Futures 43 (2011) 986–995

Contents lists available at ScienceDirect

Futures

journal homepage: www.elsevier.com/locate/futures

# From "Broad Studies" to Internet-based "Expert Knowledge Aggregation". Notes on the methodology and technology of knowledge integration☆

Vlad Tarko [a,*], Paul Dragos Aligica [b]

[a] Center for Institutional Analysis and Development, Sp. Independentei 17, Sector 5, Bucharest, Romania
[b] George Mason University, Mercatus Center, 3301N, Fair Dr., Ste. 450, Arlington, VA 22201, USA

ARTICLE INFO

*Article history:*
Available online 22 June 2011

ABSTRACT

The article takes as a starting point Herman Kahn's concept of a process-based approach to future-oriented knowledge aggregation and integration and discusses the opportunities created by Internet technologies to achieve the goal envisioned by such a model. The article briefly introduces a general framework for analyzing knowledge aggregation, the use of various aggregation mechanisms as well as aggregation systems that rely either on human facilitators (meta-experts) or on computer algorithms. Available online instruments for knowledge aggregation and expert challenge are presented and the possibility of using them for setting up "virtual think tanks" for foresight studies is explored.



There are many debatable issues in futures studies but one issue is considered commonsensical: There is no single method of thinking and approaching alternative futures. We need approaches that encompass different and divergent perspectives. The obviously complex and ambiguous problems of dealing cognitively with the future require an inter- and trans-disciplinary approach. As Ziauddin Sardar and George Cairns, the organizers of this year's Bucharest Dialogues, put it, this commonsensical view raises a natural question: can different futures and foresight methodologies be integrated to provide trans-disciplinary and culturally inclusive perspectives on alternative futures? This complex question raises in turn a series of other challenges, such as: what is the best way to integrate different methods in order to provide multiple, multi-dimensional, yet internally consistent perspectives? Is it really possible to integrate different methods? Can we develop a simple framework for integration of futures methods?

The importance of all these questions could be recognized in the fact that they were addressed from the very beginning by the founders of the field. Our article takes as a starting point precisely the work of one of these foundational contributors: Herman Kahn. We'll address the notion of board studies introduced by him as a response to some of the above mentioned challenges but also his more important idea of knowledge integration through "institutionalization of interdisciplinarity". Then, we'll move further to address the new opportunities presented by Internet, social networking technologies and artificial intelligence systems. We address these issues in the context of a rather general theoretical framework of knowledge aggregation. We conclude by pointing out the persisting problems, distinguishing between intrinsically difficult (or almost insurmountable) difficulties and the difficulties that exist merely due to technical reasons (due to the fact that truly powerful knowledge aggregation technologies are still in an incipient stage).

☆ This article was prepared in the context of the Bucharest Dialogues Workshop, organized in Bucharest, June 9–11, 2010, by Ziauddin Sardar and George Cairns. The authors would like to thank the organizers and participants.

* Corresponding author.
*E-mail address:* vlad.tarko@cadi.ro (V. Tarko).

0016-3287/$ – see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.futures.2011.06.008

The argument is based on the observation that there are two ways of dealing with the question of what is the best way to integrate different methods. The first is conceptual or meta-theoretical: To find a meta framework that could bring the methods together while creating a basis for their epistemological and logical consistency. The second is institutional and processual: to create an institutional process that brings together the dispersed knowledge that could be captured or mobilized by various methodologies, and then to organize and focus it as a function of the cognitive and forecasting task at hand. We are talking, in this second case, of a social structure and social process for what may be called a "social technology". We'll focus our attention on the second case. Interdisciplinary knowledge processes (from aggregation to verification and validation) are essential for proper foresight and this is the context in which Kahn's idea of a process of institutionalization of interdisciplinarity becomes very relevant. We'll take that as a starting point and we'll continue with current developments. Today, the development of web-based techniques have changed profoundly the nature of the game and are injecting a new life into Kahn's original ideas. "Wikinomics" [1], the "wisdom of the crowds" [2], the "army of Davids" [3], or "collective intelligence" [4,5] are names for a phenomenon that has at its core a process-based approach to knowledge and method aggregation.

## 1. "Broad Studies"

Kahn was stroked by the difference between the academic approach to a social or policy problem and the decision maker's approach (for an elaborated version of the theme of this section see [6]). The complexity and urgency of the phenomena the decision makers have to face precludes the luxury of assessing things from a unique perspective using a unique conceptual framework. Decision makers cannot afford to slice social reality on disciplinary lines or to focus on more or less arbitrarily segments or facets of the phenomena. Moreover, while researchers may eschew predictions and forecasts, the fact that the policies will affect and be affected by future developments in the distant future, makes forecasting attempts inevitable. Although predicting or even imagining the entire range of future possibilities is impracticable, the policy-makers cannot evade explicitly addressing the potential challenges of the future. Thus, decision makers, regardless of whether they like it or not, are forced by the very nature of their social function to look into the future at a "relevant range of possibilities, remembering the importance of examining possibilities which seem relatively unlikely but which would have very desirable – or catastrophic – consequences if they occurred". That happens because it is the purpose of policy "to plan for that which is more or less predictable" and to "hedge against that which is uncertain, both to be able to exploit favorable events and to guard against the consequences of unfavorable ones" ([7]: pp. 104–105).

Accordingly, decision making and policy processes pose two crucial problems that the scholar is rarely forced to face: (1) the need to have and to operate with a broad, multifaceted vision of phenomena and (2) the issue of coping imaginatively and realistically with future circumstances and conditions that can only vaguely be distinguished in the present. Kahn's approach is built on common sense: while it is impossible to fully predict the future, "it is possible to speculate usefully on many aspects of the future and even predict some". In a similar way, while it is impossible to have a complete specialized interdisciplinary personal knowledge of an issue, it is possible to assemble teams of experts and combine their knowledge in a coherent and useful manner ([8]: p.104). In other words, it is possible to have what he called "Broad Studies", a systematic mixture of conceptual and analytical approaches based on diverse technical expertise dealing with intricate and multifaceted issues from a large scale and future oriented perspective ([7]: pp. 106–7; 6, pp. 1–3). Nevertheless, this common sense simplicity is only apparent. The devil is in details and Kahn's merit is that he moved the discussion beyond the common sense level while avoiding falling into the traps of the technocratic spell or formal administrative solutions.

His treatment of the use of expertise was crucial in this respect. Kahn was a clear supporter of the employment of "experts" in dealing with forecasting problems and of the associated techniques, such as the Delphi procedure, the nominal group etc. However he also knew that although this might be a viable solution in some cases, when it came to broad and bold visions of the future, it encountered serious problems. The knowledge of empirical detail is hardly ever enough to take account of events on truly large scales. The expert group approaches work better for smaller scale predictions rather than for the large scale, aggregated, long term or other general trends, which Kahn considered to be the most important challenge to policy making in modern times. Moreover, the experts bring with them not only invaluable informational and analytic skills but also a selective professional distortion in the reception of data, the formulation of the problem, and the structuring of intuitions, i.e., they carry professional bias and parochialism ([7]: pp. 105–6).

Consequently, any attempt to deal with complex and large range issues is trapped between the danger either of being myopic and biased or of going beyond intensive specialization and risking superficiality. "Broad Studies", the fusion of interdisciplinary social science and futurology he advocated, was thus inhibited from its inception both by the risk of superficiality and by hyper-specialization. A comparison between "Broad Studies" and "narrow" or specialized studies puts their nature and the challenges facing each in a clearer light:

> "Almost all who have attempted such broad research agree that the danger of superficiality is great even when these generalizing or synthesizing efforts are carried out with high seriousness. But, just as detailed and specialized research is worth doing even if it often proves to be of narrow or sharply limited relevance, general treatments may also be worthwhile, even if they begin in superficialities. Extensive, 'shallow' studies, in order to be useful, must often meet higher standards than intensive and narrow research. The risk of total failure is much less in the narrow studies, since even a routine level of competence can usually guarantee some usefulness for the product." ([7]: p. 107).

Thus, Kahn was defining a domain of its own by contrasting it with the traditional domain of disciplinary specialization: on the one hand there were the "Broad Studies" and on the other hand, the "narrow" or specialized studies. Both had their own specific rationale and functions and both had their limits. In dealing with complex problems there was a trade off between "Broad Studies" approaches and specialized, expert approaches and although the "broad" dimension may not always be recognized, it was always present. If that was the case, then the question was why expert, specialized approaches were so overemphasized, while broad studies were so neglected. The answer to that, considered Kahn, was to be found in the intellectual climate of the epoch.

Along with the modern trend toward increasing intellectual specialization, came the academic research, especially in social sciences, that had less and less contact to the decision making process and thus it had slowly lost sight of long term and complex issues that represent the ongoing challenge of modern policy makers. The idea that solid research results can be reached only by strict specialization, dominated the epistemic climate of the century. In academic life and scholarship it was postulated that serious knowledge about any area was obtained only by highly specialized research. And also with the rise of the modern institutionalization of science there was almost "no scholarly market for integrated, overall work unless the author was a recognized 'authority' – or an elder statesman" and there was not even a non-scholarly market for such work, "except for authors with a facility for dramatizing and popularizing" ([7]: pp. 106–7).

Thus, concluded Kahn, there was relatively little motivation to seriously engage in broad and future oriented inquiries and, in these circumstances, these had unsurprisingly assumed a marginal position. By the time he started his work in this direction the difficulties the domain had to confront were huge and his work faced a very difficult uphill battle. But, as years went by, the demand for studies of such nature became more and more obvious. A literature on these lines started to grow as a living proof that they were needed. Consequently, Kahn saw his efforts [9] not only as a response to a complex and real social need related to the policy and decision making social functions, but also as a timely intervention in an intellectual climate ripe for change.

To sum up, Kahn's work in this respect could be seen as an integrated solution to a multilevel set of problems or challenges: (a) the need for broad systematic studies dealing with big problems defined in a future oriented horizon and of direct relevance for decision making, (b) in the conditions of a hyper-specialization trend in the social and policy sciences hostile to such efforts and (c) trying to undercut an increasing production of "vulgar" or popular literature that was meeting the public need for precisely such "broad" studies without meeting standards of accuracy and practicality.

## 2. Think tanks as the institutionalization of interdisciplinarity

Kahn's answer to the challenge was bold and innovative. Far from being of a mere intellectual or theoretical nature it was primarily methodological and institutional (for a more elaborated version of this argument see [6]). In fact, for Kahn the last two coincided in a very remarkable way. The starting point of his solution is the idea that indeed broad, future oriented studies are possible only as a collective effort, but a collective effort that needs a special organization and management. Thus far the idea seems pretty common: organizing groups of experts was in the air during those days. But Kahn's approach went beyond the Delphi method or other notions of how such groups should be organized and managed, that took shape by the same time he was working on the idea ([10]: pp. 134–57). Indeed, the specific mode of "organizing" the group was his particular insight. The Delphi group of experts is contingent, short term and temporary. It is organized for several rounds focusing on a specific topic or in area contingent to a particular topic and then disbars. Kahn was not envisioning a mere procedure, but a real social and organizational structure encapsulating on an enduring basis an interdisciplinary pool of talent and expertise in an institutional environment designed to amplify and stimulate the intellectual interaction between participants.

Kahn's vision of interdisciplinary collaboration went beyond the idea of putting together in an ordered way different pieces of analysis and conclusions coming from different disciplines. As "interdisciplinary workers" must almost necessarily rely on secondary sources, or on the advice of experts whom they have difficulty evaluating, Kahn thought that mere fixing together sets of data and conclusions is not going to work ([7]: pp. 117–18). Instead, he thought that this problem can be much alleviated by creating a process through which various approaches, ideas and insights are merged organically together, are played against each other and are assessed on an ongoing basis. In this respect probably it is safe to say that Kahn had in mind a sort of collective entity, something on the lines of a collective replica of a brain ("Think tank" comes from a British slang for head). This milieu was supposed to function like a social mechanism or process with each expert or participant having a loosely defined role to play in the entire arrangement designed to create a facilitating social, psychological and institutional environment. Once that overall environment had been set forth, ideas, theses, and generalizations could be produced, corrected, adjusted and tested and most of the superficiality and distortions haunting "broad" studies thus removed or at least alleviated. For Kahn it was just at this point that the interdisciplinary approach could yield important and relatively relevant results ([7]: p.107). Summing up, Kahn's solution to generate broad, syncretic studies that meet intellectual standards is primarily institutional. The institutional arrangement envisioned by him is supposed to function in two ways: first to pool information and expert knowledge and second to set into motion and maintain a milieu conductive to an ongoing process of creativity, discovery and rational criticism and to a psychological state of confidence in the process.

Once he reached that conclusion, Kahn did not limit himself to being a mere advocate, but he also followed its practical implications up to the end. Hudson Institute, Kahn's creation, has, since its formation in 1961, attempted to do exactly this kind systematic, interdisciplinary examination of topical problems that have a significant future-oriented dimension. Its

modus operandi has been precisely to synthesize and integrate very diverse specialties, skills and areas of knowledge, and to create a context for the cooperation of a varied group of researchers and specialists in a loosely defined but still organized manner. By the time it was founded, Hudson was unique in many respects and astonishing to many: a place whose key objective was to facilitate and challenge free and creative thinking by increasing the social contiguity between researchers, with loose bureaucratic procedures, with no intellectual routines or taboos, with almost no bureaucracy, with no constraining schedules, and with an entire gallery of eccentricities associated to its members, Kahn having the lead in this respect, Hudson created the public image and the legend of the archetypal think tank [11,12].

Thus, Herman Kahn's idea of a think tank was rather different from the currently dominant one. For him a think tank was not a copy of specialized academic departments. Academic departments are the institutional expression of "narrow" studies. On the other hand, think tanks are the institutional expression of "broad" and future oriented studies. Kahn's view of interdisciplinarity was not a mere endeavor of putting pre-crafted pieces of knowledge or expertise together. On the contrary, it was a vibrant and organic institution aiming at an extensive integration of expertise, where intuitions, background knowledge and critical thinking were weaved into large, multi-perspective, multidisciplinary inquiries. The image guiding Kahn was not one of isolated, silent cubicles and departments but of a wide open noisy conference hall and of continued open interactions for which the metaphor of a collective entity, something on the lines of an ongoing thinking collective entity.

Kahn's bold idea deserves to be seriously considered. It is obvious from the very beginning that its essence is in the emphasis on the institutional and processual aspect. Discussing the most common and immediate reactions and critiques that come to one's mind is a way to make a step further in its understanding. However, it is important not to confuse the ideal model imagined by Kahn with real life think tanks. In other words, to keep in mind that the organizations called "think tanks" nowadays are different from the organizational model and process described and advanced by Kahn. We can analyze and evaluate real life think tanks from the perspective of Kahn's ideal, but not the other way around. It is also important to note that Kahn did not manage to successfully promote his idea or to implement it. Moreover, Hudson Institute, his attempt to institutionalize interdisciplinarity, evolved in the direction of a standard think tank.

To explore the reasons why his idea did not take off is a project in itself. In the context of the current article it is sufficient to say that most of these reasons do not seem to be circumstantial. They are related to the limits of the human capacities to communicate, cooperate in producing, analyzing and testing knowledge and systems of ideas in a given technological environment. In other words, Kahn's ideal was not achievable due to the social and technological parameters known and available to him. The concept of "broad" knowledge was correct, but the question of how to achieve it, and of what the specific institutional process capable of delivering it are, is an entirely different issue.

However, things have changed since Kahn has first advanced his intuition. The pivotal difference is, indeed, the IT revolution. We have thus the chance to revisit Kahn's think tank idea in the Internet era.

Before providing more details about the specific ways in which knowledge aggregation systems function and about the technical opportunities provided by web-based approaches, let us mention a more general sociological aspect. This allows us to see more clearly the promise of web-based knowledge aggregation systems, in particular how they are supposed to help us overcome the aforementioned limits of Kahn's idea.

The Internet creates the possibility of expanding the circle of experts by means of virtual communities of experts. The larger the community, the better suited it is at protecting itself from biases. A virtual "think tank" may have a less clearly defined identity, but it may generate better conclusions precisely for this reason. The large online community may be permanent, but with shifting and constantly evolving members, stakeholders or interest groups. The virtual community may be seen as a pool of experts who enjoy to constantly get engaged in various tasks of practical value, but less due to being under a contractual obligation to do so, and more on a voluntary basis when they consider that they can efficiently contribute to the task. This could be even more in line with Kahn's original anti-bureaucratic conceptualization of the think tank.

One of the most appealing features of a "virtual think tank" (i.e., an Internet-based process), and which may increase the efficiency of the entire approach is that the organic development of ideas under thinner contractual obligations may diminish some of the rivalries between different intellectual paradigms. Experts may find specific "junction points" at which to insert their contributions from their own specialized perspective. It is often the case that some aspects of the social system receive only a relatively simplistic treatment from some perspective, while much more elaborated accounts are available. Nonetheless, the unified perspective is often difficult to acquire due to conflicting basic premises. A case by case operational synthesis or harmonization, especially by means of discussing empirically well defined examples, might be easier to attain and online systems may facilitate this.

## 3. Aggregation mechanisms and aggregation systems

Kahn's idea was to use the preexistent intellectual division of knowledge in order to create a system able to aggregate dispersed pieces of knowledge and facilitate a focus that better illuminates specific issues. The emphasis is on the notion of knowledge aggregation in circumstances in which expert knowledge is less effective for various reasons. Our main argument for this context and the pivotal argument of this article is that the idea of knowledge aggregation evolved since Kahn had made his point. As we have mentioned, the Internet creates novel opportunities for a better approach to Kahn's ideal. The emergence of Internet not only has stimulated thinking about knowledge aggregation, but has also created a variety of technological solutions to the technical problems posed by the processes in question.

**Table 1**
The basics of knowledge aggregation.

| Knowledge aggregation mechanisms: | Aggregation systems: |
|---|---|
| Consensus building | Human facilitators (meta-experts) |
| Voting (polling or averaging) | Computer systems (algorithms) |
| Betting (prediction markets) | |

Table 1 synthesizes in a very general manner the basics of knowledge aggregation (see also [5]). The several possible aggregation mechanisms (ways of combining the opinions of individual experts) can be incorporated in two different types of aggregation systems.

The 3 main mechanisms for knowledge aggregation are: consensus building, voting, and betting. The specific ways in which these ideas are used can differ widely, i.e., we can have very different aggregation systems built with these three basic mechanisms. As such, they are not institutional frameworks in themselves, but rather the building blocks to be utilized in various institutional designs. It is interesting to note, especially in relation to the conference's main subject, that polling (in order to be scientifically acceptable) depends on the existence of a more or less unified conceptual framework relevant to different methodologies. Consensus building systems may create such unified conceptual frameworks. On the other hand, prediction markets (which rely on betting) and related systems do not depend on the existence of a unified conceptual framework.

When thinking about institutional design and the organization of processes, especially if one intends to use web-based systems, one has to distinguish between institutional frameworks that rely on human facilitators to organize expert opinion (such as the Delphi method), and computational or automatic institutional frameworks that rely on an impersonal aggregation algorithm (such as MIT's Collective Intelligence system). The three basic mechanisms can be incorporated in either of those two types of systems. It is important to point out however that the level of personalization can differ in a gradual fashion: we can have anything from a moderator-based system (i.e., the human factor plays an important role in the aggregation process) to an artificial intelligence (AI) system with inputs from human experts. But before describing various aggregation systems, we have to briefly describe the aggregation mechanisms.

In Internet-based settings, the standard consensus building mechanism is the wiki. Unlike more traditional methods for consensus building, such as the Delphi method, wikis often lack a facilitator (although sometimes they may have moderators). This makes wikis more anarchic in nature, a characteristic that has both positive and negative consequences. There are two different types of consensus that one may aim for, which we can label "strong" and "weak" consensus. In the strong consensus case, the knowledge aggregation system aims for a consensus about a specific interpretation or prediction. The experts are supposed to debate their positions until a unified position emerges. This is very much in line with the standard approach in science, but, as mentioned before, it is an approach that may require too much time, a resource that forecasters usually lack. Consequently, the knowledge aggregation system may have a less ambitious but more realistic aim, namely building up a weak consensus – a consensus not about a specific prediction or interpretation, but about the general conceptual framework which best fits the situation, i.e., a consensus about the relevant parameters for describing a situation or for building possible scenarios. The key idea behind consensus building strategies is that the experts involved are allowed access to a lot more information than it is necessary for solving the problem [1]. This is the case because one does not know beforehand which particular piece of information will prove to be crucial. Thus, one must allow the consensus building process itself to weed out the irrelevant pieces of information and to bring to prominence what is important.

Voting (or polling) relies on the existence of weak consensus. When one builds a poll, the choice of options must be uncontroversial, for otherwise the result of the poll has little weight. In the case of the Delphi method, the entire process may end with a poll among the consulted experts, if the previous stages have not delivered a strong consensus and, for some reason, one cannot or does not want to prolong the process even longer. There are certain conditions[1] under which polling usually delivers better results than any particular expert's opinion [2]: opinion heterogeneity among the polled agents (the group must not be created in a fashion that creates an inherent bias in the group toward certain favored opinions), opinion independence (agents should not consult each other or mimic each other's positions based on, say, reputation, i.e., groupthink must be avoided), division of knowledge (agents should have opinions based on a wide range of sources of information, even competing or contradictory sources of information). The result of a vote can be utilized in various ways. This is an important question. For example one may consider the winning choice (i.e., the option with most votes), or may consider the average. To use Galton's famous ox example: do we predict the ox's weight to be the average of what everybody has said (as Galton did), or do we use the most commonly believed weight as the prediction? The best way of using the votes may differ from case to case.

Finally, prediction markets [13,14] do not rely even on weak consensus. Prediction markets set up a framework for betting on various scenarios or outcomes. The logic is straightforward. The system predicts the outcome on which participants have placed the largest total bet. Apparently, this system depends on weak consensus as it requires the bets to be framed in a certain fashion. But, as many have pointed out, this is not the case however, as different bets, sometimes about similar issues, can be framed in widely different manners. Interestingly, one could move in the other direction, from the most successful participants toward a particular conceptual framework, by assuming that the winning participants win because they have a

[1] These conditions apply not just to polling but also to other forms of knowledge aggregation such as prediction markets.

better conceptual understanding than the rest. Thus, prediction markets can be an institutional tool for testing various conceptual frameworks or foresight methodologies.

We now move in our brief overview from aggregation mechanisms to aggregation systems. In order to introduce the variety of aggregation systems, we need to refer to the issues faced by different methods, as different solutions are designed to overcome specific issues and challenges. Moreover, many of those systems have not been specifically designed for forecasting practices. Consequently, any discussion has to deal with 2 matters: (a) the issues those online systems were originally designed to overcome, and (b) the way in which the systems which have not been specifically designed for forecasting practices can be utilized, combined or changed to fit such a goal.

One of the most famous and efficient online system for knowledge aggregation is the wiki, used by Wikipedia and many other specialized encyclopedias. Online knowledge aggregators (such as wikis) work very well at pitting different points of view against each other. Sometimes, the system can be arranged as such that specific outside guidelines are imposed. However, it may often be useful to keep such guidelines to the minimum in order to gather a larger set of diverging points of view. This may be useful particularly to the task of ascertaining the intrinsically unpredictable areas as this issue requires the aggregation of expert opinion about many different aspects of social life. But there is one important issue that should be mentioned in the context of our discussion. The main problem regarding the use of a wiki for organizing the activity of a virtual think tank of many experts is that wikis are best suited for incorporating and organizing already known facts. They are not very useful for managing research activity or setting up challenges among experts. This is not, of course, a criticism of the wiki technology, as it is not designed with such a purpose in mind. A wiki may be a useful tool for aggregating the results of the virtual think tank's work, but other Internet technologies are required for the actual management of research.

Indeed, one of the problems, from the point of view of the "Broad Studies" ideal, with existing online solutions is that they rely more heavily on search rather than on challenge. That tends to defeat Kahn's original purpose for a think tank. Even in existing physical think tank settings, or in the general workings of the scientific community, an expert is often directly challenged by a colleague or another to solve or address certain issues, or is challenged to account for certain apparent consequences of her/his views. This is an essential part of the creative research activity that generates the broad and practical perspective on an issue. By contrast, most online systems rely more on the expert searching for the things that interest her/him (e.g., the arXiv.org database).

Nonetheless, some of the first online collaboration tools, the online discussion groups, are to this day successfully utilized as challenge systems.[2] An online discussion group is a community of members that agree to receive bulk messages from the other members. Many online help systems are organized as discussion groups: When a member has a particular problem, he sends a message to the group describing it; an answer is usually received from one or more other members, who know how to address the posed problem. It is often the case that multiple solutions to a particular problem thus become available. Moreover, as the answers are public, a searchable database of solutions to particular problems is thus gradually created. One of the main problems of such online challenge systems is that they can easily degenerate into spamming, i.e., members being overwhelmed by large amounts of messages from other members.

Various mechanisms have been invented to address the issue of spamming, involving moderators, voting systems, categories, and groups of pre-selected contacts, but one is left with the feeling that better solutions are still necessary. Discussion groups, forums and collective blogs rely on using categories and moderators. By contrast, other systems (e.g., Digg and Reddit) rely on voting and categories. Moderators can eliminate some of the irrelevant posted materials, before it reaches the other users (thus limiting spamming). Categories can be used by users by subscribing to only a few categories of interest. Voting can also be used to organize in a hierarchical fashion the material posted on a website, and to hide the material that reaches a certain threshold of negative votes.

An interesting alternative to categories is the use of groups of pre-selected contacts. At present, simple versions of such systems have been set up mostly for entertaining purposes. The list includes Facebook, Twitter, Google Buzz and others. The idea of relying on groups of pre-selected contacts works as follows: a user has a large number of contacts (which she/he may further organize into groups); each of these contacts posts occasional pieces of information and these pieces of information become visible to all its contacts. Consequently, the online system displays to each user a personalized feed of information coming only from the pre-selected contacts, which, presumably, has a larger probability of being of special interest to the user than universal feeds delivered by category-based or voting-based systems. A personalized feed can also include, apart from posts from pre-selected contacts, additional universal feeds delivered by other sites.

It is an uncontroversial observation that, in the context of our discussion, the potential of all the above approaches comes from the integration of their facets. To illustrate, we turn now to describing in more detail one of the most advanced online system available, which combines categories, moderators, voting and challenges, albeit lacking personalized feeds, namely the IdeaTorrent framework (as with the wiki technology, IdeaTorrent is also open-source.) This framework is successfully utilized by Ubuntu Brainstorm (http://www.brainstorm.ubuntu.com), the meeting place between Ubuntu Linux developers and end users. This is one of the most advanced online communities in terms of its success and of the complexity of the

[2] They are also utilized to other purposes, such as delivering information, but those other purposes are currently better served by other, newer systems. The discussion group still remains among the most utilized frameworks for delivering and answering to challenges (such as online help sites). Alternative, newer, specifically designed challenge systems include Yahoo Answers (http://www.answers.yahoo.com), Stack Overflow (http://www.stackoverflow.com) and others.

practical task it manages to achieve. The framework has two layers, a layer of issues or ideas, and a subordinate layer of solutions to the issues or proposed methods for implementing the ideas. All issues or ideas are placed in predefined categories, such that one can easily find what one is interested about. Moreover, and most importantly, the system has four successive development stages:

1) an ''idea sandbox'' where ideas are originally posted by members of the Ubuntu community;
2) a ''popular ideas'' discussion group, to which moderators promote ideas out of the ''sandbox'', and where various solutions are proposed, discussed (commented upon) and voted; the voting system automatically organizes in a hierarchy the proposed ideas and solutions based on their popularity, i.e., perceived importance;
3) an ''ideas in development'' stage, that concerns popular solutions which have been picked by one or more developers willing to take upon themselves the task of implementing them into the Ubuntu operating system;
4) and a final stage of ''implemented idea'' that documents the successful implementations into the Ubuntu system of various solutions.

It looks like that, at least at present, this IdeaTorrent framework is one of the best suited instruments for organizing a virtual forecasting think tank, perhaps with a wiki as the final stage for showcasing the results. The forecasting activity has the same two-layer structure that is replicated by the IdeaTorrent framework: first, one has a theme or subject of reflection and investigation, or an issue in regard to which a forecast is requested; secondly, and in a subordinate fashion to the first layer, there may be different approaches to a proposed theme, or different forecasts in regard to the same issue. The IdeaTorrent 2 layers, 4 development stages framework would allow a forecasting community to organize itself and to generate (as a result of moderation and voting) a hierarchy of issues, of preferred approaches to those issue, as well as of forecasts. Stage 1 would involve the mere proposal of various subjects or of various forecasts requests. Stage 2 would involve debating, submitting of forecasts and of solutions to issues, and voting different proposals promoted by knowledgeable moderators from stage 1. Stage 3 would involve the chronicle of actual actions taken to implement popular solutions or to mitigate/address popular forecasts. Finally, stage 4 would record the implemented solutions, noting difficulties and effects. It is also worth noting that one could also imagine a prediction market added to stage 2 as an additional knowledge aggregation mechanism, alongside the voting system, thus allowing users not just to vote on what forecasts they find most credible, but to actually bet on them.

It is important to keep in mind that the success of this type of approaches would ultimately depend, indeed, on the size of the community willing to get involved in such large scale collaboration. The actual technology capable of facilitating this online collaboration, that would approximate in process and structure Kahn's vision of a collective brain at work, is already here to a significant extent. Moreover, it is available for free and open-source. Yet, in the end the importance of incentives should not be neglected. For instance, the willingness may be lacking as forecasters and experts often perceive themselves as competitors, naturally reluctant to share information that may give them their competitive edge.

All the above aggregation systems (wikis, discussion groups, idea torrents, personalized feeds) fall in the category of personalized systems, i.e., aggregation systems that rely directly on human decisions. By contrast, on the other hand, are AI aggregation systems. For instance the MIT's Collective Intelligence system is a step in this direction. Traditionally, AI systems have aimed at replacing humans altogether, i.e., they take objective measurements as inputs and produce predictions by algorithmic means, without any human involvement. An AI knowledge aggregation system would incorporate human experts. Rather than using objective measurements as inputs it uses subjective expert opinions as inputs. Nonetheless, it would remain an AI system because these subjective inputs are analyzed by an algorithm which produces a prediction.

The question is: how far can we go in an attempt to de-personalize a system? Is it possible to create an ''objective'' aggregation system? How far could we go to eliminate human decision (and the implicit subjectivity and error that come with it)? Is it desirable? Is it feasible?

Although the mirage of pure algorithmic and objective systems persists, the evidence runs against it. The idea does not hold, neither theoretically and philosophically nor practically (For an attempt to demonstrate the point see Annex). To put it in a nutshell, there is no such thing as an agnostic, assumption free knowledge aggregation system. Of course, in case of a highly complex system, such as the IdeaTorrent framework described above, it is virtually impossible to make explicit all the numerous hypotheses it incorporates, but that does not mean that those hidden assumptions and hypotheses do not matter. This observation has a bearing on the conference's theme: ''How do we integrate different foresight methods? Can we develop a simple framework for integration of futures methods? Or, do we need yet another method to do this? What impact would integrated methodologies have on the field?'' The answer is that having a single, totally neutral framework is probably a mirage too. Such a unified framework, to be worthy of its function, should be agnostic about the actual relationships between the parameters of interest. And as far as we could tell, such an agnostic framework is theoretically impossible. The need to preserve a variety of methods is thus a fact not an option.

It does not mean that if pure systems or the ideal is not possible, we should not try to explore the possibility of having knowledge aggregation systems that are as agnostic as it is made possible by circumstances. Once we understand the structural limits of our enterprise, we could try to build our foresight approach on devising the most ''neutral'' knowledge aggregation system possible. To end with one example, the MIT's Collective Intelligence is a demonstration that such an endeavor is worthwhile as long as it is pursued in ways that are sophisticated enough to avoid the mirage of epistemic perfection. The MIT's Collective Intelligence system uses two other approaches to incorporate human experts and AI

systems: the market and Bayesian Truth Serum [15]. The free market price system has long been understood as a knowledge aggregation mechanism, a kind of social computer that aggregates the widely distributed and complex local knowledge of various agents into a set of numbers, the prices, which then act as signals [16,17]. The price system is particularly well fitted to incorporate prediction markets. By having an ecosystem of prediction markets about various issues and of challenges made from various conceptual perspectives, the MIT system has the opportunity of setting up a price system that incorporates them all and thus gives signals about which approaches are most worthy of consideration and which conceptual frameworks seem most promising. The MIT Collective Intelligence approach also assumes that experts will devise various AI algorithms to compete on such a market. In the same way as in the case of Axelrod's classic challenge [18] experts competed by submitting various algorithms, forecasting experts could compete in an ecosystem of prediction markets both by submitting their forecasts and by setting up AI forecasting systems. All in all, the MIT system is an example of a constructive and promising approach.

## 4. Conclusions

It is clear that although online solutions may counterbalance some of the problems of Kahn's original think tank idea, they are not without limits themselves. Some of them are structural and profound. Some of them are more mundane, but nonetheless important. For instance, online communication is mostly text-based. People manage to exchange more information per unit of time in spoken conversation. However, recorded spoken conversations are often not very efficient information conveyors for third parties. The technology in (more or less) near future may address these issues by integrating spoken conversations among the basic features of an online community technical design, and by incorporating automatic speech-to-text conversion of those conversations for facilitating the dissemination of information to third parties, as well as for facilitating searching and archiving under specific categories. At this point in time, however, the technical framework of online communities is far from optimal.

But, besides such more or less technical issues, the main problem is, as we have mentioned, the impossibility of devising fully agnostic knowledge aggregation systems and thus also an entirely neutral or objective meta-methodology.

We need thus approaches that encompass different and divergent perspectives. One may try to integrate different futures and foresight methodologies to provide trans-disciplinary and culturally inclusive perspectives on alternative futures but that integration can only take place processually and institutionally. The challenge is to create institutional processes that bring together the dispersed knowledge on various methodologies and organizes and focuses it in function of the cognitive and forecasting task at hand. We should be skeptical regarding the conceptual (meta-theoretical) approach and focus instead on the institutional and processual. Any meta-theoretical approach has to rely sooner or later on a set of assumptions (implicit or explicit), which will be preventing it from being truly neutral and all-encompassing. Consequently, one can try to be more objective by setting up a competition between methods and by relying on a multiplicity of conceptual frameworks.

However, we should be aware that this distinction between an encompassing conceptual meta-framework on one hand and the rules of a knowledge aggregation process on the other hand can be doubted. And the fact is that there are serious theoretical limits even to the processual-institutional approach. Although it fares better than the meta-theoretical one, it is far from perfect. For instance, the Bayesian and market-based approaches we have mentioned have their clear limits. The Bayesian method, although it seems quite well founded theoretically [19,20], it is often marred by severe technical and computational difficulties. The market approach, although it is indeed a very effective method of aggregating local knowledge [16,17,21], should not be overly idealized as it is affected by random path-dependency biases [22,23] and it cannot thus deliver truly objective results.

This being said, what we have labeled the processual-institutional approach can, presumably, provide the basis for better foresight approaches. Moreover, depending on experience, the institutional framework itself can be gradually improved in order to facilitate better outcomes or to deliver them faster. In this regard, online systems, such as the existing IdeaTorrent framework that we have introduced, can prove to be very valuable tools. In the end, they could allow a global community of forecasters to interact on a daily basis, to challenge each other, and to build a growing database of methods, solutions, forecasts and social dilemmas.

## Appendix A. Annex

Let us describe the simplest possible form an AI aggregation system can take, namely a neural network of experts (NNE). This is an extreme form of an AI knowledge aggregation system, and understanding it gives us a better perspective on more complicated intermediary quasi-impersonal systems, such as MIT's Collective Intelligence. More importantly, this knowledge aggregation system allows us to see in a principled way the fundamental reason why knowledge aggregation is such a difficult problem.

Suppose we have $N$ experts and $M$ parameters of interest. Each expert $i$ has opinion $u_{ij}$ about the value of parameter $j$. In the simplest, linear case and assuming that experts' opinions are independent, the predicted value $v_k$ of parameter $k$ is:

$$v_k = \sum_{j=1}^{M}\sum_{i=1}^{N} W_{ijk}u_{ij}, \qquad k = \overline{1,M}$$

where $W_{ijk}$ is an array of numbers that includes the information about the attributed weight of expert $i$ about parameter $j$ (our confidence about this expert's opinion in regard to each parameter) and the information about the weight of parameter $j$ in determining the value of parameter $k$. This sets up a standard neural net which can be "trained" in order to determine the values of $W_{ijk}$, prior to using the NNE for making predictions. The only difference from standard uses of neural nets is that the input values are not objective measurements, but experts' subjective opinions. This is basically an algorithmic system that makes predictions based on ad hominem reasoning (as it depends on assessments of experts credibility about various issues).

We can further use the parameters $W_{ijk}$ to set up economic incentives within the expert network, e.g., each expert is paid a sum $p_i$ proportional to the sum of all its weights within the network's predictions:

$$p_i \propto \sum_{j=1}^{M}\sum_{k=1}^{M} W_{ijk}$$

This rewards both the expert's reliability in providing good estimations and the usefulness of her or his expertise. For example an expert's opinion may be highly reliable about a particular parameter but this parameter may be relatively unimportant in determining the value of the other parameters; in such a situation, the payment $p_i$ will be relatively small. An expert would receive a large payment either if she/he provides relatively good estimations for a wide variety of parameters (strategy A), or if she/he provides highly reliable estimations for a small number of important parameters (strategy B).

The NNE functions in the following way: The NNE receives challenges (i.e., inquiries about the values of various parameters) and, based on its previously determined weights $W_{ijk}$, and on the experts' expressed opinions $u_{ij}$, it produces a prediction $v_k$. When the prediction is eventually tested, the weights are changed based on the new information. The experts in the NNE receive the information about how their weights in the NNE have changed. This change affects their payment, but also, perhaps more importantly, it gives them specific hints about how to improve their expertise. For example, an expert might realize that its weights are relatively good across the board, in which case she/he might opt for strategy A and delve more into interdisciplinary research. Some other expert might realize that its expertise is concerned with mostly irrelevant aspects, in which case she/he might opt for switching to a (perhaps related) more important area (focusing its strategy B in more fruitful directions).

As mentioned, this example of an AI knowledge aggregation system has more of a didactic value thanks to its simplicity. It is unlikely that a NNE would be a truly efficient forecasting tool for the simple reason that most social problems are non-linear (involve complex feed-back relations between parameters). Consequently, it is unlikely that experts' expertise could be reliably assessed and guided by the linear weights $W_{ijk}$. But this helps us highlight the exact nature of the difficulty of devising a good knowledge aggregation system: We like it or not, the knowledge aggregation system incorporates a hypothesis about the form of relationships between the parameters of interest. The NNE, for instance, assumes that all the relationships between parameters are linear. It seems that we need to know how the world is in order to know how to explore it. This philosophical "paradox" is well known of course, and it is simply reformulated here in terms of knowledge aggregation, rather than knowledge discovery. This "paradox" was the reason why Kant assumed the existence of a priori synthetic statements, a highly controversial claim to say the least.

## References


[1] D. Tapscott, A. Williams, Wikinomics: how mass collaboration changes everything (2006) (Portfolio).
[2] J. Surowiecki, The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations, Doubleday, 2004.
[3] G. Reynolds, An Army of Davids: How Markets and Technology Empower Ordinary People to Beat Big Media, Big Government, and Other Goliaths, Thomas Nelson, 2006.
[4] L. Iandoli, M. Klein, Can we exploit collective intelligence for collaborative deliberation? The case of the climate change collaboratorium, CCI Working Paper 2008-002, MIT Sloan School of Management Working Paper 4675-08, 2008.
[5] T.W. Malone, R. Laubacher, C. Dellarocas, Harnessing crowds: mapping the genome of collective intelligence, CCI Working Paper 2009-011 (2009).
[6] P.D. Aligica, The challenge of the future and the institutionalization of interdisciplinarity: notes on Herman Kahn's legacy, Futures 36 (1) (2004) 67–83.
[7] H. Kahn, The alternative world futures approach, in: Tugwell Franklin Search for Alternatives: Public Policy and the Study of the Future, Witrop Publishers, Cambridge, MA, 1973.
[8] R. Kostelanetz, Master Minds: Portraits of Contemporary American Artists and Intellectuals, Macmillan, 1969.
[9] H. Kahn, A.J. Wiener, The Year 2,000: A Framework for Speculation on the Next 33 Years, Macmillan, 1967.
[10] O. Helmer, Looking Forward, Sage Publications, 1983.
[11] J.A. Smith, The Idea Brokers: Think Tanks and the Rise of the New Policy Elite, Free Press, 1991.
[12] D. Stone, Time line of British and American think-tanks, in: Diane Stone (Ed.), Capturing the Political Imagination: Think Tanks and the Policy Process, Frank Cass, 1996.
[13] R. Hanson, Idea futures: encouraging an honest consensus, Extropy 3 (2) (1992) 7–17.
[14] J. Wolfers, E. Zitzewitz, Prediction Markets, J. Econ. Perspect. 18 (2) (2004) 107–126.
[15] D. Prelac, A bayesian truth serum for subjective data, Science 306 (5695) (2004) 462–466.
[16] F. Hayek, The use of knowledge in society, Am. Econ. Rev. 25 (4) (1945) 519–530.
[17] G.P. O'Driscoll, Economics as a Coordination Problem: The Contributions of Friedrich A. Hayek, Sheed Andrews and McMeel, 1977.
[18] R. Axelrod, The Evolution of Cooperation, Basic Books, 1984.
[19] E. Jaynes, How does the brain do plausible reasoning? in: G.J. Erickson, C.R. Smith (Eds.), Maximum-Entropy and Bayesian Methods in Science and Engineering, Kluwer, Dordrecht, 1988.

[20] E. Jaynes, Probability Theory: The Logic of Science, Cambridge University Press, 2003.
[21] R. Axtell, Economics as distributed computation, in: T. Terano, H. Deguchi, K. Takadama (Eds.), Meeting the Challenge of Social Problems via Agent-Based Simulation, Springer, 2003.
[22] R. Axtell, J. Epstein, Growing Artificial Societies: Social Science from the Bottom Up, The MIT Press, 1996.
[23] M.J. Salganik, P.S. Dodds, D.J. Watts, Experimental study of inequality and unpredictability in an artificial cultural market, Science 311 (5762) (2006) 854–856.