification of the conditions, which in turn will circumscribe the interpretation and treatment that constitute the response. Insofar as a specific level of competence is required for expertise, its evaluation occurs far from the typical context of use. I refer here to the process by which experts acquire professional credentials, which may involve passing specific academically and practically oriented examinations. In addition, matters of competence may be central to the adjudication of a malpractice suit against an expert—but again, with the expert peer community playing a crucial determining role. However, if the expert receives no formal complaints in

the aftermath of an engagement with a client, then the expert is presumed to have been competent, however s/he acted.

In sum, while Quast is certainly correct to prefer a role-functional account of expertise, his residual attachment to the competence-driven account is unnecessary and potentially confusing. What passes for competence in expertise is really the background support of the expert peer community, who through its own mechanisms, independent of any context of use, have come to invest their trust in the expert in question. Their willingness to vouch for the expert in what are typically open-ended conditions of practice can make all the difference in determining the appropriateness of an expert's actions. The expert community is effectively the corporate underwriter of expertise. Its power is regularly revealed these days by Pope Francis I, who frequently allows high-ranking members of the clergy to be tried for various alleged forms malpractice in civil courts after having withdrawn the immunity of sacred office. This allows the court to judge whether the acts in question would have been appropriate, had they been committed by someone unqualified in the spiritual expertise associated with holding sacred office. Comparable secular examples might involve enquiries into the treatment of (human or animal) research subjects in laboratory experiments without presuming the privilege of the scientific vocation.

My guess is that this crucial aspect of the role-functional character of expertise is missed by analytic social epistemologists, including Quast, due to the confused conception of 'reliability' in naturalized epistemology more generally. The confusion comes from trying to capture at once two senses of 'reliable' in ordinary language: on the one hand, the methodologically relevant idea of regular occurrence, and on the other, the morally relevant idea of trustworthiness that the previous discussion highlighted. As we have seen, these are quite different ideas that normally inhabit different cognitive spaces. But it would require another paper to examine how this confused conception of reliability has wreaked havoc in the recent epistemology literature. A quarter-century ago, I characterised this confusion as *phlogistemology*, in homage to that hallowed pseudo-substance, phlogiston (Fuller 1996).

References

Collins, H. (1990), *Artificial Experts: Social Knowledge and Intelligent Machines*, Cambridge/MA

Dreyfus, H./ S. Dreyfus (1986), *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*, New York

Faust, D. (1984), *The Limits of Scientific Reasoning*, Minneapolis

Fuller, S. (1988), *Social Epistemology*, Bloomington/IN

– (1993), *Philosophy of Science and Its Discontents*, 2nd ed. (Orig. 1989), New York
– (1996), Recent Work in Social Epistemology, in: *American Philosophical Quarterly* 33, 149-166
– (2002), *Knowledge Management Foundations*, Woburn
Kuhn, T. (1970), *The Structure of Scientific Revolutions,* 2nd ed. (Orig. 1962), Chicago
Langley, P./H. Simon/G. Bradshaw/J. Zytkow (1987), *Scientific Discovery: Computational Explorations of the Creative Process,* Cambridge/MA
Nickles, T. (ed.) (1980), *Scientific Discovery, Logic and Rationality*, Dordrecht
Passmore, J. (1961), Hägerström's Philosophy of Law, in: *Philosophy* 36 (137), 143-160
Putnam, H. (1978), *Meaning and the Moral Sciences,* London
Quast, C. (2020), On How Expertise Ascriptions Work, in: *Analyse & Kritik* 42(2), 399–429
Ross, S. (1973), The Economic Theory of Agency: The Principal's Problem, in: *American Economic Review* 63 (2),134-139
Turner, S. (2010), *Explaining the Normative*, Cambridge