A new paradigm of knowledge management: Crowdsourcing as emergent research and development

C.W. Callaghan

ABSTRACT

Drawing from knowledge management theory, this paper argues that the knowledge aggregation problem poses a fundamental constraint to knowledge creation and innovation, and offers a potential solution to this problem. Specific consequences of innovation failure include the failure of research and development to deliver new medicines to address threats such as widespread and increasing antibiotic resistance, the rise of airborne multidrug-resistant or totally drug-resistant tuberculosis, as well as a lack of new drugs to deal with emerging threats such as Ebola. Persistent constraints to knowledge creation exist in the form of market failure, or the failure of profit-seeking models of innovation to internalise the positive externalities associated with innovations, as well as academic failure, or the failure of academic research to provide much-needed innovations to address societal problems. However, a lack of theory exists as to how to transcend these constraints to knowledge aggregation. This paper presents a probabilistic theoretical framework of innovation, suggesting that the 'wisdom of the crowd', or emergent properties of problem-solving, may emerge as a function of scale when crowdsourcing principles are applied to research and development. It is argued in this paper that the consequences of a lack of knowledge of innovation failure are already upon us, and that a radical new approach to knowledge management and innovation is needed.

Key words: probabilistic innovation, knowledge management, innovation, crowdsourcing, crowdsourced R&D

Prof. C.W. Callaghan is in the School of Economic and Business Sciences, University of the Witwatersrand. E-mail: chris.callaghan@wits.ac.za

For some time now, the ability of pharmaceutical firms to harness innovation and produce new drugs has been declining; globally, innovation failure is imminent (Horrobin 2000). The decline in the capacity of the research and development (R&D) flow of innovations, or the flow of knowledge needed to develop new medicines, is predicted to threaten human health outcomes (McKenna 2014). Some examples, a few amongst many, of the consequences of a 'nearly broken' system of the flow of innovations (McKenna 2014) are the failure of antimicrobial drugs in the face of widespread and increasing antibiotic resistance (Halifax 2013), the rise of airborne multidrug-resistant or totally drug resistant tuberculosis (Maher 2013), and the increasing resistance of organisms such as malaria to prophylactic drugs (Plowe, Djimde, Bouare, Doumbo & Wellems 1995). The failure of the flow of innovations also potentially renders global populations vulnerable to other emerging disease threats, including Ebola (Krishnan 2014), and the re-emergence of polio, as is the case currently in Nigeria (Ajumobi 2014) and Pakistan (Zia 2014). A range of different causes of innovation failure have long been suggested, including market failure, or the failure of profit-seeking models of innovation to internalise the positive externalities associated with innovations (Martin & Scott 2000) and academic failure, or the failure of academic research to provide much-needed innovations to address societal problems (Dewald, Thursby & Anderson 1986).

In the light of these mooted causes of innovation failure, the tension between the notion of exponential global growth in information and knowledge is also not reconciled in the literature (Adair & Vohra 2003), and neither is the innovation failure associated with a breakdown in the flow of pharmaceutical innovations (Horrobin 2000; McKenna 2014).

A comprehensive review of the longstanding literature seems to offer certain theoretical rationales for the breakdown of the flow of innovations that might, in turn, underlie market and academic failure as causes of innovation failure. These include the decentralised nature of knowledge itself (Hayek 1945) and the difficulties inherent in transferring it, or its stickiness (Von Hippel 1994), as well as its tacit nature (Nonaka 1994), in that tacit knowledge cannot be separated from the individual (Polanyi 1973). These constraints have been conceptualised as a 'threshold' that constrains innovation, which cannot be transcended under the current 'paradigm' of innovation (Callaghan 2014). This notion, of a threshold constraint to knowledge creation, stresses the lack of a 'transmission mechanism' that can transmit the steady exponential increase in volumes of information and knowledge over time (Adair & Vohra 2003) towards innovative pharmaceutical products to address societal problems such as, for example, antimicrobial resistance. It is the aim of this paper to present a theoretical framework of such a 'transmission mechanism', in the form

of a probabilistic framework of innovation. This model relates to how scientific knowledge creation systems can potentially be designed to produce an 'accelerated' process of innovation. Further, the framework draws from theory that suggests that the 'wisdom of the crowd', or emergent properties of problem-solving (Mitchell 1996; Johnson 1998; Kochugovindan & Vriend 1998; Smith 2003[1776]; Surowiecki 2004; Hanson 2000), may emerge as a function of scale in crowdsourced R&D.

This article therefore seeks to build on this notion, that the flow of innovations faces an 'upper limit' to its efficacy under the current paradigm of innovation and R&D. The proposed framework seeks to theorise process relationships around innovation and R&D under conditions that are relatively more robust to the constraints associated with the decentralisation or non-transferability of information and knowledge (Hayek 1945; Von Hippel 1994).

It is argued that this research is significant, for the following reasons. First, in light of evidence of the breakdown in the flow of innovations (Horrobin 2000; McKenna 2014), knowledge management as a field can be uniquely placed to theorise around societal problems that require a multidisciplinary approach. The reason for this is that societal knowledge problems are by their very nature multidisciplinary. By applying knowledge management theory as a lens, this paper offers a perspective that may be a useful way to frame approaches to societal problem-solving. Second, the consequences of a lack of knowledge of how to surmount the problems inherent in the nature of information and knowledge (Hayek 1945; Von Hippel 1994) may ultimately result in costs, both in the loss of human life and in resources, specifically in the case of innovation failure (Horrobin 2000; Maher 2013; McKenna 2014). Third, by providing a theoretical framework that reconceptualises innovation and knowledge management systems as either probabilistic or non-probabilistic in nature, this work offers what might be a structural break in the way innovation has been conceptualised. In the same way as mass-production was made possible by notions of specialisation, it is hoped that accelerated knowledge creation may be enabled by a wider acceptance of the notion of probabilistic innovation and R&D.

It is therefore argued in this paper that the consequences of a lack of knowledge of how to transcend the constraints to innovation failure are already upon us, and that a radical new approach to knowledge management and innovation is needed to enable a new paradigm in research that offers socially beneficial outcomes. The contribution of this paper is its development of a novel and new theoretical perspective of innovation.

The paper is structured as follows. First, the nature of innovation failure is discussed, in order to understand the problem that this article seeks to address, and to justify the scale and scope of constraints to knowledge creation that impact directly on the flow of pharmaceutical innovations. Having justified the problem, propositions

are derived and a theoretical framework of certain relationships is offered. The paper then concludes with a summary of the arguments, and their relevance for theory and practice is discussed. Recommendations for further research are also offered. The nature of innovation failure is now discussed.

The nature of innovation failure

The literature offers a host of different dimensions of innovation failure that may constrain the flow of innovations. For the purposes of this work, these are all considered to be forms of innovation failure. In the review that follows, certain core dimensions of innovation failure are distilled from the literature. The purpose of this process is to highlight the extent and scope of the challenges faced by innovation. One of the dominant themes in the innovation literature is the conflict between academic research and the demands of commercial industry, or differences in institutional logics.

Organisations draw from multiple institutional logics (Jay 2013), and academic institutions are no different. In these institutions, logics of teaching are differentially commensurate with logics of research and other logics imposed by a host of different stakeholders (Callaghan 2013). In relation to the tension between academic research and commercial demands, Murray (2010), for example, relates these perspectives to a landmark case that illustrates the tensions between different institutions and logics relating to the market versus academia at the nexus of medical research: the case of a genetically engineered mouse patented by Harvard for cancer research, which was subsequently exclusively licensed to DuPont, in which the response of the scientific community was also studied. At the heart of this tension is the differentially compatible logics of academic science, which seeks to pursue knowledge for its own sake, and of commercial science, which seeks to develop ideas into private property, primarily through the exclusion of others from this knowledge so that the rewards from its capture are appropriated (Murray 2010). This work argues that the conflict between these goals may constrain innovation output, which in turn has consequences for human life.

Horrobin (2000: 341) stresses that the present and future health of populations is dependent on pharmaceutical innovation, yet that "the evidence suggests that, despite apparent optimism, pharmaceutical innovation is failing". According to Horrobin (2000: 341), to sustain average industry growth, a firm needs to introduce at least one new drug product a year and to "sell around [300 million pounds] per year for every 1–1.5% it has of the world pharmaceutical market", implying that "the industry as a whole needs about 70–100 new products per year of this size". A company the size

of GlaxoWellcome/SmithKlineBeecham would therefore need about three to seven products every year, while one the size of AstraZeneca would need at least two to four products every year; for Horrobin (2000: 341), it is therefore clear that the "problem is that research productivity is failing".

Mergers are therefore increasing because of the failure of pharmaceutical companies to successfully innovate and produce sufficient numbers of new drugs (Horrobin 2000). This volume of sales is fundamentally related to the need for profitability, or a 'profitability threshold'. Horrobin's (2000) notion of financial thresholds that constrain the development of new drugs echoes the threshold constraint to innovation posited by Callaghan (2014), a constraint that persists despite the current exponential growth in information and knowledge. On the basis of this body of literature considered above, the following proposition is offered, Proposition a, that *commercial pharmaceutical innovation faces a threshold limitation that is inherently related to the commercial, or proprietary, nature of the R&D process.* The implication that derives from this proposition is that non-commercial, or non-profit/non-proprietary forms of R&D might offer some potential to transcend this threshold. This potentiality is discussed later in this paper. Another dimension of innovation failure that seems to dominate in the literature relates to another source of knowledge creation that stands as an alternative to commercial or proprietary (profit-seeking) research, namely academic research and 'academic failure'. Across different fields, and across time, the literature reveals a range of antecedents of academic failure. The term 'academic failure' is taken from Dewald et al.'s (1986) notion of market failures in academia where innovation does not occur, because of the failure of academic research to result in knowledge creation. Included as potential aspects of academic failure are failures of research to conform to quality standards, or not to be published, and therefore for society not to realise the positive externalities that would accrue from the advancement of science. What follows now is a brief review of the longstanding literature relating to academic failure across certain academic fields, in order to justify the scope of this constraint to innovation.

In attempting to replicate a range of economics studies, Dewald et al. (1986: 587) found inadvertent errors in published empirical articles to be a "commonplace rather than a rare occurrence". The "marketplace of economic research might be expected to provide a check" against these quality concerns, yet this marketplace seems to have failed at this (Dewald et al. 1986: 589). Nevertheless, Dewald et al. (1986: 589) highlight these problems in the context of what they regard as classic examples of market failure, as the improved quality of articles is taken to share characteristics of public goods. Dewald et al. (1986: 600) also stress the particular importance of a host of other ongoing problems with research output, not least of these being collinearity

of data and correlations between coefficient estimators, where "even slight differences in data values or in the numerical precision of computer programs may produce sharply different parameter estimates".

Further, Dewald et al. (1986) stress that the incentive structure in academia does not encourage replications of existing studies. Many studies follow a path-dependent route to their completion, with multiple collaborations in the data analysis processes, including the use of graduate students to produce analytic and other portions of the research (Dewald et al. 1986). Constraints to innovative research and to the progress of science itself might, however, exist within the processes by which academic research is published and disseminated (Bornmann 2010; 2011).

Academic research in certain instances has been found to be associated with biases, including in reported uncertainties, in which actual errors are typically underestimated, as well as "persistent overconfidence...[as in the case of]... psychological research on the assessment of subjective probability distribution" (Henrion & Fischhoff 1986: 791). Nevertheless, it is in the realm of theory testing that these biases may dominate. A concern has long existed across many scientific disciplines about the processes by which science is disseminated, in that theoretical innovation can be constrained (Peters & Ceci 1982).

In relation to research on the academic knowledge dissemination process, Mahoney (1977: 161) found the "tendency for humans to seek out, attend to, and sometime embellish experiences that support or 'confirm' their beliefs", which led to a bias against the dissemination of new perspectives. For Mahoney (1977), this is fundamentally at odds with Popper's (1972) notions of falsification. Despite the challenges faced by knowledge creation that relate to the inefficiencies of the academic publication process, the literature indicates that there are serious issues related to the resistance of academics themselves to scientific discovery. Campanario (2009: 550) argues that the "important topic of scientists' resistance to scientific discovery" is under-researched.

Some of the highest cited papers, as well as certain influential books in academia, were first rejected by journal reviewers and editors, including work that eventually was awarded the Nobel Prize in the fields of Physics, Chemistry, Physiology and Medicine (Campanario 2009). Potential explanations for the rejection of innovative articles are that they do not conform to the orthodox viewpoints of reviewers, or do not fit with accepted paradigms (Campanario 2009). "This outcome of peer review raises important questions about current publishing policies which govern the dissemination of new information", according to Campanario (2009: 558). Further, according to Campanario (2009: 559), "critics often argue that peer review operates to regulate paradigmatic science (in the Kuhnian sense) rather than to welcome brand new knowledge", and there is a "real risk that evidence contrary to the established

views can be suppressed or discarded". Campanario (2009: 560) argues further that his case studies show that "scientists with something truly original to communicate often have to fight against the silence, the lack of interest, and as a result the absence of citations and recognition".

In the face of threats like the failure of antibiotics to treat emerging 'superbugs' in our hospitals, which can make simple injuries or surgery life-threatening (Halifax 2013), or the rise of pandemics like Ebola (Krishnan 2014), the costs of non-publication of ideas offering different paradigmatic perspectives may be catastrophic. Of all articles that contribute new ideas, it is possible that some might be able to make a difference, and contribute to a strengthened flow of innovations.

Another form of bias related to that of non-publication bias due to non-conformity with accepted paradigms is bias related to the non-publication of null findings. Publication bias also occurs in the form of the 'file drawer problem', where articles that report significant findings are more likely to be published, thus contributing to a false confidence in the rejection of the null hypothesis (particularly in meta-analyses and literature reviews), as studies failing to reject the null hypothesis literally remain in 'file drawers' (De Long & Lang 1992).

Nevertheless, De Long and Lang (1992: 1267) argue that "failure to report insignificant test statistics is in fact quite rare in the published literature...We find no evidence of this sort of bias". However, De Long and Lang (1992: 1268) acknowledge that "it is very possible that the bias occurs at an even earlier stage", in that "when the data strongly support the null, the paper is less likely to be written and, if written, is unlikely to be published". They also stress that all hypotheses tests are "mere approximations" (De Long & Lang 1992: 1269). For De Long and Lang (1992: 1270), authors therefore "face a catch-22: papers that fail to reject their central null hypothesis will be published only when editors think they are especially interesting, but editors will think that they are especially interesting only when the null hypothesis that they test really is false". Further, most "of us suspect that most empirical researchers engage consciously or unconsciously in data mining...[there] seems to be no practical way of establishing correct standard errors when researchers have prior knowledge of the data, or when they report only their favourite results: the distribution of the 10 highest t-statistics is not well known" (De Long & Lang 1992: 1270). Derived from the literature considered above, proposition b is offered, that *academic innovation faces a threshold limitation that is inherently related to the academic nature of the R&D process.*

It is not the intention here to exhaustively review the literature relating to innovation failure along the dimensions of market failure associated with the constraints posed by profit-seeking frameworks of innovation and academic research frameworks of innovation. Instead, these examples discussed above are considered

to be useful indicators of a problem faced by the knowledge creation processes. These and other challenges are considered to potentially act in a cumulative way to constrain innovation processes, and to contribute to innovation failure. Market failure and academic failure might be taken to contribute to a threshold constraint, where another perspective of knowledge management is needed in order to transcend it. The contribution of this article is to offer such a novel perspective, in the form of 'probabilistic innovation'.

A new paradigm of probabilistic innovation

As discussed previously, the central thesis of this paper is that the flow of pharmaceutical innovations is constrained by a host of mechanisms that have resulted in a decrease in the production of medicines, and that problems such as the increase in microbial resistance (as one example amongst many) is not being met by a corresponding increase in pharmaceutical innovation. This serious problem is framed as a challenge associated with the inherent limitations of the systems of knowledge creation that underpin pharmaceutical innovation. These systems are considered to reflect non-probabilistic processes of innovation, typically associated with linear, or non-exponential, incremental innovation outputs. The core argument of this paper follows that of Callaghan (2014), that non-probabilistic processes of innovation, associated with the paradigm of first-generation innovation, cannot deliver the innovations required by society because they are fundamentally constrained in their ability to deliver outputs beneath the level of a 'probabilistic threshold' (Figure 1).
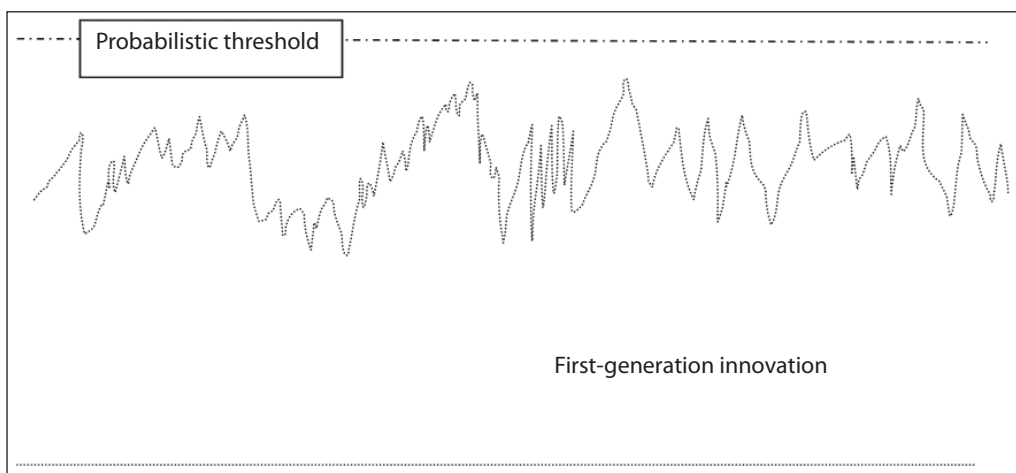

Figure 1: The probabilistic threshold constraints posed by the 'stickiness' of information
Source: Callaghan (2014: 167)

The probabilistic R&D space: Second-generation innovation

As discussed, the core argument made in this work is that the decentralised nature of knowledge (Hayek 1945) is a fundamental constraint to any system of knowledge creation that is not probabilistic, or that is able to tap into the information and knowledge of large numbers of people, spread across geographical distances. At the heart of this notion of the dispersed nature of knowledge is that knowledge is sometimes contradictory and possessed by different individuals, a condition described by Hayek (1945) as the problem of asymmetric information, a dominant problem faced by social systems, and therefore human problem-solving systems. The implications that derive from this problem are multiple. A core implication, however, is that centralised systems typically do not aggregate and distribute relevant information as effectively as decentralised systems do (Hanson 2000). The paradigm of first generation innovation is taken to represent innovation systems that do not use a probabilistic system of knowledge creation. Callaghan's (2014) 'paradigm' of second generation innovation (SGR) relates to systems of knowledge creation that harness probabilistic mechanisms to increase the probability of innovative knowledge creation. In other words, SGR utilises recent advances in information technology to allow large numbers of people to simultaneously focus on solving problems, utilising virtual networks. SGR is underpinned by a new and rapidly emerging body of literature that relates to the potential of the 'crowd' to solve problems previously taken to be inaccessible. The logic behind probabilistic innovation is that if the number of human problem-solvers focusing on a problem is increased exponentially, then the probability of finding a solution to the problem will also increase exponentially. On the basis of this logic, proposition c is derived, that *non-probabilistic systems of innovation are less likely to transcend the endogenous limitations of constraints to innovation than are probabilistic systems*. Callaghan (2014) relates this probabilistic paradigm to the rise of crowdsourcing and crowdsourced R&D, arguing that the combination of these offer new opportunities for innovation that require considerably fewer resources than first generation innovation problem-solving. He also argues that proprietary or profit-seeking crowdsourced R&D is fundamentally constrained by dyadic linkages between knowledge provider and seeker, because of the need to keep information and knowledge secret so that it can produce downstream returns (profit). He argues that if the dyadic nature of crowdsourced knowledge creation is 'collapsed', and if knowledge and information is fed back into the crowd, a 'three-dimensional' process of accelerated knowledge creation can be enabled, which might significantly increase the rate of innovation (potentially exponentially if the volumes of problem-solving contributions can be increased exponentially). However, what Callaghan's (2014) work does not

sufficiently contribute is a theoretical framework for why this new 'paradigm' is new or novel, or indeed, might have the potential that is suggested.

This article, therefore, seeks to explicate the theoretical foundations of the argument that SGR does indeed represent a new paradigm in knowledge management. To justify this argument, it is necessary to develop a theoretical framework that is robust, and that can support the implications that derive from this argument. This is undertaken as follows, by providing examples of other instances that support these arguments. The first of these relates to the problem of asymmetric information and the ability of the 'crowd', when harnessed in the form of a market, to solve problems of information and knowledge aggregation.

This problem of asymmetric information has been offered as a reason for differences in performance across a host of different contexts. For Hanson (2000: 9), large variations exist in the economic growth rates of nations, which to some extent can be ascribed to the information failures suffered by political institutions. An "important fraction of the variation, however, seems attributable to some nations adopting policies which relevant experts knew to be bad...Thus at some level bad policy seems to be fundamentally akin to a failure to aggregate and distribute relevant information" (Hanson 2000: 9). This example has theoretical relevance to the notion that the 'crowd' can be an important source of problem-solving input.

What these arguments (Hayek 1945; Von Hippel 1994; Hanson 2000; Callaghan 2014) seem to have in common is the notion that the *effective* aggregation and distribution of information, or knowledge, holds the promise of dramatic improvements to economic systems and other systems that seek to achieve objectives. For Hanson (2000), knowledge 'engines' in the form of speculative markets can be harnessed to provide [distribute] knowledge and information [that is effectively and efficiently aggregated], in some cases providing even better probabilistically-based predictions of future events than other forms of polling. In other words, according to Hanson (2000), speculative markets, which can be set up as 'decision engines', *can be more effective at solving problems than experts are.*

Similarly, another related perspective has emerged, that the problem-solving capacity of the crowd can generate exponential increases in knowledge creation, ushering in a new paradigm of 'probabilistic' innovation (as opposed to linear models of innovation output) (Callaghan 2014). In short, these arguments are akin to Surowiecki's (2004) notion of the 'Wisdom of Crowds', where collective knowledge creation can equal if not improve on expert predictions of phenomena. However, in considering this body of work, a tension seems to be present in the literature.

On the one hand, these theorists argue that problem-solving can be dramatically improved through the use of 'crowd' based mechanisms, whether in the form of

speculative markets (Hanson 2000) of probabilistic research and development (R&D) (Callaghan 2014), or of the contribution of large numbers of people, or crowds (Surowiecki 2004). On the other hand, it can be argued that these ideas are not new, and that if they were as effective as mooted, then why have they not made a dramatic impact in the world today? For example, according to a certain perspective (Callaghan 2014), many societal problems could be solved almost immediately, if the problem-solving potential of the crowd could be harnessed sufficiently, or the probability of being able to solve a problem does rise according to how many people contribute to the problem-solving effort. Simply put, this argument posits that an exponential increase in the number of problem-solvers would, ceteris paribus, be expected to result in an exponential increase in problem-solving output. However, this has, to date, not happened, at least in terms of the failure of current innovation systems (McKenna 2014), notwithstanding the longstanding use of crowdsourced R&D – as in the case of InnoCentive and similar sites (Howe 2006; Lakhani & Panetta 2007) – albeit under proprietary conditions.

Indeed, such conceptions seem intuitively simplistic, or perhaps naïve, ascribing far too much variance in the problem-solving capacity of systems perhaps to the role of asymmetric information (Hayek 1945), or the stickiness of information (Von Hippel 1994). It is this tension that this paper seeks to resolve, by offering an argument that dramatic improvements in problem-solving are possible if it is indeed possible to 'harness' the latent power of the crowd. However, the primary contribution of this article is its development of a theoretical framework of crowdsourced R&D that is inherently 'social' in nature, or that is sustainable without being underpinned by the profit mechanism (profit-seeking not as a necessary condition). The decentralised nature of knowledge has multiple implications, and the thread of logic common to these implications seems to be the notion that mechanisms that in some way harness the information and knowledge contributions of the 'crowd' can yield powerful problem-solving solutions. On the basis of this discussion, proposition d is offered, that *rates of problem-solving and levels of innovative outputs of probabilistic innovation systems are positively related to volumes of information and knowledge inputs, after a threshold point, or critical mass, is reached*. In this proposition, the threshold point, or point at which critical mass is achieved, relates to the need for large numbers of inputs. It is argued that the quality of inputs, in turn, will be a function of the volume of inputs, but these will need to be of a sufficiently large volume to provide a normal distribution of quality inputs.

Whereas the decentralisation versus centralisation debate has attracted much attention from scholars, it has been found that unregulated competition is more effective than central planning in economic terms, the latter a conception grounded

in polar opposites and in a heuristic bias (Hanson 2000). Hanson (2000) stresses the existence of a heuristic bias identified by Hayek (1973), where people typically have a biased preference for central direction. Betting markets, for Hanson (2000), can be used to harness information because they are robust to many disadvantages that markets have, as speculative markets are typically highly efficient at providing information; in other words, it is difficult to find information that has not been taken into account by these prices. Speculative markets seem to be an example of a mechanism that can efficiently harness the information and knowledge inputs of the 'crowd'.

Theory relating to the way crowds provide information in the case of speculative markets

A betting market is a speculative market, where assets are traded in a way that enables people to bet on issues of fact, or final judgements about facts (Hanson 2000). For Hanson (2000), as Hayek (1945) predicted, speculative markets are very good at aggregating information, even when those that are trading have little knowledge of other traders or even of their environment. The mechanism underlying speculative markets that aggregates information can be harnessed for other ends.

Hanson (2000) argues that betting markets can create knowledge more efficiently than academic institutions, or the academic process, and gives the example of the failure of academic research to provide effective evaluations of political institutions. Hanson (2000) offers the idea of 'futarchy', a method of political governance in which voters vote for political outcomes in a process modelled on a betting market. For Hanson (1990: np), it is not only political systems that can benefit from modelling the characteristics of speculative markets, but academia too: "academia is still largely a medieval guild...Peer review is just another popularity contest, inducing familiar political games; savvy players criticise outsiders, praise insiders, follow the fashions insiders indicate, and avoid subjects between or outside the familiar subjects". Although perhaps a little strong in his attack on academic failure, Hanson's (1990) argument is clear: not enough incentives in academic research exist to avoid bias. Hanson (1990: np) suggests that perhaps "the core problem is that academics are rewarded mainly for telling a good story, rather than for being right", and that academic incentives typically "reward being popular, fashionable, and eloquent, instead of being right".

As a solution to the problem of knowledge aggregation, Hanson (1990) takes recourse to the notion that knowledge management mechanisms that can aggregate information can transcend constraints to knowledge creation, and offers a

conceptualisation based on 'idea futures', or bets on ideas by academics, and argues that in such a market for ideas the incentives would be better aligned with scientific progress. Hanson (1990) gives examples of academics whose ideas later turned out to be right, but who were vilified for their ideas at the time they were put forward.

According to Hanson (1990), if odds were offered on these ideas, then no matter what odds were offered initially, these would be bid down or up to reach convergence; in other words, everyone involved would be incentivised to be careful and honest because they have a stake in the outcome, either way, right or wrong. The market would incentivise more research, and more knowledge creation, as new information would be rewarded. Hanson (1990) offers the examples of Wegener, who predicted the existence of continental drift, and the case of the discovery of the mass of the neutrino; in both instances these findings were initially ridiculed by the academic community. Following Hanson (1990), science bets could be used to insure against technological risk, and prices would offer valuable information about which bets to back. However, Hanson (1990) argues that the dominant constraints to the use of these mechanisms are legislative provisions that typically constrain the emergence of science betting and other uses of the speculative market mechanism to solve societal problems.

Nevertheless, Hanson's (1990) central thesis is clear, that markets have an ability to aggregate information, and that 'information markets' have been successfully used to aggregate information on topics of interest, as a host of examples attest to the ability of these markets to outperform experts. However, there are certain conditions that need to be met in order to enable this form of 'crowd-driven' problem-solving.

Information markets can only function effectively if certain conditions are avoided. For example, (i) the thin market problem, where traders need to coordinate their activities with regard to the assets they need to trade and when they will trade, but do not have a large enough market to do so; (ii) the irrational participation problem, which exists when rational agents have hedged their positions and yet continue to trade; and (iii) the thick market problem, when single consensus estimates cannot be produced due to the presence of too many different estimates (Hanson 2003). Another problem emerges when traders make bets that they gain from if they win, but that cannot be paid if they lose, termed 'money pumps' (Hanson 2003). Nevertheless, if these and other conditions are met, information markets can be effective information aggregation mechanisms if they require people to trade assets in relation to the possibility that certain events occur; diverse opinions can be combined into a single probability distribution on the basis of convergence (Hanson 2003). Hanson's (2003) work is another example of a body of theory that supports the arguments made in this paper, that at the heart of innovation failure is the failure of

the mechanisms used in academic research and in commercial R&D to maximally aggregate information and knowledge, a longstanding problem that has been the focus of seminal attention over time (Hayek 1945; Polanyi 1973; Von Hippel 1994). Using a range of different theoretical propositions, it is hoped that according to the principles of convergent and discriminant validity (Campbell & Fiske 1959), different bodies of theory provide cumulative support for the central tenets of this paper. Yet another theoretical perspective that may lend support for these arguments is provided by evolutionary theory, as it is applied to knowledge.

Insights from evolutionary theory that explain how the crowd can solve problems

Evolution, as a concept, has proved a useful optimisation tool for engineering problems, as populations of candidate solutions to problems are used to improve problem-solving (Mitchell 1996). Genetic algorithms are developed to apply principles underlying biological evolution to computer systems and to develop theory relating to adaptation based on crossover, inversion and mutation (Mitchell 1996). This approach seems to be particularly well suited to fields in which a large number of potential solutions exist, such as computational protein engineering, and where the potential for computational parallelism exists (Mitchell 1996).

Artificial intelligence research has shown that in contrast to top-down complex processes and programmes that sought to take into account detailed behaviours, simpler bottom-up processes have been relatively more successful, based on "only very simple rules, and complex behaviours such as intelligence emerge from the massively parallel application and interaction of those simple rules" (Mitchell 1996: 4). This research echoes other literature that over time has highlighted the potential for *emergent* forms of problem-solving (Forrest 1990; Johnson 1998; Kochugovindan & Vriend 1998; Smith 2003[1776]; Surowiecki 2004).

This process, as explained above by Mitchell (1996: 5), has aspects in common with evolution, as evolution is "in effect, a method for searching among an enormous number of possibilities for 'solutions'", or "a method for designing innovative solutions to complex problems", and is therefore underpinned by evolution theory. The causal mechanisms that underlie the evolutionary process of problem-solving seem to be analogous to the way emergent processes of crowdsourced R&D may reflect a search among a high number of possibilities of which the most successful 'win out'.

Searching for candidate solutions in computer science has been termed searching in a 'search space', which represents candidate solutions and a kind of distance between them (Mitchell 1996). A dominant problem in computational bioengineering

is computational protein design; for instance, a computational search for a protein, or a sequence of amino acids, that 'folds up' to a particular three-dimensional shape in order to solve the problem posed by a specific virus (Mitchell 1996: 7). The search space in this instance is therefore the collection of all possible protein sequences, which is essentially an infinite set of possibilities, or permutations of possibilities (Mitchell 1996).

The fitness landscape is the space of all possible outcomes as well as their fitness, which can be represented as 'hills', 'peaks' and 'valleys'; evolution is a process that 'moves' populations around this landscape (this computational representation differs from real-world biology, as the biological landscape cannot be separated from the population) (Mitchell 1996). It seems possible that crowdsourced R&D can benefit significantly from these ideas, as they provide a way to differentiate between solutions provided by the crowd. Essentially, using the crowd to provide very large numbers of potential solutions to societal problems such as Ebola, HIV or antibiotic resistance might replicate the mechanisms of evolutionary selection, in this case the selection of ideas.

Evolutionary theory can perhaps be taken to represent one strand of a range of different strands of theory that provide a 'bedrock' of theory that can underpin the theoretical and practical phenomenon of crowd-sourced R&D, which in turn is representative of a new paradigm of probabilistic innovation. At this nexus, what remains to be considered is theory that relates the different levels of analysis. The individual is the unit of analysis at the point at which information and knowledge is provided in the crowdsourced R&D process. However, the behaviour that emerges at the level of the crowdsourced R&D system has certain characteristics that are fundamentally different from the behaviours of the individuals involved. Certain of these behaviours are considered to be 'emergent'. It is this notion that we turn our attention to now.

The tension between the individual level and the level of the problem-solving system

Theory relating to the contributions of high numbers of individuals to a system of knowledge aggregation requires solid underpinnings in terms of knowledge of emergent effects that are not inherently related to the behaviour of individuals within such systems. In his work, Johnson (1998: 3) sought to "reconcile the traditional approaches of collective problem-solving involving cooperation and competition of globally 'aware' individuals and…examples of global problems being solved without awareness of the individual". According to Johnson, a body of literature has

developed to address this issue, particularly in the fields of biology, economics and sociology (Johnson 1998).

Insights from evolutionary biology and cognitive science suggest that even in simulations of herd formation using aggressive agents, collective cooperation behaviour is observed; a global structure that "emerges from the dynamics and special extent of the system" (Johnson 1998: 3). The characteristics of this system stand in stark contrast to approaches that assume that the locus of choice of whether to cooperate or not rests with the individual; a flaw of the game-theoretical analysis is associated with this assumption (Johnson 1998). According to this assumption, the potential exists within the individual for observed global behaviour (Johnson 1998). Further, according to Johnson (1998), this assumption is questionable because observed global behaviour cannot be assumed to be a function of individual agency. This argument has important implications for probabilistic innovation, and crowdsourced R&D, as emergent systemic behaviours may emerge that are different from the behaviours of individual contributors.

Therefore, for Johnson (1998: 4), "emergent functionality (a global property that cannot be predicted or observed at a lower level) can be achieved without embodiment of the functionality in the individual...[and]...can include, not only cooperation...but also abstract concepts of problem-solving beyond the perception of the individual". Hence, "complex global systems are a consequence of the dynamics among relatively simple agents", as posited in the fields of complexity and complex adaptive systems (Johnson 1998: 4). If abstract concepts of problem-solving that are beyond the perception of the individual can emerge from systems of probabilistic innovation, then it is possible that there are mechanisms that can be captured. It is therefore argued here that second-generation innovation theory, and its practical application in the form of probabilistic innovation, can be related to predictions derived from theory from the fields of complexity and complex adaptive systems (Johnson 1998). Having provided a perspective of theory that underlies certain of the relationships of probabilistic innovation, and specifically, crowdsourced R&D, the notion of emergent problem-solving is now extended to consider the phenomenon described as 'collective intelligence'(Surowiecki 2004).

Collective intelligence: The power of the crowd to solve (or create) problems

A theory of crowdsourced problem-solving needs to build upon a theory of problem-solving. Simon (1971: 145) stresses that a theory of problem-solving should, firstly, "predict the performance of a problem-solver handling specified tasks" and should

also explain how problem-solving takes place, or its processes and mechanisms. Secondly, it should predict the contextual relationships of the problem-solving, and the relation of these to the problem-solving process, as well as showing changes in the attendant conditions – "both changes 'inside' the problem-solver and changes in the task" alter problem-solving behaviour (Simon 1971: 145).

The mechanisms underlying the collective intelligence of markets and other forms of collective problem-solving may offer hitherto undiscovered potentialities for solving societal problems and addressing innovation failure. Historical examples exist of the ability of aggregated problem-solving, or collective intelligence, to solve complex problems (Surowiecki 2004). Examples include the almost exact estimation of the weight of an ox in a town market and the prediction of the location of the wreck of a submarine (Surowiecki 2004).

For Johnson (1998: 4), emergent functionality as a concept offers important insights into the behaviours of distributed self-organising systems. A shift has occurred, from an interest in understanding distributed self-organising systems towards an interest in the processes associated with creating these systems, partly as a result of the global shift of societies from a manufacturing focus towards a knowledge focus (Johnson 1998). In a simulated process of problem-solving using a distributed self-organising system, Johnson (1998) found the "system to exhibit chaotic dynamics at the level of the individual, but stable solutions at the level of the emergent property", in which diversity provided better collective problem-solving and solutions that were more robust.

According to Johnson (1998: 37), self-organising systems occur in nature and social systems "when the global system is too complex or the centralised problem-solver is lacking in capability or control over the system". Due to the increasing complexity of our world, and the way change occurs "faster than we can evaluate the changes, let alone respond to them, the process of collective self-organisation may be the only option", yet in time capabilities and policies will develop to support these systems (Johnson 1998: 37).

In order to understand the way 'collective self-organisation' (Johnson 1998) can contribute to the benefit of all, it is important to understand another dimension of the theory that underlies the operation of crowdsourced R&D, namely the operations of crowds themselves. A rich and seminal literature offers certain tensions that relate to the positive versus the negative influences of crowds. For Le Bon (1896), the crowd acts in direct conflict to Darwin's law of evolution, in that it typically can become atavistic or regressive. Although Le Bon (1896) relates his discussion to physical groups of people that come together to form crowds, this body of literature is considered to be important, as a theory of crowdsourced R&D needs to specify its

boundary conditions for its effectiveness, and a study of the historical literature on crowds can offer important insights into the different conditions of crowd behaviour. Le Bon (1896: np) argues that the arguments that crowds "employ and those which are capable of influencing them are, from a logical point of view, of such an inferior kind that it is only by way of analogy that they can be described as reasoning".

What is clearly evident from Le Bon's (1896) perspective is that the collective intelligence of crowds is limited and inimical to reason. In contrast with Le Bon's (1896) thesis, Surowiecki (2004: xv11) argues that if you put together "a big enough and diverse enough group of people" and ask them to "make decisions affecting matters of general interest", this group's decisions will, over time, be "intellectually [superior] to the isolated individual", notwithstanding the intelligence and knowledge of the single individual.

Surowiecki (2004) argues that complexity does not bar problems from being solved by collective intelligence, and offers three categories of problems amenable to problem-solving based on collective intelligence, namely (i) cognition problems, which might not have 'right answers' but will typically have answers that can be ranked, some being better than others; (ii) coordination problems, which require people to coordinate their behaviour, such as in the case of markets, or matching buyers and sellers; and (iii) cooperation problems, which require cooperation between people when self-interest, lack of trust or other incentives might not support such cooperation (examples of this include tax payments, the management of pollution or even achieving commonality on the definition of reasonable remuneration). However, for a crowd to be wise, *certain prerequisites exist, namely the need for it to be diverse, to have independence, and to be decentralised* (Surowiecki 2004). For Surowiecki (2004: xix), under certain conditions the crowd can act in ways identified by Le Bon, and therefore an understanding of the different conditions under which crowds operate is important, as groups benefit from communication, but too much communication, paradoxically, "can actually make the group as a whole less intelligent..." Surowiecki (2004: xix) offers examples of riots or stock-market bubbles, when "aggregating individual decisions produces a collective decision that is utterly irrational... underscoring the importance to good decision making of diversity and independence by demonstrating what happens when they are missing".

For Surowiecki (2004: xix), "[d]iversity and independence are important because the best collective decisions are the product of disagreement and contest, not consensus and compromise". An intelligent group "does not ask its members to modify their positions to let the group reach a decision everyone can be happy with...[i]nstead, they figure out how to use mechanisms – like market prices, or intelligent voting systems – to aggregate and produce collective judgements that represent not what

any one person in the group thinks but rather, in some sense, what they all think", particularly in relation to cognition problems (Surowiecki 2004: xix).

"Paradoxically, the best way for a group to be smart is for each person in it to think and act as independently as possible" (Surowiecki 2004: xx). Surowiecki (2004) offers a synthesis of the literatures, and makes the argument that if certain conditions are met, the crowd can offer solutions to problems that are, paradoxically, more likely to be correct than those offered by experts in many instances. This notion, that crowd-based solutions, or the information or knowledge gained by aggregating the responses of many different people, can offer some sort of 'convergence' with an underlying truth that is being sought, has also emerged from empirical research, most notably in the case of research conducted by Johnson (1998).

Johnson (1998) applied a simple agent-based model to solve a problem in the form of a maze. On the basis of this process, he found that the aggregated responses of agents demonstrated convergence to an optimal path through the maze. On the basis of these findings, Johnson (1998) offers the following as characteristics of cooperating collectives, namely that collectives can (i) demonstrate greater performance than individuals, (ii) face performance decline with decreasing diversity, and (iii) can perform even when exposed to extreme 'noise' and information loss.

Johnson (1998) therefore argues that emergent problem-solving offers important lessons for solving problems. He gives the example of how a new building development will best develop its pathways from the way the 'system' allows them to emerge from the collective action of many "individuals solving their own path problem, in a manner that is ultimately useful to the entire population but which is never expressed as a goal at the level of the individual" (Johnson 1998: 3).

This type of problem-solving is perhaps already reflected in the way the academic enterprise solves problems (certain academic research outcomes seem to be emergent in nature), but it is argued here that this is so only at the macro-level, and that what seems to be lacking is a micro-level 'synthesis' of the problem-solving-scape so as to create a context that makes problems more susceptible to emergent forms of problem-solving. In other words, a problem has to be formulated in a 'space' similar to the building development space, in which all that remains is to allow the emergent problem-solving to develop its own path, but using a process that harnesses the information and knowledge inputs of large numbers of people and focuses these inputs on a particular problem. It is argued here that a structure needs to be created to allow the 'wisdom of the crowd' (Surowiecki 2004) to solve problems such as Ebola, antibiotic resistance and others.

What is clear from the notion of collective emergent problem-solving is that some degree of spontaneous organisation occurs that is not directly related to the

individual level. Johnson (1998: 3) stresses that such a form of emergent problem-solving is associated with mechanisms that reinforce emerging patterns, but he argues that problem-solving can occur even without the presence of these reinforcing mechanisms.

Johnson (1998: 3) offers a more specific analysis of collective action from the perspective of evolutionary biologists and cognitive scientists. Simulations of herd formation including only aggressive agents (with "no inherent mechanism for cooperation *embodied* in the individual") have been found to reveal collective behaviour at the group level (Johnson 1998: 3). Cooperative behaviour from essentially uncooperative individuals is a "global structure that emerges from the dynamics and special extent of the system". In contrast, traditional approaches to modelling this type of cooperative behaviour typically assume that agents themselves embody the choice to cooperate or not, which effectively reduces, incorrectly, to game theoretical analysis (Johnson 1998). The explanatory value of traditional approaches to cooperation in animals and humans "is rightly questioned when the same behaviour can be observed without assuming it exists at the level of the individual" (Johnson 1998: 3).

An example of this prediction, of a perspective that takes complex global systems to be a function of the dynamics of relatively simple agents, is the notion of a self-organising social/economic system (the 'invisible hand') posited by Adam Smith (Johnson 1998). According to Smith (2003[1776]), markets operate in a way that allows them to solve problems such as allocation problems through an emergent mechanism. Certain implications derive from this notion (of emergent behaviours that are fundamentally different from the individual behaviours and choices of individuals). The attempt to develop microeconomic underpinnings to explain macroeconomic forces might not always be successful, if the dynamics of complex systems are not always derived from the properties of the individual (Johnson 1998; Kochugovindan & Vriend 1998).

The way emergent systems can solve problems, however, may be more complex than the perspectives associated with beliefs in simple market mechanisms. According to the invisible hand hypothesis advanced by Smith, individuals act out of simple self-interest, but this self-interested behaviour at the individual level "results in a coordinated overall outcome"; this is an outcome of the way complex adaptive systems operate (Kochugovindan & Vriend 1998: 54). In other words, global properties cannot be derived only from a study of individual components (Kochugovindan & Vriend 1998). A "complex system...consisting of a large number of relatively independent parts that are interconnected and interactive...is adaptive if the parts are agents that change their actions as a result of events occurring in the process of interaction";

examples of complex systems include "biological systems, immunology systems, brains, weather systems, ecologies, and societies" (Kochugovindan & Vriend 1998: 55). A crowdsourced R&D 'problem space' might share certain of the characteristics of complex systems, certain of which might be uniquely suited to effective problem-solving. However, such a system would need to preserve a sense of order, and not drift into chaos. According to Kochugovindan and Vriend (1998: 66), a complex system is not the same as a chaotic system; they "tend to evolve away from the extremes of, on the one side, absolute order and, on the other side, what appears to be complete randomness". The key theoretical concepts that relate to complex systems are "self-organisation (the formation of regularities in the patterns of interaction) and selection (through system constraints)" (Kochugovindan & Vriend 1998: 66). Selection enables self-organising systems to constantly push back to "some boundary between order and chaos", and "around this edge, these systems appear to carry out the most complex behaviour and adapt most readily to changing environments" (Kochugovindan & Vriend 1998: 66).

In general, "markets emerge as a result of locally interacting individual agents' pursuit of advantageous contacts; that is, they are self-organised"; and transactions do not therefore take place in Walrasian central markets (Kochugovindan & Vriend 1998: 66).

While individual agents adapt to the environment, aspects of the environment will adapt to the individual; a process of coevolution occurs (Kochugovindan & Vriend 1998). Applied for example to economics, therefore, "rather than analysing whether an equilibrium exists for an economy with some given structure, in this approach one analyses how structures and patterns emerge as regularities in the process of interaction of the individual agents"; it is therefore not "nineteenth-century physics but modern biology or meteorology [which] provides the relevant metaphors for this approach to the study of decentralised economies [economies that are not command economies but free-market oriented]" (Kochugovindan & Vriend 1998: 57). Patterns and regularities are identified in data that are not a function of the properties of individual units of the system (Kochugovindan & Vriend 1998). This notion is fundamentally at odds with the approach used by many in economics, for example, where microeconomic foundations are considered to fully underpin macroeconomic phenomena, an approach that is reductionist and top-down (Kochugovindan & Vriend 1998). What this body of literature suggests is that a complex system of problem-solving can be engineered, and that there might be some underlying forces in such a system of problem-solving that may result in 'emergent' forms of problem-solving. For the purposes of this paper, these forces are considered akin to some form of 'collective intelligence'. It is argued that unless these

systems are developed and tested, the consequences of a lack of knowledge about emergent 'crowd-based' complex systems of problem-solving will persist in the form of malfunctioning innovation flows and innovation processes that underperform performance thresholds. In other words, it is argued here that it is possible to harness the same forces for probabilistic problem-solving that underlie 'the invisible hand' hypothesis or that underpin financial markets or markets themselves. At the heart of the reason for this is that information and knowledge are often tacit and cannot easily be separated from individuals (Polanyi 1973), but to some extent this tacit knowledge might be captured by providing a platform for high numbers of individuals to be 'present' in a collective process of problem-solving that transcends geographical and organisational boundaries between contributors.

The application of notions of complex adaptive systems to financial markets is another example of emergent problem-solving, as the information and knowledge inputs of relatively large numbers of contributors solve the problem of signalling through the formation of market prices, although this process is not perfect. There has typically been a tension in the literature in this respect, as standard theory predicts that investors are identical, sharing rational expectations of the future prices of assets and using all market information to adjust their expectations (Kochugovindan & Vriend 1998). The implication of standard theory is that speculative profits cannot be earned in a sustainable way, and technical trading has no real scope, as volatility is not related across periods; in other words the market is rational and efficient (Kochugovindan & Vriend 1998). However, many financial market traders believe that speculative opportunities do exist, and that technical trading can be effective, while a market psychology holds sway over investors, and market crashes and bubbles can occur independently of market information; and further, that serial correlation is present (Kochugovindan & Vriend 1998). The interactions between individual-level effects and systemic relationships in the context of problem-solving have also been studied using agent-based modelling.

Agent-based modelling is a computational method that attempts to 'map' macroscopic regularities and organisations that arise from the individual interactions of agents (Epstein 1999). The agent-based approach "invites the interpretation of society as a distributed computational device, and in turn the interpretation of social dynamics as a type of computation" (Epstein 1999: 41). The outcome of this body of work suggests that social systems can be represented in generativist ways; in other words, that from the interactions of individuals, complex patterns in social systems can arise: a decentralised system arises from the behaviour of heterogeneous autonomous agents. What is not known, however, is the potential of complex systems of problem-solving to solve societal problems that themselves arise from the problem-solving

behaviour of individuals that contribute information and knowledge to a probabilistic problem-solving system. This tension, between the behaviours of individuals and the behaviours of systems, has also been modelled using computational analysis, in the form of emergent computation.

Emergent computation is a term used to describe systems that produce emergent behaviour that is a computation, derived from individual inputs that produce global behaviour that emerges from many local interactions (Forrest 1990). Systems that exhibit emergent computation are typically associated with certain characteristics, namely self-organisation ("the spontaneous emergence of order from an initially random system"), collective phenomena ("in which there are many agents, many interactions among the agents, and an emphasis on global patterns"), and cooperative behaviour (i.e. "that the whole is somehow more than the sum of the parts" (Forrest 1990: 8). At the crux of any notion that probabilistic problem-solving systems can be highly effective in problem-solving is the question of how much more effective such a system would be over and above the sum of the contributions of its individual contributors. Harnessing collective forces for problem-solving might offer a novel way to advance science, using mechanisms that exist across different contexts in nature. Forrest (1990: 8) illustrates this concept with reference to the way the collective organisation and behaviours of ant colonies are "highly sophisticated, including such activities as mass communication and nest building" and where, in the "absence of any centralised control, the collective entity (the colony) can 'decide' (the decision itself is emergent) when, where, and how to build a nest – self-organising, collective, and cooperative behaviour in the extreme".

However, this emergent property can also have detrimental consequences. In the case of internet transmissions, messages have been found to start as randomised information flows but to self-organise into higher-level structures, or token-passing rings (Forrest 1990). On the basis of the discussion above, proposition e is drawn, that *emergent patterns of problem-solving behaviour will emerge from the application of probabilistic innovation, which will provide different and new opportunities for solving problems*.

Figure 2 diagrammatically illustrates the theoretical framework of probabilistic innovation. In this diagram, knowledge inputs are shown to cross the emergent threshold (indicated by the broken line), at which point the relationships predicted by the theory reviewed above are expected to manifest themselves. In other words, the currently uncharted territory of emergent problem-solving is expected to exist past this threshold. Crowdsourced R&D inputs are processed and are fed back into the crowd in order to further increase the *velocity of knowledge creation*. The velocity of knowledge creation is a measure of the speed at which the system can generate quality

innovative problem-solving outputs, which in turn are expected to be a function of the volume of inputs, or to be the tail end of an expected normal probability curve representing all outputs. This framework suggests a 'transmission mechanism', in the form of a probabilistic framework of innovation, between the high levels of information and knowledge available in decentralised locales, some of it relatively tacit and embedded in individuals, and the 'front-line' of societal problem-solving. This theoretical framework therefore relates to how scientific knowledge creation systems can potentially be designed to produce an 'accelerated' process of innovation. The framework suggests that crowdsourced R&D can, under certain conditions, be systematically robust to the constraints associated with the decentralisation or non-transferability of information and knowledge (Polanyi 1973; Hayek 1945; Nonaka 1994; Von Hippel 1994) by allowing inputs from large numbers of people to be focused on specific problems.

What has not been discussed, however, is the technical interface, or the systems that process these inputs and feed them back into the crowd. A discussion of these is beyond the scope of this paper. Nevertheless, certain of the principles considered above have also been applied to the use of computational systems. The application of probabilistic innovation might be well complemented by these systems.

Principles of evolutionary theory have also been applied to computational systems. Computational systems can also supplement systems of probabilistic innovation. Genetic algorithms (GAs) are computational models of the evolution of artificial-life systems; these are taken by some to be an abstraction of biological evaluation (Mitchell & Forrest 1994). Populations of 'chromosomes' are modelled, or bit strings, that represent potential solutions to problems, and selection is taken to occur, using crossover exchanges and mutations, to develop a 'fitter' system (Mitchell & Forrest 1994). This process begins with a randomly generated population of these chromosomes, or potential solutions to a problem, and then the fitness of these chromosomes is calculated as selection and genetic operators (the processes of crossover and mutation) are then applied to create a new population, and the process is then repeated continuously (Mitchell & Forrest 1994). Repeated in steps, termed generations, certain highly fit chromosomes can develop (Mitchell & Forrest 1994). GAs have applications in optimisation, automatic programming, machine and robot learning, and in modelling innovation processes, economic markets, immune system effects, ecological systems, population genetics, interactions between evolution and learning, and social systems, such as the evolution of cooperation and communication (Mitchell & Forrest 1994). It is therefore possible that computational methods may be able to supplement systems of probabilistic innovation. This is an avenue of discussion, however, that is beyond the scope of this article. Further research might consider how

computational systems might synergistically interface with probabilistic systems of knowledge management to further accelerate the velocity of knowledge creation.
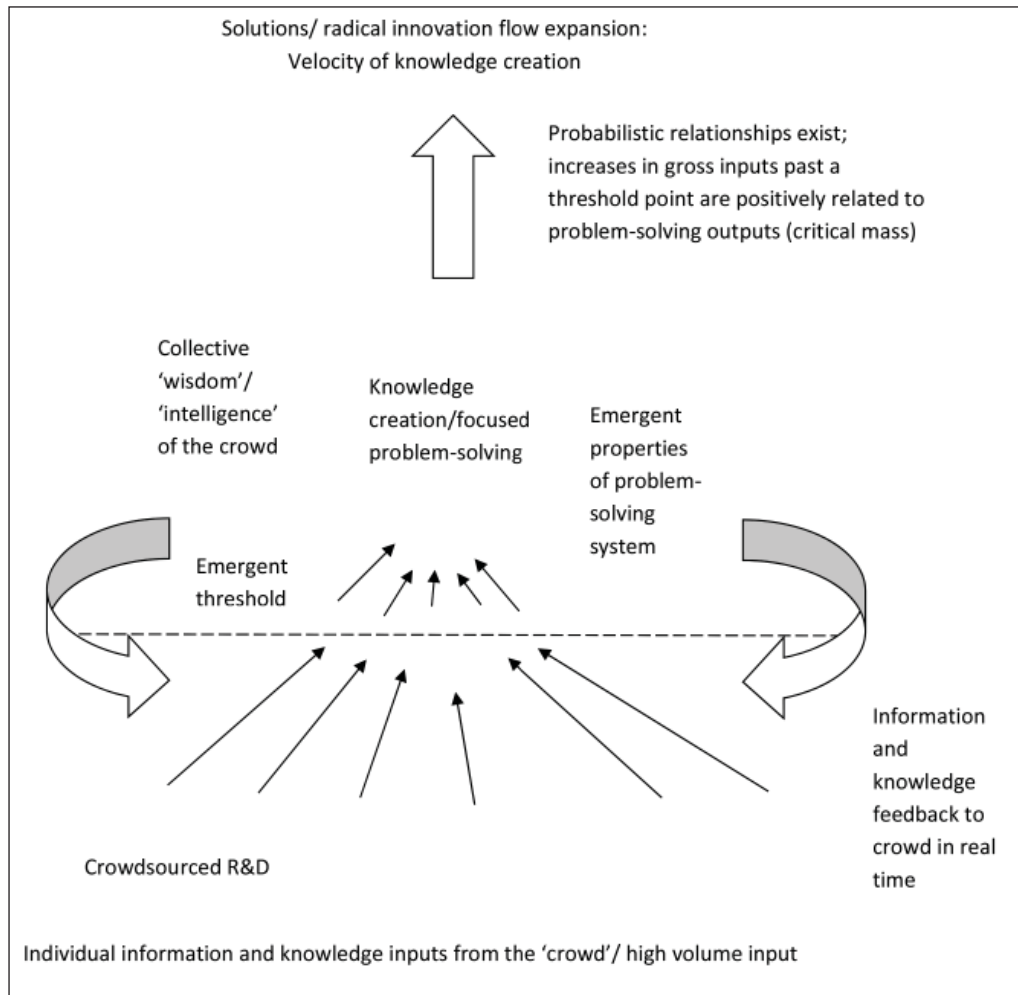

Figure 2: Framework of probabilistic innovation

Conclusions and implications for theory and practice

The objective of this paper was to offer a conceptual framework of probabilistic innovation. A theoretical argument was offered, that due to a host of different constraints, the flow of innovations (proxied by the lack of new innovations in the pharmaceutical industry as a case in point) faces a 'probabilistic threshold', as innovation systems typically apply non-probabilistic knowledge management

processes. An alternative paradigm of innovation, in the form of SGR, was discussed, and a range of different testable propositions was derived from an analysis of the literature. A framework that relates certain of these relationships was then presented. On the basis of the analysis and discussions, certain implications for theory and practice seem evident. This analysis has offered improved insights into the boundary conditions to theory that relate different forms of market failure to innovation failure. Further theorising might use the probabilistic lens offered by this work to interrogate the different conditions under which probabilistic innovation may best be used to accelerate innovation. The analysis suggests that practitioners engaged in pharmaceutical and medical research may benefit from using more innovative knowledge management systems such as crowdsourced R&D, and may also benefit from using these systems to bypass bottlenecks in their R&D efforts. Those seeking non-proprietary, or non-profit, forms of innovation such as state or non-profit medical research entities may be emboldened to use crowdsourced R&D as a social form of R&D that is less resource dependent than its proprietary counterpart. The promise of probabilistic innovation, as argued in this article, might be overstated, because it is difficult to imagine a social innovation system that can accelerate R&D to quickly and effectively solve problems such as Ebola, HIV and antibiotic resistance. Nevertheless, if the ideas offered in this paper serve to encourage theorists to further engage with the increasingly dangerous consequences of innovation failure, then this work will have made an important contribution.

References

Adair, J.G. & Vohra, N. 2003. 'The explosion of knowledge, references, and citations', *American Psychologist*, 58(1): 15–23.

Ajumobi, K. 2014. War against polio: Nigeria reports single case in 2014. [Online] Available at: http://businessdayonline.com/2014/04/war-against-polio-nigeria-reports-single-case-in-2014/. Accessed: 6 April 2014.

Bornmann, L. 2010. 'Does the journal review select the "best" from among the work submitted? The state of empirical research', *IETE Technical Review*, 27(2): 93–96.

Bornmann, L. 2011. 'Scientific peer review', *Annual Review of Information Science and Technology*, 45: 199–245.

Callaghan, C.W. 2013. Organisational culture, individual values and research productivity. Unpublished doctoral thesis, University of the Witwatersrand. [Online] Available at: http://wiredspace.wits.ac.za/handle/10539/104/browse?value=Callaghan%2C+Christian+William&type=author. Accessed: 30 December 2014.

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.

Campanario, J.M. 2009. 'Rejecting and resisting Nobel class discoveries: Accounts by Nobel Laureates', *Scientometrics*, 81(2): 549–565.

Campbell, D.T. & Fiske, D. W. (1959). 'Convergent and discriminant validation by the multitrait-multimethod matrix', *Psychological Bulletin,* 56: 81–105.

De Long, J.B. & Lang, K. 1992. 'Are all economic hypotheses false?', *Journal of Political Economy*, 100(6): 1257–1272.

Dewald, W.G., Thursby, J.G. & Anderson, R.G. 1986. 'Replication in empirical economics: The *Journal of Money, Credit and Banking* project', *American Economic Review*, 76(4): 587–603.

Epstein, J.M. 1999. 'Agent-based computational models and generative social science', *Complexity*, 4(5): 41–60.

Forrest, S. 1990. 'Emergent computation: Self-organizing, collective, and cooperative phenomena in natural and artificial computing networks', *Physica D*, 42: 1–11.

Halifax, J. 2013. Superbugs immune to antibiotics a threat to humans as big as terrorism or climate change. [Online] Available at: http://www.birminghammail.co.uk/lifestyle/health/superbugs-immune-antibiotics-bigger-threat-3410907. Accessed: 30 May 2013.

Hanson, R. 1990. Could gambling save science? [Online] Available at: http://li.mit.edu/Stuff/CNSE/Paper/Hanson90.pdf. Accessed: 1 December 2014.

Hanson, R. 2000. Shall we vote on values, but bet on beliefs? [Online] Available at: http://hanson.gmu.edu/futarchy2000.pdf. Accessed: 1 December 2014.

Hanson, R. 2003. 'Combinatorial information market design', *Information Systems Frontiers*, 5(1): 107–119.

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

Hayek, F.A. 1973. *Law, Legislation and Liberty*. Chicago: University of Chicago Press.

Henrion, M. & Fischhoff, B. 1986. 'Assessing uncertainty in physical constants', *American Journal of Physics*, 54(791): 791–798.

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*, 14/2006. [Online] Available at: http://sistemas-humano-computacionais.wdfiles.com/local—files/capitulo%3Aredes-sociais/Howe_The_Rise_of_Crowdsourcing.pdf. Accessed: 27 May 2013.

Jay, J. 2013. 'Navigating paradox as a mechanism of change and innovation in hybrid organizations', *Academy of Management Journal*, 56(1): 137–159.

Johnson, N.L. 1998. Collective problem solving. Los Alamos National Laboratory Working Paper LA-UR-98-2227. [Online] Available at: http://collectivescience.com/Documents_1/Johnson%20Collective%20Problem%20Solving%20AB%20paper%20v11f%201998.pdf. Accessed: 6 October 2014.

Kochugovindan, S. & Vriend, N.J. 1998. 'Is the study of complex adaptive systems going to solve the mystery of Adam Smith's "invisible hand"?' *The Independent Review*, 3(1): 53–66.

Krishnan, J. 2014. Panic as deadly Ebola virus spreads across West Africa. [Online] Available at: http://www.independent.co.uk/news/world/africa/panic-as-deadly-ebola-virus-spreads-across-west-africa-9241155.html. Accessed: 6 April 2014.

Lakhani, K. R. & Panetta, J. A. 2007. 'The principles of distributed innovation', *Innovations: Technology, Governance, Globalization*, 2(3): 97–112.

Le Bon, G. 1896. The crowd: A study of the popular mind. [Online] Available at: http://www.gutenberg.org/cache/epub/445/pg445.html. Accessed: 28 September 2014.

Maher, K. 2013. 'The drugs don't work', *Time Magazine*. [Online] Available at: http://www.bbc.co.uk/news/health-22649922. Accessed: 30 May 2013.

Mahoney, M.J. 1977. 'Publication prejudices: An experimental study of confirmatory bias in the peer review system', *Cognitive Therapy and Research*, 1(2): 161–145.

McKenna, M. 2014. Sneak peek: What the White House is thinking about antibiotic resistance. [Online] Available at: http://www.wired.com/2014/04/pcast-resistance/. Accessed: 6 April 2014.

Mitchell, M. 1996. *An Introduction to Genetic Algorithms*. London: MIT Press.

Mitchell, M. & Forrest, S. 1994. 'Genetic algorithms and artificial life', *Artificial Life,* 1(3): 267–289.

Murray, F. 2010. 'The oncomouse that roared: Hybrid exchange strategies as a source of distinction at the boundary of overlapping institutions', *American Journal of Sociology*, 116(2): 341–388.

Nonaka, I. 1994. 'A dynamic theory of organizational knowledge creation', *Organization Science*, 5(1): 14–37.

Peters, D. & Ceci, S.J. 1982. 'Peer-review practices of psychological journals: The fate of published articles, submitted again', *Behavioral and Brain Sciences*, 5(2): 187–195.

Plowe, C.V., Djimde, A., Bouare, M., Doumbo, O. & Wellems, T.E. 1995. 'Pyrimethamine and proguanil resistance-conferring mutations in *Plasmodium falciparum* dihydrofolate reductase: Polymerase chain reaction methods for surveillance in Africa', *American Journal of Tropical Medicine and Hygiene*, 52(6): 565–568.

Polanyi, M. 1973. *Personal Knowledge: Toward a Post-critical Philosophy*. London. Routledge & Kegan Paul.

Simon, H. & Newell, A. 1971. 'Human problem solving: The state of the theory in 1970', *American Psychologist*, 26(2): 145–159.

Smith, A. 2003[1776]. *The Wealth of Nations*. New York: Bantam Dell.

Surowiecki, J. 2004. *The Wisdom of Crowds*. London: Little, Brown.

Von Hippel, E. 1994. '"Sticky information" and the locus of problem solving: Implications for innovation', *Management Science*, 40(4): 429–439.

Zia, A. 2014. Newly diagnosed child brings polio cases to 40. [Online] Available at: http://tribune.com.pk/story/691219/newly-diagnosed-child-brings-polio-cases-to-40/ Accessed: 6 April 2014.