AOSIS

The probabilistic innovation theoretical framework

CrossMark
click for updates

Author:
Chris W. Callaghan[1]

Affiliation:
[1]Department of Management and Human Resources Management, School of Economic and Business Sciences, University of the Witwatersrand, South Africa

Corresponding author:
Chris Callaghan,
chris.callaghan@wits.ac.za

Background: Despite technological advances that offer new opportunities for solving societal problems in real time, knowledge management theory development has largely not kept pace with these developments. This article seeks to offer useful insights into how more effective theory development in this area could be enabled.

Aim: This article suggests different streams of literature for inclusion into a theoretical framework for an emerging stream of research, termed 'probabilistic innovation', which seeks to develop a system of real-time research capability. The objective of this research is therefore to provide a synthesis of a range of diverse literatures, and to provide useful insights into how research enabled by crowdsourced research and development can potentially be used to address serious knowledge problems in real time.

Setting: This research suggests that knowledge management theory can provide an anchor for a new stream of research contributing to the development of real-time knowledge problem solving.

Methods: This conceptual article seeks to re-conceptualise the problem of real-time research and locate this knowledge problem in relation to a host of rapidly developing streams of literature. In doing so, a novel perspective of societal problem-solving is enabled.

Results: An analysis of theory and literature suggests that certain rapidly developing streams of literature might more effectively contribute to societally important real-time research problem solving if these steams are united under a theoretical framework with this goal as its explicit focus.

Conclusion: Although the goal of real-time research is as yet not attainable, research that contributes to its attainment may ultimately make an important contribution to society.

Introduction

In a global context experiencing exponential growth in stocks of knowledge and information, the 'breakneck pace of genome-technology development' has 'revolutionised bioscience research', and the scientific fields that underlie medical research and development (R&D) are advancing as never before (Hayden 2014:294). Notwithstanding these developments, Kaitin (2010) stresses the lack of progress associated with contemporary pharmaceutical innovation:

> Forged in the early 1960s, the paradigm for pharmaceutical innovation has remained virtually unchanged for nearly 50 years. During a period when most other research-based industries have made frequent and often sweeping modifications to their R&D processes, the pharmaceutical sector continues to utilize a drug development process that is slow, inefficient, risky and expensive. (p. 356)

The number of 'new drugs approved per billion US dollars spent on R&D has halved roughly every 9 years since 1950, falling around 80-fold in inflation-adjusted terms' notwithstanding the 'huge advances in many of the scientific, technological and managerial' factors that are inputs into the R&D process (Scannell et al. 2012:191). This deficiency is particularly acute in contexts of rapidly ageing populations who are experiencing rising levels of chronic illness amidst dramatically increased societal health budgets [World Health Organization (WHO) 2015], as well as in contexts of resource scarcity, where many cannot afford the medicines they need (Nathan 2007). It is argued in this article that a deficiency in the knowledge management literature exists, as the current dysfunctional medical R&D paradigm, which currently lacks the capacity to deal effectively with potential global pandemics such as Ebola, Middle East Respiratory Syndrome (MERS), or with rapidly developing antibiotic resistance (Callaghan 2014), poses a serious potential threat to society. Specifically, what is absent from the literature is knowledge of how the current dramatic advances in the fields that underlie medical science can be transmitted to pharmaceutical R&D outcomes. This article seeks to offer a theoretical framework to develop this knowledge, and a synthesis of theory is undertaken, based on the argument that the current

paradigm of pharmaceutical innovation faces a threshold limit, and that there may be a methodology that can overcome these constraints.

Some authors, such as Kaplan and Haenlein (2011:253), have drawn a parallel between the rapid spread of lethal viral epidemics and the 'viral' spread of social media; such a comparison highlights the development of new technologies that can support a relatively new and rapid form of problem-solving that utilises the 'crowd', or large numbers of people, to improve the probability of success in R&D, part of a body of literature termed 'probabilistic innovation' (Callaghan 2015). The rapid mobilisation of knowledge flows using large numbers of people, or the 'crowd', is the subject of this growing body of literature, which may have important implications for how crowds can be used to solve problems in real time, or quickly enough to stop certain disasters unfolding.

Viral communication (Kaplan & Haenlein 2011) is one example of how technology has enabled a radical acceleration in the speed at which crowds of people communicate. This article seeks to stress the importance of drawing insights from the way viral social media and internet communications operate, as well as the way flash mobs work, in order to add to a synthesis of theory to develop a theoretical model of how the crowd can be used to solve knowledge problems on the scale of global epidemics. The rationale that underlies this process is that the spreads of viruses or epidemics have certain characteristics that are common to the spread of viral communications, and that a transdisciplinary meta-theoretic synthesis is necessary to develop systems of real-time problem-solving, or problem-solving research that can be undertaken under intense time constraints. It is argued that the development of this theoretical synthesis can be taken to represent the genesis of a new paradigm within knowledge management, or a 'second generation' of innovation theory, termed 'probabilistic innovation' (PI) (Callaghan 2015). The term PI is taken to reflect the way the probability of solving problems can be increased exponentially if the number of problem-solvers, or problem-solving inputs into problem-solving, can also be increased exponentially. Similarly, the speed at which problems can be solved is also taken to be a function of the extent to which the same probabilistic mechanisms can be harnessed in support of real-time problem-solving. Key to the development of this field, however, is the development of its theoretical foundations.

To understand how to unearth the underpinning body of theory that explains how crowds work and how crowd-based distributed knowledge management systems can contribute to real-time research systems, it is necessary to draw together theory from a wide range of literatures. This article seeks to do this, and hence to make a contribution to the knowledge management literature in the following ways.

Firstly, with the relatively recent emergence of crowdsourcing (Howe 2006) as an area of study in the academic literature, the field of crowdsourced R&D seems to be advancing rapidly, but seems to lack a coherent theoretical structure that relates its constituent elements. PI as a developing field seeks to offer this structure. The core theoretical structure of the crowd-based problem-solving literature devoted to PI currently has its roots in: (1) economic theory and the problems posed by decentralised (Hayek 1945), or tacit (Von Hippel 1994), knowledge; (2) theory of how crowds in the form of markets solve pricing problems (Fama 1970, 1995; Smith 1962, 2003[1764]); (3) swarm intelligence theory (Bonabeau & Théraulaz 2000; Dorigo, Bonabeau & Theraulaz 2000; Garnier, Gautrais & Theraulaz 2007); (4) collective intelligence theory (Malone & Klein 2007; Woolley et al. 2010), as well as (5) the use of the crowd as applied to disaster management (Zook et al. 2012) as well as how crowdsourcing can be used to reorganise work to support accelerated problem-solving (Niederer & van Dijck 2010).Whereas these constituent dimensions contribute useful knowledge to this emerging field, what is lacking is their integration into a single theoretical framework with a clear rationale which can inform the further development of the field as a unified area of study within the area of knowledge management. This article attempts to contribute to this vision, and offers a theoretical framework in support of this.

Secondly, given wide acknowledgement of the decreasing returns to investment of pharmaceutical R&D (Gassmann & Reepmeyer 2005; Grasela & Slusser 2014; Horrobin 2000; Martin & Scott 2000; Munos 2009; Nathan 2007; Scannell et al. 2012), and its failure to address problems such as Ebola outbreaks, rapidly increasing antibiotic resistance or microbial drug resistance in general within a global context of dramatically rising healthcare costs (which are particularly problematic in certain regions with ageing populations or severe resource constraints), knowledge management theory is needed to address these knowledge problems in a way that is useful to R&D practitioners and theorists. In light of these challenges, this article argues that a new paradigm in R&D is necessary to solve the knowledge problems of a new era. In light of these challenges, this article builds on theory relating to the PI paradigm and second generation innovation, which is differentiated from first generation innovation on the basis of its use of probabilistic mechanisms (Callaghan 2016), which in turn is derived from the open innovation literature (Chesbrough 2011), in an attempt to provide useful insights into how real-time research problem-solving might be enabled.

Thirdly, the field of knowledge management, through its specialised focus on knowledge management systems, requires a theoretical 'bridge' from knowledge theory across to other fields to which it can contribute theoretical and practical problem-solving insights. This article seeks to build linkages between crowd-based knowledge management theory and its application in the scientific field in general and in medical or pharmaceutical R&D in particular. Given the serious problems outlined above, theory from a wide range of disciplines needs to be integrated into a framework that is useful to scientists in general. This article therefore seeks to integrate (into the body of PI theory described above) further insights that relate to the phenomena of flash mobs, viral

marketing and congestion theory to provide an overarching framework that can serve certain knowledge needs of scientists in many fields. This framework therefore builds on the use of crowdsourcing, which is increasingly being applied to solve research problems in the medical field (Adams 2011; Armstrong et al. 2012; Foncubierta-Rodriguez & Müller 2012; Sims et al. 2014; Yu et al. 2013).

Having outlined the objectives and rationale behind the paper, it proceeds as follows. Firstly, proof of concept as it relates to crowd-based problem-solving is considered to demonstrate the salience of the growing body of literature that forms the basis for the developing PI theoretical framework. Certain challenges to the development of the field are then acknowledged, and the importance of a constant and ongoing synthesis of theory is stressed. Next, insights from the way flash mobs are able to mobilise members of the crowd are identified as examples from which PI theory can be developed, as are the ways flash teams can offer insights into how teams can be formed almost instantaneously and be designed to grow or to break up in response to the requirements of the problems they solve, or in response to the 'problem landscape'. The focus then shifts from the mobilisation of the crowd in support of problem-solving to operationalisation, and the role of artificial intelligence (AI) and theory relating to congestion is briefly considered as a prelude to the discussion of an overarching PI theoretical framework, as well as certain arguments in support of its theoretical synthesis. The paper then concludes with a summary of its arguments and recommendations for further research.

Proof of concept and the need for a unifying theoretical framework

The field of crowd-based problem-solving is developing rapidly. Different terminologies have emerged, that relate to essentially the same domain. At the heart of this domain is the notion that large numbers of people, or crowds, can be harnessed to solve problems. Zhai et al. (2011:879) offer the notion of expert-citizen engineering, or citizen engineering, described as 'a concept that engages a cohort of physically dispersed citizens connected by the Internet to collaboratively solve real-world problems through massive cooperation'. Zhai et al. (2011) explain this concept as follows.

> With advances in information technology, we can build transformative cyber-infrastructures to effectively leverage the 'wisdom of crowds'… Regarding the citizen engineers who function as the main contributors, there is a wide spectrum of human resources that that crowdsourcing system designers can harness- from amateurs/hobbyists, lacking practical experience, to experts/licensed engineers, with years of professional training. As such, we are encouraged to investigate proper approaches to design CEs that can sufficiently engage and support expert citizens who have unique needs that may be different from those of amateur citizen engineers. (p. 879)

The implementation of second generation innovation systems and the principles of PI has been made possible by recent advances in information technologies. Zhai et al. (2011) also offer the following vision of these recent information technology advances.

> Emerging information technologies empower us to build transformative cyber-infrastructures. Characterised by broadband networks, high performance processors, these novel technologies have facilitated expansive collaboration among users scattered across many physical and institutional locations. (p. 879)

On the back of rapidly developing technological capabilities, a host of different platforms have developed that can enable crowd-based work and problem-solving. A range of different platforms in the form of online marketplaces have emerged which allow for crowdsourced work, such as Amazon's Mechanical Turk (AMT); this type of platform, however, is limited to tasks that are mutually independent, of shorter duration, and less cognitively challenging (Zhai et al. 2011). Zhai et al. (2011:880) argue that while the development of Wikipedia required only about 100 million brain hours to develop, much more than this is spent by crowds on leisure activities per year, and a cognitive surplus exists within the crowd which can be captured by citizen engineering, where 'researchers are encouraged to develop well-designed mechanisms and methodologies to channel and motivate humans to solve challenging problems that computers cannot yet handle well'. Examples of this include the development of Mozilla Firefox and the Apache Web Server, and other examples of 'proof of concept' relating to crowdsourced projects in the literature include eBird, Galaxy Zoo, Foldit, People-Centric Sensing, Knowledge Collection, Stardust@home, Human Search Engine, Crowd Photo Tagging, Participatory Risk Management (PRM), and Online Team Gaming (Zhai et al. 2011).

Despite the widespread success of crowd-based platforms in solving complex problems, the field is still new and its potential relatively untapped; however, certain challenges exist that will need to be addressed for the field to move forward, particularly in terms of crowdsourced R&D, which will typically require some proportion of expert input from the crowd. According to Zhai et al. (2011:880), there are three primary challenges faced in attempts to develop expert-citizen engineering projects: (1) task complexity (associated with the need for high human intelligence and the need for advanced levels of skills as well as the need to 'conduct a whole range of experiments to provide objective, insightful and trustworthy consultancy'), (2) recruitment difficulty (because of the complexity inherent in tasks, 'available human resources are limited and membership eligibility is rather selective, compared to traditional crowdsourcing tasks') and (3) resource requirements (complicated tasks can require sophisticated analysis tools and computational resources; for example, current analysis and design methods, such as 'nonlinear finite element analysis and design methods, such as nonlinear finite element analyses of complex structures, can overstress in-house computational capabilities of many firms and laboratories and far exceed the resources of most citizen engineers').

The system design functions of crowdsourced R&D systems can also be provided by the crowd itself. However, to do this system designers face a further challenge as crowd-based users have diverse backgrounds and malicious users can also create challenges; practicable workflows are necessary to achieve an effective aggregation of results and to maintain quality control (Zhai et al. 2011). Zhai et al. (2011:886) stress that to 'leverage the expertise from skilled citizens, we need to develop new principles and theories that can guide system designs to satisfy the unique needs of high level users'. The PI field will need to borrow theory from wherever it can to supplement the theory it develops in order to develop system design processes that can manage very large numbers of problem-solving inputs in real time. At this point in time these challenges seem daunting. However, given the rapid speed of technological development it might be possible to accelerate progress towards this end as long as an overarching theoretic framework can exist to guide these developments.

The mobilisation of the crowd

An important dimension of theory development for the PI field relates to challenges associated with the need to mobilise the crowd. A flash mob is formed by groups of people that semi-spontaneously form in public space, typically for the purposes of performance or as part of a guerrilla marketing strategy, a process that is made possible by social media (Grant, Bal & Parent 2012). The literature related to flash mobs and the development and management of 'flash teams' offers further insight into how the PI field might integrate useful insights from these phenomena. Brejzek (2010) offers the following description of the effects of flash mobs.

> Since 2003, stunned commuters, shoppers, sales staff and politicians have been confronted unexpectedly with flocks of seemingly unrelated people congregating in the central business districts (CBDs) from Leipzig to Teheran, London to Munich. Performing nonsensical actions, the individuals tend to disperse shortly after their action has taken place without so much as a word to each other. (p. 112)

As a cultural phenomenon, flash mobs have been described as a manifestation of the physicalisation of viral culture (Brejzek 2010; Wasik 2009). What sets this phenomenon apart from other social media applications is the way physical action is enabled and directed through a powerful effect that captures the imagination of the crowd. The ability to capture the imagination of crowd participants is an important dynamic, and the PI literature should not neglect this stream of its literature development. To invest effort in the crowdsourced R&D process and to produce almost instantaneous research results will require the mobilisation and motivation of large numbers of people. The flash mob conception can be related to the more technical aspects of managing problem-solving teams through the use of the concept of 'flash teams'.

Retelny et al. (2014:75) offer a framework for assembling and managing paid experts from the crowd, termed 'flash teams', which 'advance a vision of expert crowd work that accomplishes complex, interdependent goals such as engineering and design' according to sequences of 'linked modular tasks and handoffs that can be computationally managed'. Flash teams can be defined as 'computationally-guided teams of crowd experts supported by lightweight, reproducible and scalable team structures' which seek to 'embed the techniques of high performing offline teams within a model that can take advantage of computation's ability to abstract, scale, and visualise progress' (Retelny et al. 2014:77). The use of reproducible and scalable team structures are an example of systems that can be used to manage processes associated with very large numbers of problem-solvers as they populate a problem space. The ultimate goal of this process is to have the crowd also provide direction and interactive systems to manage R&D work as it unfolds in real time. The work of Retelny et al. (2014) is now given special attention in order to highlight these concepts, which are considered especially important to the emerging PI literature.

Interactive systems are used to plan and reconfigure the structures of these teams, so as to create larger organisational structures in response to user requests, immediately hiring in reaction to needs, and to 'pipeline intermediate output to accelerate completion times' (Retelny et al. 2014:75). Retelny et al. (2014:75) present an end-user authoring platform, namely 'Foundry', which allows uses to initiate modular tasks and manage teams though the stages up to handoffs of intermediate work; in this way crowdsourced design prototyping, course development and film animation is enabled, 'in half the work time of traditional self-managed teams'. This example is but one of many which can be used to highlight practical solutions to the problem of managing very large numbers of problem-solvers in real time, and to the problem of coordination.

According to Retelny et al. (2014:75), crowdsourcing systems 'coordinate large groups of people to solve problems that a single individual could not achieve at the same scale', and microtasking systems typically use 'highly-controlled workflows to manage paid, non-expert workers towards expert-level results'. However, this becomes more difficult when it comes to more complex real-world tasks requiring deep domain knowledge that is less easy to decompose into independent microtasks which any member of the crowd can complete; with regard to their suggested system of structured collaborations between experts from the crowd as a solution to this problem, Retelny et al. (2014) offer the following vision to:

> enable anybody with a napkin sketch of a design idea to ask the crowd to follow the user-centred design process and create a user-tested, high-fidelity prototype of that idea within twenty-four hours? (p. 75)

The importance of this body of work is in its focus on design; to have the crowd design its own solutions to design, system and quality assurance problems would be an important dimension of the PI literature.

Retelny et al. (2014) suggest that the use of expert crowd work can be designed around the concept of flash teams, organised around sequences of linked tasks, which have the same coordinating strength as more lightweight team structures while at the same time being used to leverage and support collaboration, automatically create teams, manage the size of these teams and instantaneously combine teams into larger organisations.

To do this, each task would need an input and an output, and end users would need to link modular tasks and each task's output becomes the input for the next task, as web applications are used to monitor the workflow as computational systems leverage this structure and create crowd dynamics; in this way, work can be pipelined, and in-progress work that is helpful to downstream tasks is passed along to support these tasks (Retelny et al. 2014). Research into these processes is important, as flash teams have the potential to leverage the scale of paid crowdsourcing for expert work, going further than volunteer crowd systems, and the scaling process can be ramped up through computational management of an elastic work force, resulting in complex work at crowd scale as the structures of traditional organisations are automated (Retelny et al. 2014). The vision here is to have an organisational structure, or structures, that morph in real time to the contours of the problem space. This process is illustrated in Figure 1.

The phenomenon of swarm intelligence (Bonabeau & Théraulaz 2000; Dorigo, Bonabeau & Theraulaz 2000; Garnier, Gautrais & Theraulaz 2007) can be taken to provide further insights into the process whereby very large numbers of problem-solvers can 'populate' the problem space of a particular problem landscape. The importance of swarm intelligence in this process relates to how large numbers of problem-solvers can work on a problem in the absence of central direction. For example, ants build complex architectural structures, complete with passages and antechambers, without central coordination, using a process termed stigmergy, where each individual ant reacts to its point of contact with the 'problem space' and 'erodes' this problem space at that point on the front-line of the problem landscape. This process is also illustrated in Figure 1, which also seeks to illustrate the process whereby flexible organisational structures can form, grow or break up in real time as needed (Retelny et al. 2014), according to the unique configuration of the problem landscape. An example of this process is the case of proteomics research, where an almost infinite number of permutations of protein strings represents the problem landscape and 'first generation' innovation or R&D processes (which do not use probabilistic mechanisms) are simply not able to populate this landscape with the many thousands of researchers required to make a difference in such a large problem space. PI may offer an opportunity for real-time research in large problem space contexts such as proteomics and genetics, or in the 'new' pharmaceutical space dominated by (large molecule) protein research.

Figure 1 therefore seeks to illustrate two dimensions of the challenges PI needs to solve to become a successful field, namely the problem of coordination of large numbers of the crowd engaged in large scale problem-solving in the absence of (or under constrained conditions of) central direction, and the way individuals can come together almost instantaneously
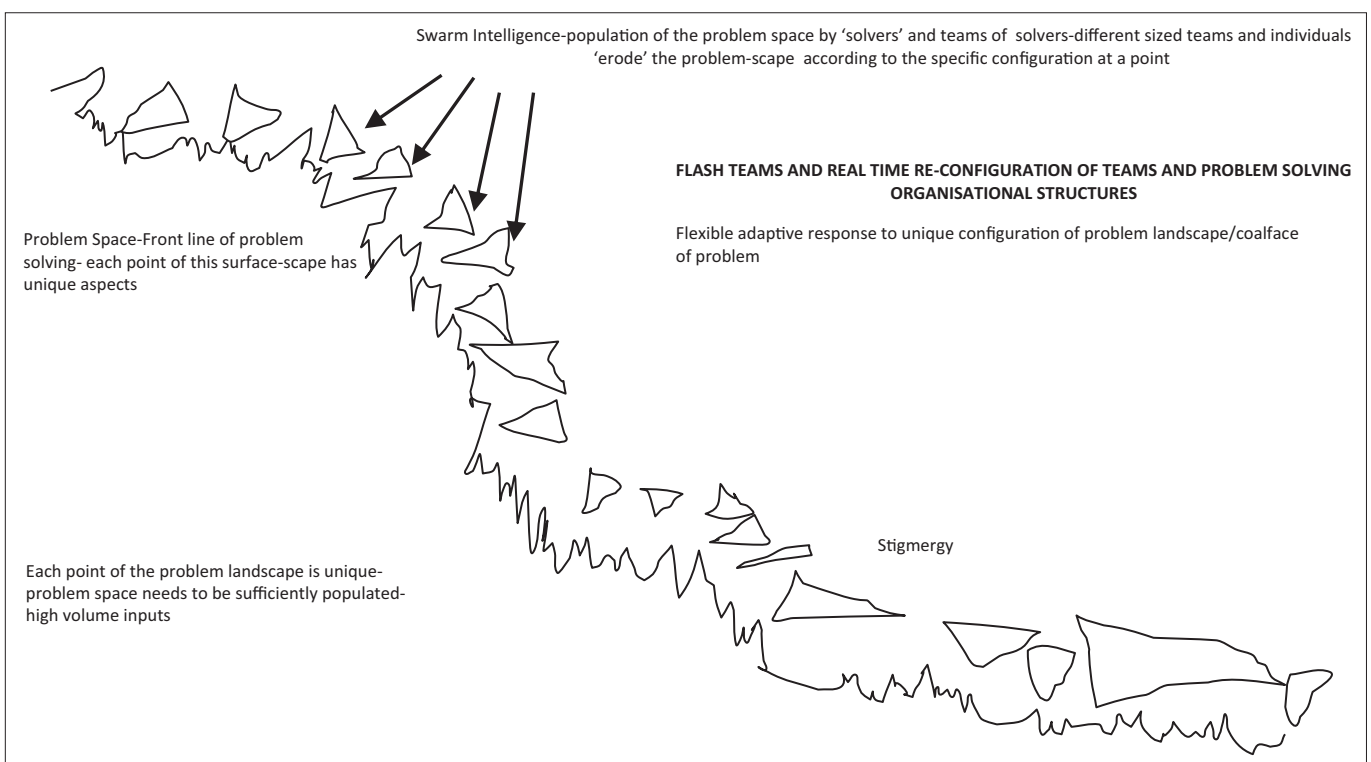

FIGURE 1: Erosion of problem space driven by swarm intelligence configurations.

and form micro-organisations in real time in response to specific problem-solving needs at a specific point on the problem landscape configuration. It is acknowledged that each point of contact with the problem landscape may have unique characteristics relevant to problem-solving.

As stressed previously, the mobilisation of the crowd in support of problem-solving is the focus of a rapidly developing body of literature. Crowdsourcing can successfully mobilise large numbers of problem-solvers, yet to date platforms have developed that are well suited to tasks requiring few skills, such as AMT, and those using amateur input, such as Foldit, and most crowdsourcing workflows and algorithms now aim to produce expert-level performance from non-expert contributions (Retelny et al. 2014). Examples of success in this process include MapReduce frameworks which channel crowdsourced inputs into encyclopaedia entries, document editing, translation and visual question answering, which can be optimised using AI (Retelny et al. 2014). The transition of crowdsourced problem-solving, from successfully addressing problems that require a limited skill set to being able to solve problems that are extremely complex, may be the core theoretical and practical problem that the emerging PI literature needs to resolve.

Whereas certain insights can be gleaned from how flash mobs successfully capture the imagination of the crowd and are successful in mobilisation, further insights can be drawn from the use of different platforms that seek to operationalise the problem-solving potential of the crowd.

Post-mobilisation: Operationalisation

According to Retelny et al. (2014), problem-solving crowdsourcing that recruits experts has to date typically been restricted to single-expertise areas and have been 'one offs', but what is needed are platforms such as Foundry, which can support large numbers of tasks on demand at using higher-level work flows based on expert knowledge. The capacity of crowds to undertake expert work using non-experts and to leverage the skills of experts can be enabled using visual workflow and management tools.

More complex problem-solving processes can be enabled using visual workflow tools and management tools like Gantt charts which are based on a visual timeline language; these processes can be designed to facilitate worker interest, honesty, and motivation, and a process of visible collaborators with clear goals underpinned by class hierarchies which integrate with business processes in a team-based context (Retelny et al. 2014). Theory from the field of organisational behaviour already provides insights into how challenges to effective team coordination, such as geographic dispersion, technology-mediated communication and dynamic changing team membership can be managed in expert crowdsourcing (Retelny et al. 2014). Similarly, theory from the field of AI can offer complementary perspectives of how the crowdsourced R&D process can be managed.

AI can therefore offer further important insights into how the crowd can contribute to real-time problem-solving. In practical terms, propositional methods of planning algorithms can be used to convert planning challenges into 'propositional conjunctive normal form formulas for solution using systematic or stochastic' methods; interleaved planning and execution is now part of the AI landscape (Weld 1999:93). The speed at which classical planning problems can be solved has increased exponentially over time (Weld 1999). Theory related to the planning and management of flash teams is therefore especially salient in three areas: (1) in how encoding can be used to align responsibilities and coordination using structures that manage shared space and shared work around clearly defined work roles, (2) in how management modularity theory predicts how loosely coupled system components with standardised interfaces can be used in multiple configurations and (3) in how multiple integration mechanisms, including pipelining, structured handoffs, and directly responsible individuals (DRIs) can address weaknesses in key points of coordination (Retelny et al. 2014).

Key to the success of expert crowdsourced teams is the need to quickly comprehend issues related to shared work, interdependencies and roles, which, when complemented by team structures, modularity, and coordinating mechanisms allow for the leverage of automation, computation, the economies of scale of the crowd and the flexibility of the crowd (Retelny et al. 2014). However, if these economies of scale are achieved, the management of congestion and overcrowding at the surface of the problem space also needs to be considered.

A consideration of how very large numbers of the crowd can solve problems in real time would therefore not be complete without a discussion of congestion, and theory related to this is potentially also an important dimension of the PI theoretical framework. With regard to congestion and transport systems, Vickrey (1969) makes the following observations.

> Investment in transport facilities necessarily begins by being largely investment in the provision of new routes or new services under conditions of substantial indivisibilities and increasing returns to scale. Under these conditions the usual profitability tests for determining the desirability of specific investments lead generally to under- rather than to over- investment in transportation facilities. At this stage, cost-benefit analysis needs to include substantial elements of consumers' surplus on the benefit side in order to arrive at correct evaluations. As investment proceeds, however, larger and larger proportions of transportation investment are made primarily, or at least in large measure, to relieve congestion on existing routes and to expand overall capacity. In such instances criteria based on apparent profitability may be seriously misleading in the opposite direction, and when notions of consumers' surplus are narrowly applied without regards to the overall situation, the errors may be compounded. (p. 251)

In his seminal work, Vickrey (1969:251) offers six categories of congestion: (1) simple interaction, multiple interaction, bottleneck, triggerneck, network and control, and general

density. A synthesis of theory that seeks to show how to exponentially increase the numbers of problem-solvers that populate a problem space will need to draw insights from the categorisation of different types of congestion to manage what will, if successful, become a congested 'space'. At some point, a critical mass might be reached in PI processes that successfully mobilise very large numbers of problem-solvers, and the critical success factors of the process might shift towards managing congestion in the problem-solving system. This challenge might be particularly acute under a real-time temporal constraint.

These categories are now briefly outlined, using Vickrey's (1969) transport analogies: (1) single interaction relates to the case where two units pass each other in a way that requires a delay so as to avoid collision; these are typical to light traffic, and congestion delay 'tends to vary as the square of the volume of the traffic' and a driver will experience a similar delay to what he or she causes; (2) multiple interaction refers to conditions at high levels of traffic density, but short of capacity (between 0.5 and 0.9 in capacity), under which one additional vehicle can contribute a multiple of the congestion it experiences; (3) the pure bottleneck situation relates to the presence of a relatively short route segment with a fixed capacity that is not sufficient to meet demand; (4) the triggerneck occurs when queues at a bottleneck interfere with other traffic not on route through the bottleneck; (5) network and control congestion occurs when peak traffic requires control measures in order to manage it, which in turn may slow traffic; and (6) general density of traffic can lead to long-run increasing costs, as more routes need to be planned and constructed.

The rationale behind the inclusion of this congestion typology here is simply to suggest a central place for theory related to congestion in the developing theoretic framework of PI, given the challenges of coordination and management of very large numbers (high volume traffic) of problem-solvers. Intuitively, the PI framework can seem to be a vision that is a 'bridge too far', or a vision that seems difficult to attain at this current time, but it is argued that with a comprehensive incorporation of theory this vision will be possible to attain, and the management of congestion in the crowd may have a central place in this framework. Understanding the mechanisms whereby congestion can be managed by some form of internalising pricing mechanisms, or by imposing differentiated costs on congestible behaviours, is one dimension of this.

According to Vickrey (1969:258), congestion can be managed by allowing pricing to allocate transport flows, as pricing 'makes it possible to exclude the low-value uses and base the magnitude of the improvement primarily on the uses that are valued sufficiently highly so that they warrant the marginal cost of the final increment to the magnitude of the improvement'. As with other flows of traffic, if airport landing fees reflect congestion costs then relief from congestion is possible; where 'charges for the use of alternative routes fail to reflect congestion costs at the margin', then congestion will be problematic (Vickrey 1969).

These charges also have an informational role, as they can provide information on capacity as well (Vickrey 1969). Understanding the relative value of different types of problem-solving inputs might provide a way to ensure that the costs of congestion are balanced with benefits; further research might do well to build on this work in anticipation of the success of the crowd mobilisation process.

Having considered certain aspects of how knowledge of flash mobs can contribute to knowledge of how to mobilise crowds, more specific knowledge was considered, of how crowds, once mobilised, could be managed in such a way as to change their configurations and their permutations of team structures, from individual to micro-organisational, according to the dictates of the problem space. Under conditions of successful crowd mobilisation, however, the management of congestion arguably becomes an increasingly important challenge, and Vickrey's (1969) congestion theory was therefore briefly reviewed. At this point, it is necessary to provide an integrative logic to the discussion, and Figure 2 is used for this purpose.

Synthesis

Figure 2 illustrates a theoretical framework that can be taken to underlie the development of the field of PI at this point in time. To capture PI advantages such as extremely large economies of scale in relation to accelerated problem-solving that approaches real-time capabilities, certain knowledge is first required as to how large scale crowd-based problem-solving systems work. The theory discussed thus far in this article is now included in a broader theoretical synthesis, which seeks to provide an overarching overview of how different streams of literature fit together to enable the PI vision. The seminal basis of this body of work remains the work of Smith (2003 [1764]).

The solving of the problem of resource distribution on a large scale by the market mechanism, or the 'invisible hand' (Smith 2003 [1764]) is an example of how very large 'crowds' of people solve the knowledge aggregation problem, notwithstanding the existence in certain cases of market failure. At the heart of the challenges facing real-time crowdsourced R&D is the knowledge aggregation problem (Hayek 1945; Von Hippel 1994), or the problem of how to distribute knowledge and how to bring together knowledge that is geographically dispersed and difficult to find. To some extent, the solution to the knowledge aggregation problem as it relates to crowdsourced R&D has two aspects: a global knowledge aggregation, or macro, dimension and a more localised, or micro, dimension. The former relates to how very large numbers of people can be mobilised and motivated to contribute inputs into the problem-solving process. The latter relates to how the knowledge inputs are integrated and processed at the individual level, so that they can be channelled into real-time R&D outputs (while managing the congestion associated with a deliberate 'overpopulation' of the problem space.

The 'macro' process holds the key to capturing the probabilistic effects associated with large numbers. The central limit of
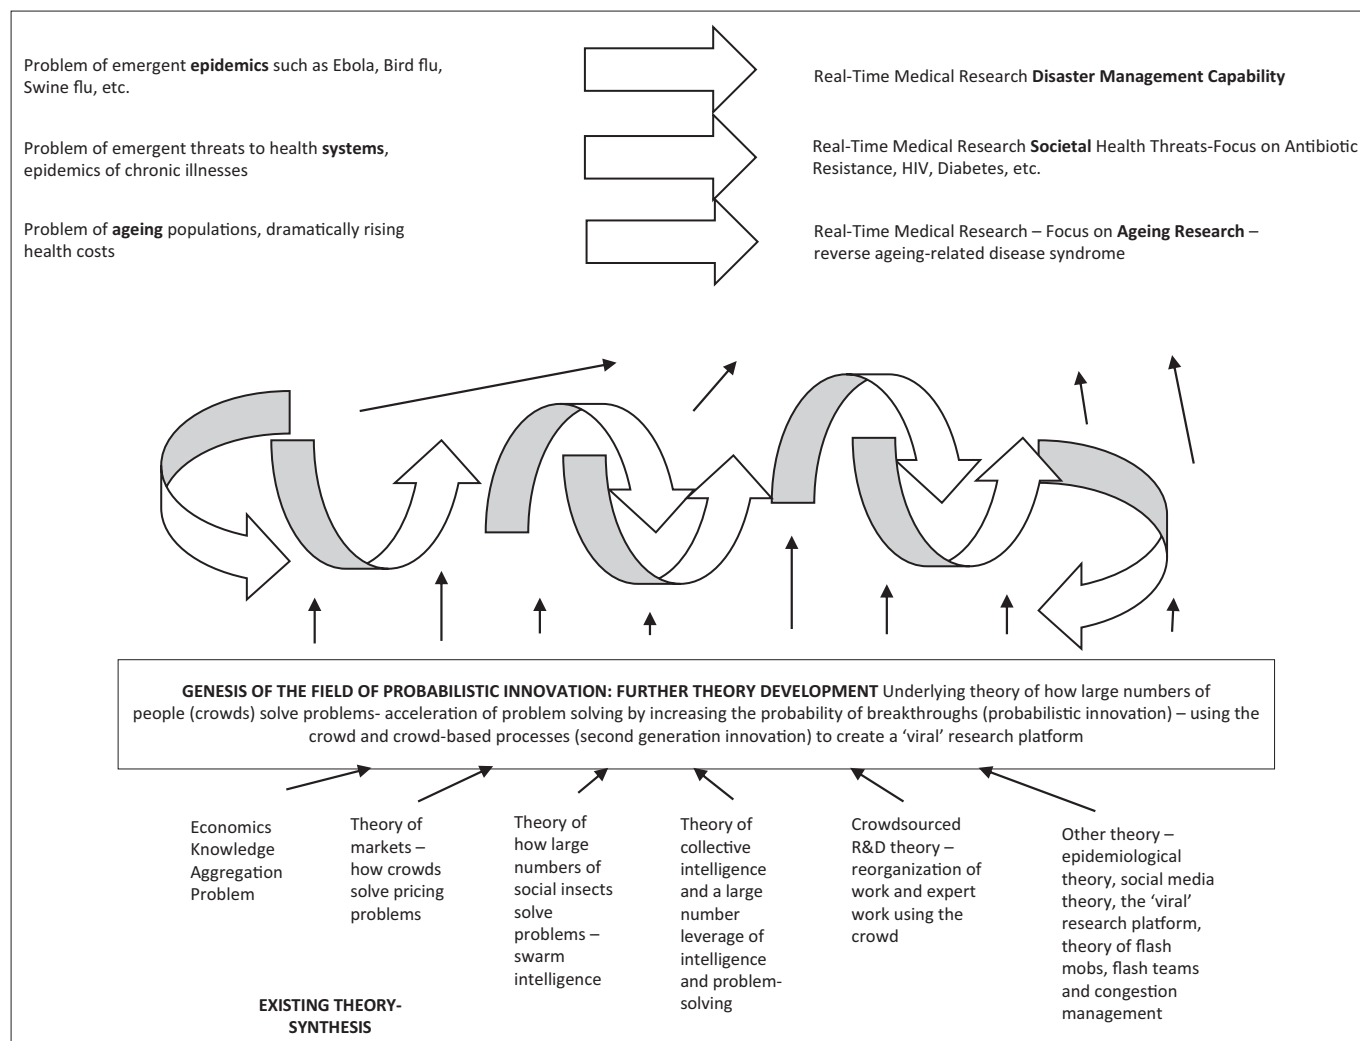
FIGURE 2: From theory synthesis to theory development: The field of probabilistic innovation and what it offers.

probability theory (Bernstein 1944) suggests that as numbers in a sample increase, certain distributional effects occur. Galton (1907) drew attention to the way crowds, if large enough, can be effective in solving certain types of problems, offering the example of a crowd of 800 at a stock and poultry exhibition who provided individual estimates of how much an ox would weigh after it was processed (and who were able to, on median aggregate, predict this within 0.008 of its value: the median estimate was only nine pounds off the weight of 1198 pounds). While evidence abounds of the effectiveness of the crowd to solve scientific research problems, as in the case of InnoCentive (Howe 2006), it is argued here that little is known of the upper limits, or ultimate potential, of the collective crowd in problem-solving; different literatures seem to offer up different pieces of this puzzle, certain of which are included in Figure 2. However, it is argued that this body of theory might be just 'the tip of the iceberg' and part of a wave of literature that is building up, and which will 'break' at some point in the near future, and will yield benefits in human health, medicine and science itself.

Galton (1907) stresses the effectiveness of democracy as another example of what he terms 'the wisdom of crowds'. An objective of PI as an emerging field would be to apply a

decomposition analysis to this phenomenon (the wisdom of crowds), to understand the causal effects that make crowds effective at solving certain problems, over and above how the knowledge aggregation problem is solved through the use of the crowd. However, to understand the wisdom of crowds, one first has to take recourse to other bodies of theory that illustrate crowd problem-solving in its different forms. Galton's (1907) example of a betting market also echoes seminal work on how crowds solve pricing problems in larger betting markets such as stock markets (Fama 1970, 1995; Smith 1962). Galton's (1907) betting market analysis has been taken up in work such as Hanson's (1995, 2000, 2003), which argues that betting markets are very effective at solving knowledge problems and that the principles that underlie them can also be applied to the research process itself as well as many other applications. Collective problem-solving can be found in other contexts, and biological examples exist, such as the use of collective intelligence on the part of insects.

Swarm intelligence (Bonabeau and Théraulaz 2000; Dorigo, Bonabeau & Theraulaz 2000; Garnier, Gautrais & Theraulaz 2007) offers insights into how social insects like ants can build advanced architectural structures as a swarm without central direction or central planning (and no blueprints).

This has relevance for crowdsourced R&D because when very high numbers of contributors provide knowledge inputs into problem-solving, they need to 'populate' the problem space in a way that can be independent of the need to direct the problem-solving process- to be effective even in instances where the processes of central direction are not, or cannot, cope with congestion or volumes of inputs.

This is an important theoretical stream of literature as applied to PI, which is expected to grow in its importance over time, as exponentially larger numbers of problem-solvers become part of the KiehraSearch process (to be discussed shortly). Theory that draws from swarm intelligence is part of the bedrock of theory that is taken to contribute to the development of the field of PI. In Figure 2, certain theoretical frameworks are shown, which are drawn from different fields that contribute to the 'bedrock' of the larger PI meta-theoretic platform on which further theory development needs to take place. The bodies of theory which form part of this 'bedrock' include collective intelligence theory (Malone & Klein 2007; Woolley et al. 2010) which seeks to leverage the problem-solving abilities of groups, and disaster management principles (Zook et al. 2012) which can be used to guide the PI process in attaining real-time capabilities to solve serious problems under serious time constraints.

Theory relating to how crowdsourcing can reorganise work is also important, as a rapidly growing area of literature is devoted to new work systems and to the way work is being reorganised to be performed by crowds (Niederer & van Dijck 2010); this body of literature offers novel micro-level insights into how large scale problem-solving can be effective. This literature also offers insights into how to differentiate processes between expert work and non-expert work, and how far this 'front-line' can be extended; how far non-experts can be enabled to perform expert work or to support expert work. However, to recruit and manage exponentially increasing numbers of knowledge inputs, or capturing 'viral effects' also requires multiple platforms, or a proliferation of problem-solving platforms. An example of such a platform is therefore included in Figure 2.

In Figure 2, the KiehraSearch platform is taken to represent an example of an operational platform, or website, dedicated to solving knowledge problems and developing real-time research processes to do this. The KiehraSearch platform is an example of a crowdsourced R&D platform that is under development but is similar to other platforms such as InnoCentive, the difference being that it is non-profit in nature; it is used here simply to illustrate the need for such platforms to proliferate to also populate the problem space on the macro level. In other words, platforms that seek to address these problems need to also be developed in large numbers in order to support the exponential increase in problem-solvers required for PI to develop.

The non-profit nature of these platforms can become a problem-solving advantage in that knowledge inputs can be fed back into the crowd to develop a 'three-dimensional'

problem-solving space (Callaghan 2014), as this might enable and accelerate innovation and knowledge creation exponentially. As a research methodology, PI is perhaps a necessary complement to existing research systems, and should not be used to supplant current R&D systems and the profit-seeking model of innovation, as PI systems are ideally suited only to large scale foci and the extensive mobilisation of resources, which may only be appropriate for target problems which have certain characteristics.

Figure 2 shows a clustering of three types of problems that might be uniquely suited to PI interventions: (1) emergent epidemics require the development of a disaster management capability, and it is argued that PI is well suited to this; (2) the large scale threats posed by antibiotic resistance and the chronic disease burden on society is another area that may be well suited to the large scale resource mobilisation associated with PI; and (3) the problem of ageing populations in a context of rapidly rising health costs poses another societal problem that might be uniquely suited to problem-solving using probabilistic platforms. While it may never be possible to reverse ageing, it might be possible to solve many of the diseases that hasten the ageing process, and thereby to decrease the pressures on national health budgets through increasing levels of health in spite of the ageing process.

Conclusion

The objective of this article was to outline a theoretical framework for the development of PI as a field of academic enquiry, and to identify different literatures as candidates for inclusion in this framework. Firstly, dangers associated with the absence of a global real-time problem-solving system in the face of unsolved knowledge problems such as epidemics like Ebola as well as other threats such as increasing antibiotic resistance were highlighted. The concept of expert-citizen engineering was introduced, to locate the crowdsourced R&D literature that followed in relation to notions of expert crowd problem-solving. Literature was then considered that was related to the mobilisation of the crowd in support of real-time problem-solving, with a specific focus on the temporal aspect of flash mobs and the use of flash teams as a potentially important stream of future research. Next, the role of AI, swarm intelligence and congestion theory within this overarching framework was discussed. The paper concluded with the development of an overview of the theoretical framework of PI as a field currently, and three practical targets of the PI process, namely the development of: (1) real-time research disaster management capability suited to managing global epidemics like Ebloa, (2) the development of real-time research with the potential to tackle societal health threats, or solving the current 'epidemic' of chronic illnesses such as diabetes as well as problems such as rising levels of antibiotic resistance, and (3) real-time research capability to focus on ageing research, as rising health costs associated with ageing populations pose threats to societies. In all, this article sought to offer some sort of 'road map' for the development of the PI field, in a way that sought to include the theory-search process, or the process whereby

theory-search and incorporation contributes to a synthesis of literatures, and a vision of potential outcomes for the field. Although there is perhaps little in common between the spread of lethal viral epidemics and the viral spread of social media-enabled communication, what does seem to link them is the notion that these are both highly effective processes; and it is hoped that the fledgling field of PI can offer concrete benefits to science by continuing to develop theory based on insights from the study of other highly effective processes, or processes that explicitly harness the probabilistic forces of innovation.

Acknowledgements

Competing interests

References

Adams, S.A., 2011, 'Sourcing the crowd for health services improvement: The reflexive patient and "share your experience" websites', *Social Science and Medicine* 72(7), 1069–1076. https://doi.org/10.1016/j.socscimed.2011.02.001

Armstrong, A.W., Harskamp, C.T., Cheeney, S., & Schupp, C.W., 2012, 'Crowdsourcing for research data collection in rosacea', *Dermatology Online Journal* 18(3), 15.

Bernstein, S.N., 1944, 'Extension of the central limit theorem of probability theory to sums of dependent random variables', *Uspekhi Matematicheskikh Nauk* 10, 65–114.

Bonabeau, E. & Theraulaz, G., 2000, 'Swarm smarts', *Scientific American* 282(3), 72–79. https://doi.org/10.1038/scientificamerican0300-72

Brejzek, T., 2010, 'From social network to urban intervention: On the scenographies of flash mobs and urban swarms', *International Journal of Performance Arts and Digital Media* 6(1), 111–124. https://doi.org/10.1386/padm.6.1.109_1

Callaghan, C.W., 2014, 'Solving Ebola, HIV, antibiotic resistance and other challenges: The new paradigm of probabilistic innovation', *American Journal of Health Sciences* 5(2), 165–178. https://doi.org/10.19030/ajhs.v5i2.8972

Callaghan, C.W., 2015, 'Crowdsourced 'R&D' and medical research,' *British Medical Bulletin* 115, 1-10. https://doi.org/10.1093/bmb/ldv035

Callaghan, C.W., 2016, 'Knowledge management and problem solving in real time: The role of swarm intelligence', *Interdisciplinary Journal of Information, Knowledge, and Management* 11, 177–199.

Chesbrough, H., 2011, 'Pharmaceutical innovation hits the wall: How open innovation can help', *Forbes,* viewed 11 April 2015, from http://www.forbes.com/sites/henrychesbrough/2011/04/25/pharmaceutical-innovation-hits-the-wall-how-open-innovation-can-help/

Dorigo, M., Bonabeau, E. & Theraulaz, G., 2000, 'Ant algorithms and stigmergy', *Future Generation Computer Systems* 16, 851–871. https://doi.org/10.1016/S0167-739X(00)00042-X, https://doi.org/10.1016/S0167-739X(00)00041-8

Fama, E., 1970, 'Efficient capital markets: A review of theory and empirical work', *The Journal of Finance* 25(2), 383–417. https://doi.org/10.1111/j.1540-6261.1970.tb00518.x, https://doi.org/10.2307/2325486

Fama, E., 1995, 'Random walks in stock market prices', *Financial Analysis Journal* 76, 75–80. https://doi.org/10.2469/faj.v51.n1.1861

Foncubierta-Rodriguez, A. & Muller, H., 2012, 'Ground truth generation in medical imaging: A crowdsourcing-based iterative approach', in *Proceedings of the ACM Multimedia 2012 Workshop on Crowdsourcing for Multimedia*, ACM (Association for Computing Machinery), Nara, Japan, October 29, pp. 9–14.

Galton, F., 1907, 'Vox populi (The wisdom of crowds)', *Nature* 1949(75), 450–451. https://doi.org/10.1038/075450a0

Garnier, S., Gautrais, J. & Theraulaz, G., 2007, 'The biological principles of swarm intelligence', *Swarm Intelligence* 1, 3–31.

Gassmann, O. & Reepmeyer, G., 2005, 'Organizing pharmaceutical innovation: From science-based knowledge creators to drug-oriented knowledge brokers', *Creativity and Innovation Management* 14(3), 233–245. https://doi.org/10.1007/s11721-007-0004-y

Grant, G.S., Bal, A. & Parent, M., 2012, 'Operatic flash mob: Consumer arousal, connectedness and emotion', *Journal of Consumer Behaviour* 11, 244–251. https://doi.org/10.1002/cb.384

Grasela, T.H. & Slusser, R., 2014, 'The paradox of scientific excellence and the search for productivity in pharmaceutical research and development', *Clinical Pharmacology and Therapeutics* 95(5), 521–527. https://doi.org/10.1038/clpt.2013.242

Hanson, R., 1995, 'Could gambling save science', viewed 1 December 2014, from http://li.mit.edu/Stuff/CNSE/Paper/Hanson90.pdf

Hanson, R., 2000, 'Shall we vote on values, but bet on beliefs?', viewed 1 December 2014, from http://hanson.gmu.edu/futarchy2000.pdf

Hanson, R., 2003, 'Combinatorial information market design', *Information Systems Frontiers* 5(1), 107–119. https://doi.org/10.1023/A:1022058209073

Hayden, E.C., 2014, 'Technology: The $1,000 genome', *Nature* 507, 294–295. https://doi.org/10.1038/507294a

Hayek, F.A., 1945, 'The use of knowledge in society', *The American Economic Review* 35(4), 519–530.

Horrobin, D.F., 2000, 'Innovation in the pharmaceutical industry', *Journal of the Royal Society of Medicine* 93, 341–345.

Howe, J., 2006, 'The rise of crowdsourcing', *Wired Magazine*, 14 June, viewed 27 May 2013, from http://sistemas-humano-computacionais.wdfiles.com/local--files/capitulo%3Aredes-sociais/Howe_The_Rise_of_Crowdsourcing.pdf

Kaplan, A.M. & Haenlein, M., 2011, 'Two hearts in three-quarter time: How to waltz the social media/viral marketing dance', *Business Horizons* 54, 253–263. https://doi.org/10.1016/j.bushor.2011.01.006

Kaitin, K., 2010, 'Deconstructing the drug development process: The new face of innovation', *Clinical Pharmacology and Therapeutics* 87(3), 356–361. https://doi.org/10.1038/clpt.2009.293

Malone, T.W. & Klein, M., 2007, 'Harnessing collective intelligence to address global change', *Innovations* 2(3), 15–26. https://doi.org/10.1162/itgg.2007.2.3.15

Martin, S. & Scott, J.T., 2000, 'The nature of innovation market failure and the design of public support for private innovation', *Research Policy* 29, 437–447. https://doi.org/10.1016/S0048-7333(99)00084-0

Munos, B., 2009, 'Lessons from 60 years of pharmaceutical innovation', *Nature Reviews Drug Discovery* 8, 959–968. https://doi.org/10.1038/nrd2961

Nathan, C., 2007, 'Aligning pharmaceutical innovation with medical need', *Nature Medicine* 13(3), 303–308. https://doi.org/10.1038/nm0307-304

Niederer, S. & Van Dijck, J., 2010, 'Wisdom of the crowd or technicity of content? Wikipedia as a sociotechnical system', *New Media Society* 12(8), 1368–1387. https://doi.org/10.1177/1461444810365297

Retelny, D., Robaszkiewicz, S., To, A., Lasecki, W.S., Patel, J., Rahmati, N. et al., 2014, 'Expert crowdsourcing with flash teams', *UIST 2014 - Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*, Association for Computing Machinery, Inc., Honolulu, HI, October 05–08, 2014, pp. 75–85. https://doi.org/10.1145/2642918.2647409

Scannell, J.W., Blanckley, A., Boldon, H. & Warrington, B., 2012, 'Diagnosing the decline in pharmaceutical R&D efficiency', *Nature Reviews Drug Discovery* 11, 191–200. https://doi.org/10.1038/nrd3681

Sims, M.H., Bigham, J., Kautz, H. & Halterman, M.W., 2014, 'Crowdsourcing medical expertise in near real time', *Journal of Hospital Medicine* 9(7), 451–456. https://doi.org/10.1002/jhm.2204

Smith, A., 2003, *The wealth of nations (1776)*, Bantam Dell, New York.

Smith, V., 1962, 'An experimental study of competitive market behavior', *The Journal of Political Economy* 70(2), 111–137.

Vickrey, W.S., 1969, 'Congestion theory and transport investment', *The American Economic Review* 59(2), 251–260.

Von Hippel, E., 1994, '"Sticky Information" and the locus of problem solving: Implications for innovation', *Management Science* 40(4), 429–439. https://doi.org/10.1287/mnsc.40.4.429

Wasik, B., 2009, *And then there's this: How stories live and die in viral culture*, Viking, New York.

Weld, D.S., 1999, 'Recent advances in AI planning', *AI Magazine* 20(2), 93–123.

WHO, 2015, *Ageing and life course*, World Health Organization, viewed 25 April 2015, from http://www.who.int/ageing/en/

Woolley, A.W., Chabris, C.F., Pentland, A., Hashmi, N. & Malone, T.W., 2010, 'Evidence for a collective intelligence factor in the performance of human groups', *Science* 330(6004), 686–688. https://doi.org/10.1126/science.1193147

Yu, B., Willis, M., Sun, P. & Wang, J., 2013, 'Crowdsourcing participatory evaluation of medical pictograms using Amazon Mechanical Turk', *Journal of Medical Internet Research* 15(6), e108. https://doi.org/10.2196/jmir.2513

Zhai, Z., Sempolinski, P., Thain, D., Madey, G., Wei, D. & Kareem, A., 2011, 'Expert-citizen engineering: "Crowdsourcing" skilled citizens', *Dependable, Autonomic and Secure Computing (DASC), 2011 IEEE Ninth International Conference*, IEEE, Sydney, Australia, December 12–14, 2011, pp. 879–886. http://doi.org/10.1109/DASC.2011.148

Zook, M., Graham, M., Shelton, T. & Gorman, S., 2012, 'Volunteered geographic information and crowdsourcing disaster relief: A case study of the Haitian earthquake', *World Medical and Health Policy* 29(2), 7–33.