Response to the Royal Society Call for Evidence on

The Fruits of Curiosity:
innovation and future sources of wealth

by

The Society for the Study of
Artificial Intelligence and Simulation of Behaviour

Introduction

This response is by the executive Committee of The Society for the Study of Artificial Intelligence and Simulation of Behaviour (SSAISB or AISB), which is the main academic society in the UK for the discipline of Artificial Intelligence.[1] Founded in 1964, the society has an international membership drawn from both academia and industry, and from a wide range of disciplines, such as Philosophy and Psychology, as well as Computing. It is a member of the European Coordinating Committee for Artificial Intelligence (ECCAI).[2]

Some of the ways in which the issues in the *Fruits of Curiosity* enquiry engage AI are shared with ways in which they engage the discipline of Computing. This response therefore leaves some issues to be addressed in responses such as the the one by the UK Computing Research Committee (UKCRC).

We use the name *Computing* in this document to refer to the science concerned with computation, not to the activity of using or managing software tools, or to routine software-tool development that does not require advances in deep knowledge. We use the term in preference to *Computer Science* in order to align our terminology with the UKCRC, the CPHC (Council of Professors and Heads of Computing)[3] and the projected Academy of Computing under the banner of the British Computer Society.[4]

Also, the distinction between Computing on the one hand and Information [and Communication] Technology (IT or ICT) on the other is crucial though unfortunately obscured in the educational, professional and governmental landscapes—see the Question 7 section below, subsection on Education in Schools.

AISB broadly endorses the views expressed in the UKCRC response. The present response concentrates largely on issues that engage in a special and notable way with AI—but the comments that we are led to make nevertheless illuminate the general issues of the consultation, not just AI-specific versions of them.

While Computing is the main discipline within which AI is structurally lodged within universities and funding councils, AI also strongly intersects disciplines such as Philosophy, Brain Science, Psychology and Linguistics. The comments in this response therefore highlight the issue of cross-disciplinarity (not only of the AI sort and but also generally).

[1] `http://www.aisb.org.uk`. Comments on this response should go to the main author and society chair, Prof. John Barnden (J.A.Barnden@cs.bham.ac.uk), School of Computer Science, University of Birmingham, B15 2TT, tel: 0121–414–3816. The response has been informed by valuable input from various Fellows of the society as well as from the society Committee.

[2] `http://www.eccai.org`

[3] `http://www.cphc.ac.uk`

[4] Now known as "BCS, The Chartered Institute for IT" (`http://www.bcs.org`). The Academy was announced in a press release by the BCS Press Office, 21 September 2009. Accessible at time of writing at `http://www.bcs.org/server.php?show=conWebDoc.32253`.

The cross-disciplinarity of AI extends to Humanities as well as scientific disciplines. An example of this is the concern in AI with artistic activity and with creativity more broadly.

Nor is AI confined to emulating purely human qualities (in cognitive, perceptual, emotional, artistic, ... domains): there is a strong interest also in related qualities that are exhibited by other organisms or that emerge through collective activity of organisms.

The sections below and the question numbering used follows the organization and numbering of the questions posed in the Call for Evidence. Only those questions on which we wish to make a distinctive comment are included.

The Long-Term Direction of Policy for Science

Q1: What role should curiosity-driven research play in the UK science base in the next 15-25 years?

Short answer: A huge role.

AI has always been a curiosity-driven research discipline *par excellence.* The main targets of curiosity have been large, beyond-blue-skies questions such as what thought and consciousness are, how the human mind works, and how we could create artefacts that emulate the mind. Yet the discipline has led to: large industrial sectors of global significance (e.g., automated translation of documents, automated mining of product-user opinions expressed in vast collections of documents, intelligent image understanding applied, for instance, to face detection and skin disease diagnosis); systems that affect many if not most ordinary people's direct engagement with life (e.g., intelligence in search engines, with Google in particular being based from the start on AI technology; and AI in computer games, virtual worlds, etc.); and systems of major practical and economic import but that are more hidden from the public (AI technology used in fraud detection by credit card companies, etc.). AI is also of central importance in other areas of science such as space exploration, where automated planetary explorers must intelligently plan their own detailed activities.

AI therefore stands as one example of a discipline where having a basis in pure curiosity goes hand-in-hand with making advances of major economic and societal benefit. Indeed, it would be dangerous to imagine that those benefits could have arisen from some much less curiosity-driven version of the discipline. The more ambitious the aims of a discipline then, as long as deep and detailed work is done in pursuit of them, the more impressive the outcomes (pure and applied) are likely to be. Short-termism and lack of ambition in aims, with too much adherence to current assumptions about the world and to short-term industrial, economic and political goals, is likely to lead to advances of only short-term significance. Connection to industrial and economic goals is of course valuable and has inspired much important, fundamental AI research; rather it is the short-term nature of some goals and of many cycles in industry and elsewhere (e.g., product-development cycles, governmental spending review cycles, planning cycles in universities) that is the issue.

Q4: How should science be governed to maximise benefits to society while acknowledging public questions, uncertainties and concerns?

Short answer: (A) Science should be overseen by independent learned societies that nevertheless include: the pro-active engagement with understandable public concerns as a main aspect of the societies' governance; in particular, alerting of the public to new dangers; and the acknowledgment of scientists as themselves part of the public. (B) Some types of currently unregulated application of AI and other sciences should be regulated by scientifically, politically and industrially impartial ethical bodies where appropriate.

There are legitimate concerns about how science, including AI, may affect the world and in particular how it might be misused, and it is important for the AI and other scientific communities to engage with them and not react defensively. By "legitimate" we include not only concerns that are scientifically justified but also concerns that are perfectly understandable even if not actually justified: the issues may be too difficult for the public to be able to grasp without major effort and without extensive guidance from scientists. While the public may often be misled by misinformed, biased, mischievous or opportunist media coverage, some public concerns, including about AI, are at least understandable because of genuine difficulty, and some concerns are profoundly justified.

For the purposes of this consultation we can put aside science-fiction scenarios of humankind being destroyed or enslaved by humanoid robots, and we can put aside feelings of humiliation when faced with the works of future automated scientists and poets. But the use of AI in ongoing efforts to develop automated, armed guards, soldiers and other military units—or even just automated aids for human soldiers—is surely a legitimate cause for concern.[5] Another legitimate concern is to ensure that AI in mass surveillance of email and voice communications does not lead to misuses and abuses. Other more traditional, legitimate concerns include issues of responsibility and accountability in the use of AI in such areas as Medicine and Law.

A new problem arising, or at least only now revealing its dire effects in the current economic and financial crisis, is that of automated agents operating within real markets, doing automated stock trading or bank risk management, for example. The consequent speed and complexity of transactions, and opaqueness to human understanding, is liable to produce massive, difficult to manage, *qualitatively profound changes in the nature of economic/financial reality*, if it has not done so already. The use of more and more sophisticated AI techniques in such agents can only exacerbate the problems.

The following quotation from an article by Martin Taylor, the former chief executive of Barclays bank, is pertinent to the use of Computing in general, and by extension AI, in banking, at least without appropriate expert human supervision:

> In the last decade, alas, [the function of determining the overall risk appetite, etc.] has come to be seen as "technical," or subject to computer modelling. Human judgment has not been applied, sometimes not even to the decision not to apply human judgment.
>
> Bank board members, usually experienced leaders from other business sectors, are ill-equipped to understand the risk-growth tradeoff. ... They quite literally do not know what they are doing, as much empirical evidence suggests.[6]

It would be perfectly appropriate therefore to have a national ethics committee—impartial scientifically, politically and industrially—that has at least an advisory capacity in the nature and application of AI research, and similar bodies for other currently unregulated research. (A current example of such a body, at least approximately, is the HFEA.[7]) Such a body could, on the applications side of a science, have statutory power to require certain types of public enquiry, legal process or parliamentary process to take place in order to decide whether to allow product-development or product-installation work to proceed when the societal dangers may be great. The body would also need to be in dialogue with other similar bodies and with bodies such as financial regulators, utility regulators (given the emerging use of AI in utility network design and distribution), etc.

Such a body would not only help ensure that AI is used wisely but, by pre-alleviating public concern, ultimately give AI researchers and their collaborators in many disciplines more overall freedom to pursue investigation.

These investigations could even include theoretical or experimental investigations that could in principle lead towards products that would not be legal or would not pass the scrutiny of the ethics committee. After all, and apart from any other consideration, the AI technology involved in such investigations would share much with that needed for parallel developments within virtual worlds (e.g., military units within them). Presumably the regulation of what is developed for virtual worlds should be looser than that for the real-world case. Of course, when the virtual world impinges upon or intersects the real world—such as when valued items in virtual worlds also have real economic value—then the real-world regulation needs suitably to encompass parts of virtual worlds.

The nature of AI itself implies that an ethical body would need to be multidisciplinary, at least in its advice-seeking and preferably in its make-up. Expertise in all the AI-intersected disciplines is potentially needed. But this raises the question of whether ethical bodies for other scientific areas could benefit from being multidisciplinary as well. Being multidisciplinary may be necessary anyway because of the nature of the ethical problems faced, which cannot be relied upon to respect traditional disciplinary boundaries. But as a side-effect the multidisciplinarity could also help to bring down inter-disciplinary barriers within science, as scientists would need to understand more about each others areas in order to engage fully with relevant ethical bodies.

Sciences *must not* be *governed* by government, and it is not clear they should be *governed* by anyone, if this word is taken in any strong sense. They are, or should be, *led* by excellent scientists, with learned societies sometimes also playing an

[5] Such concerns have been articulated by many within AI and without, and notably by Noel Sharkey (Professor of AI and Robotics and Professor of Public Engagement, University of Sheffield). See for instance the Opinion Interview with him in *New Scientist*, 29 August 2009, pp.28–29, and article "Danger, Danger," *Times Higher Education*, 13 August 2009, p.22. (We should stress, however, that the AISB does not endorse Sharkey's sceptical views about the in-principle capabilities of AI.)

[6] At p.35 in "Reflections on the Crunch," *Prospect* magazine, August 2009, pp.34–36.

[7] The Human Fertilisation and Embryology Authority, `http://www.hfea.gov.uk`.

inspirational role but more generally playing facilitatory, oversight, advisory, communicative, and (in limited, specialized respects) regulatory roles. Learned societies should work within a framework of well-conceived, democratically-managed ethical regulation that has statutory operational impartiality towards scientific communities, industry and political groups (while of course taking information and advice from those sectors).

On the other hand, there are issues for the scientific societies and professional bodies themselves. We mention this because of a subtle, unfortunate feature of Q4 as posed, namely the use of the defensive-sounding verb "acknowledge." The verb "address" would have been better. The flavour of Q4 as it stands is of scientific bodies pressing forward somehow within an obstacle course of public objection. But it would be better for scientific bodies to *cooperate visibly, sincerely and proactively* with the public. Indeed, it would be beneficial for the governance of all learned societies to include structures and mechanisms geared to addressing, cooperatively when at all possible, the concerns and uncertainties of the public. (Of course, this will often involve disagreeing with the concern — and therefore suitably educating the public.) It is important that this aspect of the governance should be *clearly evident to the public*.

Some learned societies no doubt already do a good job in these respects, but all too often scientists are at least perceived by the public as being defensive in the face of legitimate concerns, rather than really taking the concerns on board from the start and from the heart.

Quite apart from the case of the public raising a concern and demanding an answer, there is the case of an issue about which the public should be concerned but are not, because they do not yet have the right information. Learned societies (AISB included) should play the proactive role of alerting the public and relevant ethical bodies to dangers discussed among scientists, not just reacting when dangers are pointed out by others.[8] Such a proactive function would itself do much to allay public mistrust of science and scientists.

One point the public may need to be alerted to is that, while fanciful fears about AI tend to focus on physical robots, and while the comments about military AI above were addressed mainly at AI-controlled physical hardware, the more pressing dangers for society at large are arguably in the form of intelligent software agents on the web.

Of course, regulatory bodies can also, and sometimes do, serve an alerting role, cf. the annual Horizon Scanning process performed by Human Fertilisation and Embryology Authority.[9] But we hold that learned societies themselves should also engage in alerts.

Finally, scientists themselves are part of "the public" with regard to areas outside their own, and are capable of exhibiting deep misunderstanding and ignorance of these, despite the commendable efforts of multidisciplinary organizations like The Royal Society. (Hence part of the difficulty of doing cross-disciplinary work.) It behoves the scientific community at large to do more about this problem. Perhaps the lessons learned would help with dealing with the non-scientific sector of the public.

In short, in counterpoint to the often-sung theme of the *Public Understanding of Science*, we need voices carrying three other motifs. First, we need: *Embrace by Scientists of Public Worries about Science*. Secondly, we need: *Alerting by Scientists of the Public about New Worries*. Thirdly, we need: *Understanding by Scientists of Self as Public*.

[8] A similar point has independently been made by Noel Sharkey—see for instance the recent *Times Higher Education* article in footnote 5.

[9] See the press release at `http://www.hfea.gov.uk/5397.html` and *Scientific Horizon Scanning at the HFEA: Annual Report 2008/09* at `http://www.hfea.gov.uk/docs/Horizon_Scanning_Report.pdf`. The Introduction says: "Horizon scanning is an early-warning system which allows the HFEA to consider the legal, ethical and regulatory implications of any techniques that researchers or clinicians may wish to use in the future in HFEA-licensed research or treatment. The Authority can then be prepared with information to make a decision on the potential licensing of techniques, or have guidance in place to ensure that new treatments are carried out safely and appropriately. The HFEA can also ensure that patients and the wider public are suitably informed."

Investing in Tomorrow's Talent in [all education sectors]

Q7: What are the future challenges for STEM education at primary, secondary and tertiary levels? How should these challenges be addressed?

Short answer [Schools]: AI has a special role to play in efforts to show primary and secondary students that Computing is an exciting, deep research area, not just a matter of developing (or, worse, merely using) bread-and-butter IT products, and has fascinating connections to other disciplines such as Psychology, Brain Science and Evolutionary Biology.[10]

Short answer [Universities]: A major challenge at the tertiary level is cross-disciplinary courses, both amongst different STEM subjects and between those and others. Additional care is needed, including on the part of scientific learned societies and professional bodies concerning the educational constraints they impose via course-accreditation policies.

Education in Schools

It is widely acknowledged that the teaching of IT and Computing in schools in the last few decades has failed to engage students. Students are given the impression that computers and computation are merely useful tools, to be used but not studied in depth. This is partly engendered by an emphasis on "skills" (worthy as these are in themselves). Recently the Council of Professors and Heads of Computing (CPHC) produced a report revealing worrying problems in the perception and/or ignorance of Computing at schools, particularly around up to the GCSE level and amongst students who do not go on to do AS or A level study in Computing, partly based on an extensive survey performed in 2006.[11] It is worth quoting segments of the report at length:

> [On p.1:] [There is] a more serious and potentially damaging issue over the ICT and Computing curricula at Key Stages 3 and 4, leading to GCSE examination. Feedback from school pupils and University students [references cited] indicate that enthusiasm for Computing as a subject is being stifled at this point in their studies, with the result that they find the subject "boring" and eliminate it from consideration for further study. This can be evidenced from the decline in pupils taking A-level Computing and applying to study Computing/Computer Science at University, as reported by CPHC [reference] and e-Skills UK [reference].
>
> ... Perhaps most worryingly, the evidence is that pupils are emerging from Primary School with enthusiasm and interest in ICT and Computing, both genders, and within three years the majority have lost that enthusiasm, find the subject boring, and only a relatively small percentage, predominantly male, will continue to study it at A level and on to degree level. Teachers need the learning resources, content and skills to provide pupils at Key Stages 3 and 4 with subjects that are exciting, motivating and really show the power and creative potential of the tools that are available to them, and it is clear that this is not currently happening.
>
> ... [A]nother major issue identified is the conflation of the ICT curriculum with Computing, with the effect that pupils see no difference between the two. The analogy is often made, in considering this issue, with the use, maintenance, design and development of motor cars [roughly, ICT as driving, A-Level Computing as design and construction].
>
> [p.2:] From the foregoing, it should be clear that the view held by the CPHC community is that, while there is some crossover and mutual support between the subjects, ICT and Computing should be seen and treated as separate subjects. Perhaps most importantly, this distinction should be made early in the curriculum, to ensure that pupils are aware of the distinction and don't believe that the use of application packages, such as Word, Excel and PowerPoint, represents the major focus of Computing.
>
> [p.4, quoted from a survey report:] "Many of the students who were taking neither CS nor ICT at AS/A-level stated that their negative experience of GCSE ICT was a primary reason for this. Common themes in criticisms of GCSE ICT included: ... The content is repetitive ... The content is uncreative or mundane ... ICT involves doing things 'that you can do already' ... its main function was to support studies in other areas."

[10]The connection to Evolutionary Biology is that one technique used in many successful AI applications is algorithms that automatically evolve solutions to problems and designs for robots and for software entities, using processes inspired by DNA recombination and mutation (Bäck, T., Hammel, U., Schwefel, H.–P. (1997), "Evolutionary computation: comments on the history and current state," *IEEE Trans. Evolutionary Computation 1*(1): 3–17).

[11]*A Consideration of the Issues relating to the Revision of the GCSE ICT and Computing Curricula*, CPHC, February 2009. Available at `http://www.cphc.ac.uk/docs/cphc-gcse-ict-curriculum.pdf`.

Also, the lack of awareness amongst school students of what Computing involves, and negative public perceptions created by government security lapses and large-IT failures, are discussed in an earlier British Computer Society report.[12]

Surely, given the targets for curiosity mentioned in answer to Q1 (nature of thought, etc.), together with the actual and potential practical outcomes and deep impact on society, AI is in a position to inspire school students at all levels. Further, it is in a specially favoured position here compared to areas of Computing whose deep and long-term significance is more indirect or requires considerable technical or scientific knowledge to appreciate. Therefore, in future advances on the Computing aspects of STEM (and other) education, it would be beneficial consciously to give AI a major role. This does not exclude other areas of Computing being given such a role—another candidate could be computer security and electronic voting, especially as the latter has important consequences for matters such as social inclusivity and the underpinning of democracy in insecure or corrupt environments.

The cross-disciplinary ramifications of AI into areas such as Psychology, Linguistics, Philosophy, Literature and Art should also be stressed. AI could be an integral part of a broad, humanities/social-science rich curriculum. It would need to be stressed that AI can be an exciting subject even for students who do not have a technical/scientific bent.[13]

However, the changes implied by these comments would require considerably more flexibility in school curriculum design. Because of the high level of diversity of topic, technique and disciplinary background within AI, making it unreasonable to expect teachers to be uniform in their grasp of the subject, national curricula and policies would need to allow considerable divergence between schools as to the portions of AI covered, the time spent on AI, and the impact on the teaching of other subjects.

We should stress that we do not see schoolteachers as primarily to blame for deficiencies in Computing education at schools. The essence of the problem lies within curricula and misunderstanding by policy makers and the public at large.

Education in Universities

Funding pressures, budgetary competition between disciplines, course-definition bureaucracy and excessive regard to the philosophy of modular delivery and credit-awarding can make it hard at universities to sustain minority courses. This is especially so for multidisciplinary courses, which tend to have complex multi-department timetabling problems, to upset a department's module prerequisite structure and to push beyond normal credit-total limits. These issues are familiar to most people who have been involved in the process of developing such courses.

However, an additional serious problem whose origin is external to universities as such is the constraints on curricula and employment imposed by accreditation philosophies of bodies associated with particular disciplines. While these philosophies have good intentions, they can impose constraints that make it difficult to define multidisciplinary courses that either satisfy the accreditation requirements or that, if they do not, remain attractive in practice to students in view of the deleterious consequences for employment, whether those consequences are real or perceived. We give one example of the problem, though in anonymized form to avoid pointing the finger at individual organizations, something that is not our aim.[14] A thriving half-and-half AI-and-X course run by a Computing department and a department of X (a science) at a particular university collapsed in the early 2000s because it no longer became possible to acquire accreditation for it from what we will call here the BXS (the British combined learned society for X and professional body for the practice of X). This was because of the amount of material that did not formally count as X. This meant that the X department was no longer willing to put resources into the course, and students were not attracted to the course because of X-related employment opportunities requiring BXS-accredited degrees, or simply because the course lacked any formal accreditation. This is despite the fact that it would in fact have equipped students well for other types of employment. We can certainly see the need for stringent accreditational links between education and some forms of employment. But it appears that there is a general problem amongst accrediting bodies in confining the notion of accrediting a degree for the purpose of its being a qualification in a particular line of professional work rather than accrediting a degree because it is a well-designed, valuable educational experience—an experience whose particular professional usefulness cannot be accurately foreseen but that is therefore all the better suited to novel types of career and to innovative emerging and future developments in society.

[12] McGettrick, A., Boyle, R., Ibbett, R., Lloyd, J., Lovegrove, G. & Mander, K. (2004), *Grand Challenges in Computing - Education*, The British Computer Society. Available at `http://www.bcs.org/upload/pdf/gceducation.pdf`. Additional commentary, including explanation of the Computing/ICT distinction, can be found at the website of the *Computing at School Working Group*, `http://www.computingatschool.org.uk`.

[13] More on this matters and others in this section can be found in A. Sloman, "Teaching AI and Philosophy at School?" (`http://www.cs.bham.ac.uk/~axs/courses/ai-syllabus.pdf`) to appear in the *Philosophy and Computers Newsletter* of the American Philosophical Association, Fall 2009. Sloman is a Fellow of the AISB.

[14] Details can be supplied confidentially on request.

Q8: How do we ensure that adequately qualified science and mathematics specialists are attracted into the teaching profession at all levels of education?

Short answer: Complementary to the Schools answer to Q7. AI is especially promising for grabbing the imagination of teachers.

However, to ensure that AI succeeds in being a usable imagination-grabbing tool, teachers need to be convinced that exciting teaching about AI can be done without a great deal of technical knowledge of Computing or even IT.

Again, it would also be useful to stress to potential teachers the cross-disciplinary ramifications of AI. In fact, AI could be one main avenue for delivering cross-disciplinary education and for conveying the interplay of sciences with each other and of sciences with the humanities.

The Ecology of Research Funding

Q16: What would an ideal research funding landscape look like in 20 years time? How would funding be allocated? What would the funding bodies look like? How would they relate to one another?

Short answer: Cross-disciplinarity and curiosity need to be much more systematically and comprehensively involved in funding policy and practice, and in the very philosophy of what universities and government research funding are for.

Our comments here are inspired in part by the extremely cross-disciplinary and curiosity-driven nature of AI, as outlined above. However, the comments are not specific to the case of AI.

There are several dimensions as regards the nature of possible funding schemes, notably the following: the time-period of the grant; whether it is in responsive mode or targeted in some way; whether it is "project-level" — given to an individual researcher or small team — or "organization-level" — to a department, university or consortia of universities and (possibly) other bodies; whether it is for a specific, agreed research topic or for whatever the recipients wish to investigate as time goes on; and the extent to which it covers some of the time of academics on existing open-ended contracts versus supporting specially-acquired research assistants or students. These dimensions are in principle largely independent of each other—a point that itself needs emphasizing—although practicalities dictate that some regions in the multidimensional space are more justifiable, practical or effective than others; and a given region is likely to be advantageous from some points of view and disadvantageous from others.

In the following we mainly address the project-level/organization-level dimension and the question of lead investigators' own time versus that of assistants and students, with less attention to distinctions along the other dimensions.

Ordinary individual research projects (with a small number of investigators over a small number of years), especially in responsive mode where there is no push to thematic cohesion between different grants, are usually at too small a grain and too small a scale to be genuinely assessable on likelihood of real, long-lasting societal/economic benefit, or for it to be reasonable to expect them to have it. Rather, it is universities as wholes, not individual projects, that might be expected in the long term to produce a decent amount of economic/societal benefit and to be able to put forward a well-evidenced, plausible case that their research programmes will provide such benefit. Thus, *to the extent that* government funding for research is tied to having some assurance of economic/societal benefit in the foreseeable future, it is logical to focus on university-level funding, whether through QR-style funding, grants under targeted programmes such as that for Digital Economy Hubs, or other mechanisms. The roughly-QR-style case of general support for staff research is an important element in the mix here as it allows for innovative research beyond what the government has already foreseen, and has a regular, ongoing quality that is important in view of the fact that innovative research ideas can be inspired at any time and can need long periods of mental incubation.

But, in a similar way, QR-style funding is also important for curiosity-driven research. It is much easier for a university as a whole to make a case that its purely curiosity-driven research is something the country should fund than it is for an individual grant proposal to justify the worth of the specific curiosity-driven research it proposes. A university can point to examples over many years of where its curiosity-driven research has led in the past. Again, the point is one of grain and scale: it is a curiosity driven research as a whole, or at least in large chunks, that can best be argued to be in the nation's interest.

Furthermore, by its very nature curiosity-driven research is risky, and especially in its initial stages is best served by ensuring that academic staff in universities are given substantial blocks of time in which to do their research as part of their paid contract, without having to secure funds from outside to cover that time. Securing these funds involves making a case for worthiness that may in principle not be able to be made and that is dependent on the lottery of the particular academics elsewhere chosen as reviewers having the right sort of vision to approve the work. It is up to the university to attract and nurture excellent staff who can be expected to do excellent research. Some of that research can be, or can become, linked to programmes in the university that are aimed at particular types of expected societal impact. Some of it can be, or can become, riskier and more open-ended. And the balance here in the case of any specific academic is not entirely predictable on appointment. Indeed if it were predictable one might question the wisdom of the appointment.

Better built-in funding for research at staff-initiated research at Universities would remove current tremendous pressures on early-career academics in research universities to spend precious time on chasing scarce funding, time that should be devoted to expanding their research horizons (not to mention becoming excellent teachers, administrators, and outreach agents). One important time-consuming activity here is studying related disciplines.

Cross-disciplinarity strengthens all these points. The panoply of possible mixtures of disciplines in cross-disciplinary projects (in AI or elsewhere) is too great for it to be reasonable to expect national bodies to service it properly and fully. It is universities, not government, that are — in principle at least — in the best position to set up (internally or by inter-university links) flexible, context-sensitive, imaginative arrangements for cross-disciplinarity, where these arrangements may be unpredictable in structure and may go beyond the scope of mechanisms that a national body has thought to set up. This is not to negate the important role that grants from national bodies could play. For instance, apart from some of the cross-disciplinary initiatives that RCUK has at times commendably set up, we are attracted by the idea of *Structured Interdisciplinary Research Programme* grants from the government, as proposed in the (independent) response to this consultation by John Fox (Professor of Engineering Science, University of Oxford, and Fellow of the AISB).

We believe these considerations suggest is the great importance of, roughly speaking at least, enhanced QR-style funding: organization-level, per-staff-member funding enabling universities standardly to have their academic staff devote large portions of their working time (over each year) to research devised on their own initiative. (Of course, there is a need for suitable monitoring practices to ensure diligence.) Universities should not be dependent on external, additional funding to support this mode of operation. Such research activity by its staff is *a large part of its central mission and should be paid for as part of the staff's normal salary.* We are therefore calling for more of a recognition that it is the *nationally crucial job* of universities to undertake research, and to include risky, imaginative research alongside research of more assured, more obvious and prompter societal impact. Finally, the needed recognition is not just on the part of government and other national policy leaders but also by university administrators and academics themselves.

Of course, to the extent that staff need assistants or other additional resources for their research, there remains a role as at present for applying to research councils, charities, companies, etc. for the necessary funds, whether in targeted programmes or under responsive mode, and whether the research is strongly curiosity-driven or not.

One upshot is that, against the current trend towards full economic costing of the lead investigators' time, this time should in many cases already be provided by better core funding of universities. This does not of course itself prevent academics from applying for external funds to buy additional research time.

Another corollary is that in responsive-mode project-level government funding there should be little if any pressure to demonstrate likely economic or societal impact. Such impact would be more the responsibility of (some) targeted programmes, (some) university-level funding, and third-sector funding. The likely impact that should be required to be argued is impact within the scientific or other disciplines involved in the research.

We are aware that the ideas expressed here do not sit well with the current economic circumstances of the country and in particular the current threat of spending cuts in HE, but we rely on this consultation truly being about the next 20 years not the next two. In this connection we should comment on the views reportedly expressed by David Lammy, HE Minister, at the Universities UK Annual Members Conference in September 2009.[15] We can certainly support Lammy's call for more UK universities to emphasize their economic and social contribution in their publicity material. However, his statement that [as reported] "Research funding will depend increasingly on demonstrating the potential impact [of an economic or social kind, presumably] of academic work to panels of businesspeople and policymakers" is highly questionable if the envisaged panels are not just for restricted types of funding but extend more generally. Why business people and policymakers? There is no evidence that businesspeople and policymakers are in more of a position to judge likely economic/social impact than academics are, in long-term cases at least. Indeed, many academics have authoritative knowledge about economic and social impact of advances in their field. Moreover, the independent advice of academic historians of commerce, society and science would be important. In any case businesspeople and (politically biased)

[15] "Image is everything, Lammy tells v-cs," *Times Higher Education*, 17 September 2009, pp.10–11.

policymakers have their own vested interests that it would be difficult to believe they would simply put aside, especially when it comes to long-term impact (cf. the comments above on short-term cycles in the subsection on Question 1).

Finally, whether the above ideas are followed or not, the better support of curiosity-led and/or cross-disciplinary research may require changes in culture and practice within universities themselves, both amongst the administrative bodies and amongst the academics. In particular, while deep changes in funding ecology may be called for, these will probably require deep changes in funding digestion—how universities manage the use of whatever funds are acquired. Although the issues are too large to go into here, there are consequences for university structure, budget-apportionment practices, other aspects of internal governance, and promotion practices. We would point in particular to the dangers of having rigidly organized research groups, a large role for individual funding success in academic promotion, and lack of imaginative management provision for bottom-up creation of cross-disciplinary initiatives.