Towards an Intelligent Network for Matching Offer and Demand: from the sharing economy to the Global Brain

Francis Heylighen
Global Brain Institute
Vrije Universiteit Brussel
fheyligh@vub.ac.be

Abstract: We analyze the role of the Global Brain in the sharing economy, by synthesizing the notion of distributed intelligence with Goertzel's concept of an offer network. An offer network is an architecture for a future economic system based on the matching of offers and demands without the intermediate of money. Intelligence requires a network of condition-action rules, where conditions represent challenges that elicit action in order to solve a problem or exploit an opportunity. In society, opportunities correspond to offers of goods or services, problems to demands. Tackling challenges means finding the best sequences of condition-action rules to connect all demands to the offers that can satisfy them. This can be achieved with the help of AI algorithms working on a public database of rules, demands and offers. Such a system would provide a universal medium for voluntary collaboration and economic exchange, efficiently coordinating the activities of all people on Earth. It would replace and subsume the patchwork of commercial and community-based sharing platforms presently running on the Internet. It can in principle resolve the traditional problems of the capitalist economy: poverty, inequality, externalities, poor sustainability and resilience, booms and busts, and the neglect of non-monetizable values.

Keywords: global brain; Internet; sharing economy; distributed intelligence; offer networks

Introduction

The relentless innovations accompanying the development of the Internet and related ICT may seem overwhelming. Hardly a week passes by without a revolutionary new technology, social network, or lucrative new business application being announced. Initially, in the 1990's, the applications of the Internet were primarily focused on publishing and retrieving information. The first decade of the new millennium saw a proliferation of so-called Web 2.0 collaborative communities and social networks. The present decade seems most in thrall with the sharing economy and the applications of Artificial Intelligence (AI).

With this explosion in the number of new applications and trends, forecasting the long-term future of our information society appears like a daunting prospect. Yet, there is a paradigm that offers hope for a more integrated vision: the *Global Brain* (Goertzel, 2002; Heylighen, 2002, 2008; Last, 2014). The analogy underlying this perspective is that the

Internet increasingly starts to play the role of a brain for the planet, i.e. a distributed, intelligent network that supports humanity in solving its problems and coordinating its actions. The initial inspiration for the Global Brain model of Internet development (Heylighen & Bollen, 1996; Mayer-Kress & Barczys, 1995) came from its role as a medium for the exchange of information. However, this perspective seems less applicable to its more recent social and economic functions, where the emphasis is on exchanging goods, services and experiences rather than mere information.

The present paper wishes to propose a broader interpretation of the Global Brain (GB) metaphor—one that is directly applicable to the sharing economy and to the more distributed forms of social organization that accompany it. It will do so by synthesizing the notion of the GB as a distributed, intelligent network with Ben Goertzel's newly proposed concept of an *offer network*, i.e. an architecture for a future economic system that is not centered on the accumulation of money, but on the direct matching of offers and demands (Goertzel, 2015).

The intention is to show that the main applications of the Internet are likely to become integrated into a single, universal system for coordinating all the activities of the people and machines on this planet. Such a system would immensely reduce the confusion, friction and waste caused by poorly aligned activities, while boosting the synergy of collaboration. Complemented by on-going technological innovation, the resulting increase in productivity would create an *economy of abundance* (Diamandis & Kotler, 2012; Dugger & Peach, 2015), where all needs can be satisfied at negligible costs. The combination of abundance with an intelligent, bottom-up system of coordination should eventually produce a solution for all the major problems that plague humanity, including global warming, poverty, inequality and conflict. This utopian but realizable scenario for the mid-term future has been called "return to Eden" (Heylighen, 2014a).

To get there, we will first review the abstract conception of the Global Brain as a distributed mind, and some of the concrete applications of the sharing economy. We will then elaborate a generalized concept of an offer network that applies the cognitive capabilities of a distributed mind to the practical opportunities and needs of our society. Finally, we will argue that the various experiments with a sharing economy that we are witnessing are converging towards such a universal network, and that such a network would be able to solve our present problems of poverty, inequality, sustainability and resilience.

The Global Brain as a distributed mind

A mind in the most general sense can be defined as an intelligent, autonomous system that collects and processes information to assess its situation and then decides how to act in order to realize its desires. In other words, a mind is a *sense-making agent*: it interprets and evaluates the phenomena it perceives in order to extract their meaning with respect to its value system, and then it acts based on that interpretation in order to further its values.

For an agent that is the product of evolution, such as an organism, these values are rooted in survival and growth, because agents that do not hold these values are eventually eliminated by natural selection. Artificial agents can in principle be programmed with different values, such as serving their designer, but it seems unlikely that they would last long in a complex environment without at least some inbuilt "survival instinct". Moreover, their lack of autonomy may disqualify them as true "minds".

The activity of a mind can be summarized by the basic cybernetic feedback loop:

perception of the situation → interpretation with respect to preferences or desires → action to bring the situation closer to the desires → new perception to ascertain in how far the action was sufficient → new interpretation → new action → …

A perceived situation will elicit action if it entails a *problem*—i.e. a (threat of) deviation from the desired situation—, an *opportunity*—i.e. a possibility to advance towards the desired situation—, or some combination of the two. Situations that elicit actions can be called *challenges*: they challenge the agent to remedy the problem or to exploit the opportunity (Heylighen, 2012, 2014b).

The intelligence of the agent resides in its ability to recognize and effectively address the most relevant challenges. This requires knowledge, in the sense that the agent must be able to recognize different categories of situations (which we will call *conditions*), and to associate each condition with the action most appropriate to deal with it. The elements of such knowledge can be expressed most simply as "condition-action rules" or "production rules", with the following form:

$a \rightarrow b$

This is to be read as: IF condition a is perceived, THEN perform action b. For example: *banana → eat*, *tiger → flee*, *tired → rest*.

Such conditions that immediately lead to actions are merely the simplest form of knowledge. More complex situations require a process of *inference*, in which perceived aspects of the situation (perceptions) imply more abstract conditions (conceptions), which in turn imply even further conditions, until the process settles on a particular action. Here is a simple example:

tiger → predator
predator → danger
danger → flee

More complex processes of inference will moreover take into account conjunctions of conditions and actions (which we will denote by the "+" symbol), e.g.:

striped + large + cat-like + animal → tiger
tiger + jungle → flee
flee + jungle → locate tree + climb tree

A situation will normally be characterized by several recognizable conditions. These perceived conditions will trigger different rules that infer additional conditions, which in turn trigger further rules, and so on. Thus, the initial perception will be processed through the application—in parallel and in sequence—of the different rules that constitute the agent's knowledge, until it settles on an interpretation and, if need be, an associated course of action.

Such rule-based processing of information becomes even more flexible when the different rules have different "strengths", denoting their relative importance or probability of being correct. Strength can be represented by a number between 0 and 1, where 1 denotes a rule whose conclusion is absolutely certain. The principle is that if several rules compete for execution, the one with the highest strength will be chosen. Alternatively, if several rules act in parallel, then their contribution to the final interpretation will be proportional to their strength. These strengths should be able to adapt to experience: the more successful a rule has proven to be, the larger its strength should become. This is the basic mechanism of *reinforcement learning* that rewards good rules and weakens less good ones (Woergoetter & Porr, 2008).

These elements (conditions, actions, and their conjunctions, rules expressing elementary inferences, and adjustable strengths) correspond to those of the "production rule systems" (Anderson, 2014) that are used as general-purpose representations of knowledge and inference in Artificial Intelligence (AI) and cognitive science. They basically allow us to recover the flexibility of the neural networks that process information in the brain (Heylighen, 2014c; McLeod, Plunkett, & Rolls, 1998).

The present analysis of the functional components needed to build a mind is on purpose so general that it can apply to very different kinds of minds—including those exhibited by human brains, by AI computer programs, but also by the collectives of human and technological agents that together would form a "global brain". Let us see how these different components are realized in society.

First, society is an autonomous system: it is able to survive and grow by solving the problems or exploiting the opportunities that it encounters—and this without need for outside direction. It recognizes such challenges through its (largely implicit) value system, which evaluates certain conditions as beneficial (e.g. education, shelter, drinkable water, peace, …) and others as harmful (e.g. hurricanes, pollution, crime, disease, …) to its development. It tackles these challenges by analyzing and interpreting the situation, and by initiating actions to deal with it (e.g. purifying water, building houses, curing people from disease, preventing crime, …). This means that society possesses an implicit store of knowledge and an intelligence that applies that knowledge by making the necessary inferences and eventually reaching decisions about the actions to take.

This intelligence is not localized in a central executive, such as a king or government, that would command and control the social system. It is rather distributed over billions of individuals, organizations, networks, documents containing specialized knowledge and regulations, computer programs, and machines that perform actions. Thus, the intelligence of society is similar to the one inherent in the brain (Minsky, 1988), where knowledge and inference processes are distributed across billions of neurons and their connecting synapses. But that insight is not yet sufficient to understand how this societal intelligence functions, how it can be improved, or how it is likely to further evolve. Let us therefore continue our analysis at the functional level rather than at the level of the physical components performing these functions.

Matching offer and demand

The values of society present themselves as a collective *demand* for better conditions (Heylighen, 1997). Whenever a deficiency is noted or a possibility for improvement is conceived, some individual or collective agent will express a desire to get such improvement. The best-known instantiation of this desire is the economic demand for a particular type of goods (such as cars or apples), services (such as house cleaning or medical treatment), or professionals (such as engineers or interpreters). But demand in the economic sense is measured merely by the price agents are willing to pay for the thing they desire. This ignores the desires of agents that lack money, or their desires for things that cannot be bought—such as friendship, clean air, or democracy.

A well-functioning society should be able to satisfy the basic needs of all of its members. The only restriction is that desires should not be mutually inconsistent—like when two people desire to own the same property, or when some desire to turn a patch of forest into agricultural land, while others wish to safeguard the species living on that patch. Part of the function of intelligence is to resolve the unavoidable inconsistencies that arise in any complex value system, by either determining a good trade-off between values, or, even better, finding a win-win solution that satisfies them all.

We can include demand in our production system representation of intelligence by adding requests or queries. A *query* is a production rule of which the input condition (left-hand side of the rule) remains unspecified:

$$? \rightarrow b$$

This can be read as “what is needed to achieve condition b?”. The strength of the rule corresponds to the importance of the demand, i.e. its value relative to other queries.

A demand will be satisfied if an answer to the query is found. This answer has the form of a potential *offer*, i.e. a condition a that can be produced somewhere in the system, and that would result in the desired condition b via one or more intervening production rules, e.g.

$a \to b$, or

$a \to c + d$,
$c + e \to f$,
$f \to b$.

When the offer is immediately available, it functions as an opportunity or *supply*, i.e. an existing resource that merely needs to be recognized and exploited. For example, someone looking for a date may just need to check the profiles of people who expressed their interest in dating on a website and select the most fitting one. In the most general case, satisfying a demand will require some action in order to realize the necessary condition. For example, someone who is looking for a date at a party will need to chat with different potential candidates in order to convince one to go on a date. In this case, the supply of a solution is merely potential, as additional conditions will need to be realized by performing the appropriate actions that would mobilize the needed resource.

In either case, the satisfaction of a demand happens through the *matching of offers and demands*. This is a process of inference that requires intelligence: it needs to seek out and make sense of all the different actual and potential conditions that could satisfy the demand, and then find the most efficient route, demanding the minimum investment of action, that would lead from the present condition to the desired condition. In the production systems used in AI, this process is typically performed by an *inference engine,* which uses a variety of heuristic algorithms to find the best match using the least search.

Perhaps the simplest heuristic is *backward chaining*. Here, the inference engine starts from the query (final condition to be satisfied). It then searches for one or more production rules that would produce this final condition. If the initial conditions for these rules are not fulfilled, it repeats the search for preceding production rules whose end condition would match the beginning condition of the first batch of rules. This moving "backwards" to find prior, preparatory rules is continued until it finds a rule whose initial condition is already satisfied (i.e. there exists an offer in the system for producing this condition). If it cannot find an existing condition, it will determine the minimal actions necessary to create such a condition.

Let us illustrate this strategy with the dating example. Suppose you are looking for a date, but do not know anybody available. This defines the initial query. You ask your friend Carl whether he knows a possible candidate. If he can bring you into contact with a good date the problem is solved. But Carl may not know anybody, and therefore he consults his friend Lisa. Lisa in turn may consult her friends Peter, Tom and Beth, one of whom may eventually propose a good match to your query. Thus, your query had to move "backward", increasingly farther away from your present demand, until it met a match.

Another strategy is *forward chaining*. Here the system starts from the existing conditions (offer) and applies the production rules until it produces a condition that is in demand. For example, without knowing about your query, Beth introduces one of her friends to Lisa who introduces him to Carl who eventually introduces him to you, where he is pleased to discover that you are willing to go on a date with him. Here, your query did not need to travel away from you; it was rather the solution that traveled towards you. There exist a variety of more complex strategies that combine forward and backward chaining with other methods (such as constraint satisfaction) in order to find the best matches between queries and solutions to the queries, or more generally between offers and demands.

The truly intelligent systems will furthermore learn from their experience in finding matches, typically by strengthening the most effective rules and heuristics and weakening the less effective ones. In addition, they will generate potentially better rules through the variation and recombination of effective rules, and then test them out so as to keep only the best ones (Holland, 2012; Holland, Holyoak, Nisbett, & Thagard, 1989). In this way they become ever more intelligent, i.e. better at matching offers and demands.

When we go back from considering society as an abstract intelligence to society as a network of interacting agents, we see that these processes are implemented through what may be called *challenge propagation* (Heylighen, 2014b; Heylighen, Busseniers, Veitas, Vidal, & Weinbaum, 2012). Production rules are normally executed by agents: people, organizations, or technological systems. For example, a person who encounters a challenging condition (problem or opportunity) will typically act on that condition using her or his (cognitive) skills and (physical) abilities in order to gain some value or benefit. This will change the condition so that it better satisfies the person's value system. For example, if your friend is lonely (challenge), then you will introduce him to other friends (action), so that he (and by extension you) may feel better (gain in value).

But a single agent normally cannot satisfy a broad, societal demand. Agents have limited abilities. That means that the new condition the agent produces, while offering some value, will in general not be a full solution to the problem or a full exploitation of the opportunity. On the other hand, different agents have different abilities, and therefore the one may be able to tackle some of the challenge left unresolved by the other. Thus, the as yet unsatisfactory condition defines a new challenge for one or more agents with complementary abilities that may produce further improvement. For example, the people to whom you introduced your friend may not really be the ones he was hoping to meet, but they may introduce him to further people they know, and so on, until he perhaps meets his true soul mate…

Thus, the initial challenge propagates from agent to agent across a network of social connections, each one applying its own unique skills in meeting demand, until it is fully resolved (i.e. no further improvement is worth investing in) (Heylighen, 2014b). Such challenge propagation is the common mechanism through which distributed cognition (Hutchins, 2000), problem-solving or query-resolution takes place across a social system.

Social media and the sharing economy

In a traditional social system, the processes of challenge propagation tend to be slow, laborious, and unreliable. If challenges are passed on from individual to individual via existing links of acquaintanceship or collaboration, it is unlikely that they will reach the agent(s) best positioned to tackle the challenge (i.e. exploit the offer or satisfy the demand) within the shortest time. The result is a waste of potential synergy between the different agents.

The Internet is a powerful medium for cutting short such meandering, haphazard searches for good matches. First, information about a challenge posted on the Internet can in principle reach any agent anywhere in the world without delay. Second, such information can be registered very precisely, without any distortion or forgetting of details, for as long as it remains relevant. Finally, the information can be analyzed, aggregated and routed to the right destination by intelligent software.

This makes it much easier to find good matches between the billions of challenges that confront the world population—in particular by proposing offers to the agents most likely to need them, and demands to the ones most ready to satisfy them. In other words, the Internet can play the role of a highly intelligent mediator, which coordinates all the actions needed to tackle challenges, in such a way that problems are resolved and opportunities are exploited in the most direct manner. This avoids such sources of friction as queries that do not find an answer, offers that are not taken up because no one knows about them, effort needlessly spent in searching for things that are readily available, poor matches where what you get is not really what you want, and unnecessary delays where agents suffer while waiting for an available solution that has not yet reached them.

An enormous variety of protocols, platforms, systems and communities have arisen over the past two decades to tackle such difficulties. Let us quickly survey some common applications. *Email* mediates between individuals so that the one can ask a question or propose a solution to the other without being limited by distance, time or cost. *Search engines* mediate between a user who has a query and the vast amount of information already available on the web that may answer that request. Information sharing sites, such as Wikipedia, collect, organize and publicize potentially useful answers, thus making them easier to find. *Social media* automate the propagation of offers and demands across extended social networks. They also help you to maintain your social connections, such as friendships or professional acquaintanceships, by keeping track of how you can reach the people in your social network if you need to—even if you have not seen them in years.

Web stores or *marketplaces*, such as Amazon or eBay, mediate between clients wishing to purchase some item, and vendors ready to sell it. They provide access to a much larger variety of items than what is available in a traditional store, while their algorithms can help you to find the best match for the lowest price. *Crowdsourcing* systems (such as Amazon Mechanical Turk) and freelancing sites (such as oDesk) mediate between people or organizations that need some work to be done and people willing to do it. In the case of crowdsourcing, intelligent algorithms aggregate the (relatively small) contributions from a large number of workers into a coherent result (Doan, Ramakrishnan, & Halevy, 2011). This allows the system to tackle truly large-scale challenges, such as mapping out the need for relief after a major earthquake (Gao, Barbier, & Goolsby, 2011). *Crowdfunding* platforms mediate between agents who need capital in order to start up some promising project and the millions of people who may be willing to offer some of their money to invest in such a venture (Agrawal, Catalini, & Goldfarb, 2013; Mollick, 2014). Because of the aggregation of countless offers, only a tiny contribution may be needed from each individual investor in order to tackle a great challenge.

Such approaches to matching offer and demand also impact more specific niches. For example, demand for transportation can be satisfied by websites that find the most efficient combination of bus, train or plane journeys to lead you from your present position to your desired destination at the desired moment. This is still relatively inflexible, because it starts from the existing offer of public transport, which may not satisfy your specific demand. Involving privately driven cars can solve this issue. A flexibly coordinated taxi service, such as Uber or Lyft, matches drivers and passengers depending on proximity, time and price. However, this is still rather costly because the driver needs to fully invest car, time, and fuel into satisfying an individual traveler's demand. Ridesharing is in principle more efficient, as several people with similar destinations or overlapping journeys can simultaneously profit from the same vehicle, driver and fuel. However, this requires a more intelligent coordination algorithm applied to a larger pool of offers and demands in order to be effective (Agatz, Erera, Savelsbergh, & Wang, 2012; Furuhata et al., 2013).

Another recently popular application domain is finding temporary accommodation. Like in the case of ride sharing, there are commercial platforms (such as AirBnB), where people are offered a room in a place they are visiting in return for a fee paid to the room's owner, and non-commercial ones (such as Couchsurfing), where people can stay freely in someone's room on the sole condition of leaving it in a good state. In both cases, the platform mediates between local owners offering rooms and travelers demanding rooms, while seeking an optimal match along criteria such as location, time, fee and reputation of owner or traveler. Another example is neighborhood sharing, where people living in a particular community use a website to advertise the rarely used equipment (such as lawnmowers, ladders, or drills) they are willing to lend out to other community members—on the sole condition that they too can borrow something from the community when they need it. The resulting large and diverse offer removes the need for most individuals to buy such equipment, thus strongly reducing costs and resource consumption.

The drive for disintermediation and universality

A general strategy to increase the efficiency of a socio-economic system is *disintermediation* (Gellman, 1996)—a process more popularly known as "cutting out the middleman". The more links there are in a chain along which challenges propagate between offer and demand, and the more restrictive, demanding or inefficient these intermediate links, the more time, money and effort is needed to reach the end of the chain, and the less value will remain. Reducing the number of intermediate stages can strongly increase the throughput of the chain. The Internet has boosted productivity in part by greatly reducing the number of middlemen, in general by allowing producers of goods and information to connect directly with the consumers. For example, authors can now immediately present their manuscripts or articles to their readers, instead of having to go through the chain of editors, publishers, printers, transport companies, and bookshops.

However, the proliferation of the various "Web 2.0" sites and companies that we surveyed has created a host of new intermediaries—such as Uber, Facebook, Apple, and AirBnB—which skim part of the profit generated by better coordination. These intermediaries moreover compete with each other. This forces them to spend a major part of their profits in advertising and marketing, and in preventing competitors from getting a foothold in their market segment. They do the latter by trying to "lock in" users to their system, so that they are not tempted to switch to other systems. Therefore, they tend to resist attempts to make the different systems interoperable—i.e. able to seamlessly exchange data. This creates a range of "silos" separated from each other by rigid walls. For the users, the effect is complexity, confusion, and a bombardment with distracting and sometimes misleading information. The overall result is poor coordination and a waste of mental and physical resources.

A similar situation existed when the first computer networks appeared in the 1980s. Instead of the single Internet that we have now, there was a proliferation of local area networks in organizations, community-based bulletin board systems (such as The Well), commercial providers (such as CompuServe and America Online), and academic networks (such as Usenet and Bitnet). In practice, it was very difficult to get any document from one network onto another one. The Internet was initially promoted as a "network of networks", i.e. a medium that made it possible to *inter*connect all *net*works. It achieved this by converting their different communication schemes into the very flexible and robust TCP/IP protocol. This flexibility and universality proved to be so beneficial that the other networks either switched to TCP/IP themselves, thus becoming part of the Internet, or disappeared altogether—because their users abandoned them for the cheaper, more reliable and more transparent Internet.

A similar state-of-affairs again arose in the 1990s when several protocols for the retrieval of documents competed on the Internet, including UUCP, WAIS, FTP and Gopher.

Tim Berners-Lee resolved the problem by defining the URL (Uniform, or Universal, Resource Locator) protocol, which could identify any document residing on any Internet-connected computer (Berners-Lee & Fischetti, 1999). Moreover, he created the HTML protocol for formatting such documents, so that any computer could read them and connect the user transparently to any linked documents, wherever they resided. Thus, HTML and URL resolved the proliferation of incompatible ways of retrieving documents, while defining the World-Wide Web as a universal interface for such retrieval. Shortly afterwards, the first graphical browsers implementing these protocols appeared, thus making "web surfing" so easy that anyone could do it. This was the final breakthrough that created a massive, global interest for the Internet in the late 1990s, with the subsequent dotcom boom and explosion in social media use.

It seems likely that a similar boom will occur in the sharing economy once a universal protocol is developed that would transparently and flexibly match all offers and demands, no matter what kind of goods, services, solutions or information are proposed or requested by whom. Presently, the proliferation of mutually incompatible commercial and community-run platforms makes it nearly impossible for a challenge to propagate from one platform to another.

For example, suppose that I want a ride from my home to my travel destination, while I am willing to offer my room for someone to stay in while I am traveling. But for ride sharing I have to advertise my demand on one website, while for room sharing I have to register on a different website with different rules and protocols. Moreover, the person who takes up my offer for a room has no connection with the person who can satisfy my demand for a ride, and so I cannot earn any credit for my offer of a room, while remaining in debt for taking up someone else's offer of a ride. In addition I may have to pay fees to the different sites just for advertising my offer or demand, while having to invest quite some effort in making my offer or demand fit in with the idiosyncratic requirements of each platform. Finally, any platform I use will only reach a tiny fraction of all the people in the world that could potentially take up my challenge.

All these factors hinder the free propagation of challenges. As a result of such frictions, the process of coordination between all the offers and demands that exist in the world remains highly inefficient. This creates a selective pressure (Heylighen, 2014a) for the development of a universal coordination protocol that would interconnect and eventually absorb or replace all existing sharing platforms—just like the Internet and its TCP/IP protocol did with all preceding computer networks. Thus, such a protocol would eliminate all the intermediaries, enabling immediate peer-to-peer transactions via an open, free, public medium. Such a medium for synergetic interaction would provide the foundation for what has been called the "collaborative commons" (Rifkin, 2014) or "automated commons" (Last, 2016), i.e. a vast virtual space in which anyone can share anything. This would completely change the rules of the present economy. But to show how revolutionary this change would really be, we need to examine how such a protocol could bypass a last, ubiquitous intermediary, namely *money*.

The shortcomings of money

The *invisible hand* of the market is a well-known self-organizing mechanism for the matching of supply and demand. Supply is aligned with demand via the *price mechanism*: as demand for some commodity increases relative to supply, the price of the commodity increases. This increases the potential income for its suppliers, thus motivating them to supply more—until the unfulfilled demand is satisfied and the prices go down again. This constant adjustment through negative feedback keeps supply more or less in balance with demand, without the need for any centralized planning (Heylighen, 1997). However, the market mechanism cannot function without a currency that functions as a standardized unit to express prices in. This quantitative measure of value becomes tangible in the notion of *money*: a supply will only be delivered in return for an amount of money equivalent to the price.

In spite of its flexibility and apparent universality, a money-based, capitalist economy has a number of well-known shortcomings. Most generally, money or price as a one-dimensional variable is much too coarse to express how various offers and demands differ in value according to multiple criteria—such as size, ease of use, beauty, sustainability, or social impact. In particular, there are fundamental values—such as happiness, love, or clean air—that intrinsically do not have a price, because they cannot be transferred from person to person and therefore have no exchange value. Moreover, human psychology does not treat value or preference as what mathematicians call a *total order*, but as a *partial order*: some things cannot be unambiguously ordered according to their value; in some contexts we prefer the one, in other contexts the other. Therefore, value can in general not be reduced to a universal, one-dimensional quantity.

This has led different authors to propose multidimensional, localized measures of value, such as “qualified money” (Helbing, 2014) or various “community currencies” (Richey, 2007; Seyfang, 2001). These express different aspects of value that are important for the local community, independently of the “universal currency” that is money. The problem is that this again creates a proliferation of mutually incompatible platforms, because it is impossible to establish a consensus on what all the relevant dimensions of value are. Attempts at interconnecting the exchange systems of communities that use different kinds of alternative currencies tend to founder in complexity and confusion (Jean-François Noubel, personal communication, 2015).

While this may tempt us to go back to the simplicity of a single currency, there is another fundamental shortcoming of money. Because it can be exchanged for about any good or service, it is in practice the most powerful resource to possess. Moreover, its power increases proportionally to the amount you have: with double the amount of money, you can buy double the amount of things. Its linear, numerical character further means that it can be accumulated unrestrictedly without losing any value. In this it is different from most other values, such as food, shelter or rest, which are subjected to diminishing returns: once you have had enough, you do not benefit by getting more. This has turned money into the dominant value of our society. Thus it has pushed aside values such as sustainability, community, happiness, or wisdom that are more fundamental, but more difficult to quantify or to accumulate. As a result, for many people money has become its own reward—rather than a means to a more basic end such as well-being, self-actualization, or social harmony.

The dynamics of accumulation moreover leads to *inequality*, according to the well-known positive feedback of the *rich getting richer*: the more money you have, the easier it is to invest surplus money in property, stocks or innovations that will get you even more money (Piketty, 2014). This leads to the concentration of wealth in an ever-smaller group. An unfortunately common side effect is the *poor getting poorer*: the less money you have, the more difficult it is to invest in any business or form of self-development, such as education or health care, that would increase your income. What is worse, you will also have more difficulty to pay your general cost of living, and to find the resources, time or energy to tackle any unexpected problem that may occur (Mullainathan & Shafir, 2013). This makes it likely that you end up in a spiral of accumulating debt and the concomitant poverty.

This *Matthew effect* is boosted by globalization and the Internet: the easier movement of money, goods and information increases the reach of the most successful, thus increasing their profit, while reducing the profit of the less successful—even when the things they offer are virtually identical (Brynjolfsson & McAfee, 2014). The positive feedback of the Matthew effect in a sense negates the negative feedback of the invisible hand, where increase of demand is supposed to increase supply and thus again decrease demand (Heylighen, 1997). People who fall into poverty, e.g. because of illness, unemployment or debt, actually have a greater need for things like medical care, accommodation, and transport than people who live comfortably in their own home from the interests on their savings. But that demand is ignored by the invisible hand because poor people have no money to pay for the things they need. Therefore, as their situation worsens because they do not get what they need, their real demand increases rather than decreases.

Another form of positive feedback is produced by *speculation*, which is a demand for things that are bought not for their intrinsic worth, but for the sole reason that they are expected to increase in price. Thus, increasing demand can lead to an expectation of prices increasing further and therefore to even higher demand. This vicious cycle produces the bubbles, booms and busts typical of a capitalist economy, and the concomitant market volatility and economic crises. This instability too is amplified by the easier movement of capital across the Internet (Heylighen, 2007a).

In sum, while money appears like a simple and flexible medium for matching offer and demand, in practice it produces a strongly distorted picture of the actual demand, ignoring many fundamental needs and values of society, while creating speculative demand for things that are not intrinsically valuable. Thus, it amplifies minor differences into growing inequality and instability. That is why Goertzel proposed his concept of offer network as a system that would match offer and demand without money as an intermediate (Goertzel, 2015).

From barter economy to gift economy

At first sight, Goertzel's system would function like an ICT-enhanced *barter economy*, i.e. a platform through which people would indicate which goods or services they are willing to exchange for which other goods or services. As a foundation for a complex economy, the mechanism of barter is highly inefficient, because it assumes that you should match one person with an offer of A and a demand for B with one other person who has a complementary demand for B and offer of A, so that A can be exchanged for B. Most offers will not come from people who have demands precisely complementary to the one of the person who would take their offer. Therefore, it is intrinsically difficult to find a precise match that could lead to an exchange.

To prevent this well-known problem, Goertzel proposes to model all the possible offer-demand pairs, independently of the person who makes them, as production rules. These have the form $A \rightarrow B$, meaning that someone is willing to exchange B for A, i.e. offer B on the condition of receiving A. As we saw, the inference engines developed in AI provide a variety of intelligent algorithms that search for matches across long chains of such condition-action rules, thus "closing the loop". This means that the demand of any person in the loop can be fulfilled by the offer of some other person in the loop, without need for one-to-one exchanges between pairs of individuals. The larger the number of people participating, and the larger the number of condition-action or offer-demand rules they enter into the system, the easier in principle it becomes to find a match that satisfies most of the demands. Thus, an intelligently designed global offer network seems as if it might be flexible enough to replace a money-based market economy.

Yet an offer network is at present little more than a promising idea. Its implementation is likely to be very complex, because a protocol will have to be developed that allows all offers and demand to be formulated in a precise, unambiguous, and universal manner. Otherwise, the inference engine will not be able to recognize which combinations of offers precisely satisfy which combination of demands. Without the universal intermediate of money to absorb the differences, it will be very hard to find exact matches. Still, there are other ways to create sufficient slack in the system so that it can absorb imperfect matches.

This may be understood by going back in history to pre-monetary economic systems. While most people assume that before the invention of money exchanges took place through barter, anthropologists have observed that communities without money rather rely on *gifts* (Cheal, 1988; Graeber, 2014). The principle of a *gift economy* is that someone who has an offer X that s/he does not need simply gives it away to someone with a demand for X—*without a priori condition*.

At first sight, this seems like a recipe for abuse, allowing free riders to exploit the gifts of others without doing anything in return. But in a small enough community everybody knows everybody, so that the community members have a good idea of who has given what and who has received what. People who give more than they receive gain prestige or reputation. On the other hand, people who take out more than they put in lose reputation. People with high reputation are more likely to receive what they want if they express a demand. This motivates them to increase their reputation further by giving more. On the other hand, if someone's reputation would drop beneath a certain level, then people would stop giving anything, thus effectively excluding that person from the community and its benefits. This is an effective measure against free riders. Still, the needy—such as children, old or ill people—will not suffer loss of reputation even if they cannot reciprocate, because the community knows that they are not willfully abusing the system, and have made or will make contributions in other periods of their life.

Such a gift economy can only work in a small community where there is a sufficient collective knowledge of people's overall needs and contributions to keep the exchanges approximately in balance. This ability to survey the whole system—which Jean-François Noubel has called "holopticism" (Rousseaux, Saurel, & Petit, 2014)—enables a form of distributed intelligence that flexibly coordinates offer and demand. Our present global economy is much too complex for any individual or group to acquire such an overview of all the exchanges that occur. Nevertheless, thanks to the power of ICT the holoptic and balancing functions can be delegated to the emerging Global Brain, using the kind of production system algorithms we have discussed.

To make this practically feasible we need to implement two additional features of a gift economy. First, we need the equivalent of a *reputation system*, which keeps track of the aggregate quality and quantity of the offers that each individual has made through the offer network, increasing that person's reputation proportionally, while reducing it whenever the individual is observed to abuse the offer network. Computer-supported reputation systems have proven to be very effective in keeping vendors honest and constructive in online marketplaces (Jøsang, Ismail, & Boyd, 2007) and in motivating volunteers to make high-quality contributions to open collaboration networks (De Alfaro, Kulshreshtha, Pye, & Adler, 2011; Mamykina, Manoim, Mittal, Hripcsak, & Hartmann, 2011). While it will not be a trivial exercise to extend such systems to the level of a global offer network, it is likely that this will eventually happen through the same "disintermediation" dynamics that we have repeatedly encountered, in which local, mutually incompatible systems are eventually subsumed into a universal coordination protocol.

The second feature of a gift economy that needs to be added to a basic offer network is an initial surplus of gifts (i.e. unconditional offers) to produce the necessary buffer or slack that would make it easy to fulfill demands without precise matching of conditions. This would avoid the intrinsic difficulty of satisfying all the constraints of a conditional offer networks, where each demand is to be matched with a corresponding offer, which, however, first requires the satisfaction of another demand, and so on. With plenty of unconditional offers to fill in the gaps, balanced matches should be much easier to find. Unconditional offers would moreover allow people who have nothing to offer yet to build up the resourcefulness that will allow them to eventually reciprocate by offering valuable things themselves (Mullainathan & Shafir, 2013). This would create a virtuous cycle of "the poor getting richer".

There are several sources that can feed such a surplus of offers. The simplest one is that with our present level of wealth, people typically have more than they need. Most people would rather get rid of those clothes that do not fit anymore or those gadgets that were replaced by a more recent model by giving them to someone that can use them instead of accumulating them in their wardrobes or garages. A second source of free offers is *altruism*: most people are willing to unconditionally give things to others that are in need, as testified by the generous donations to charities, the enormous amount of volunteer work, and the success of crowdfunding initiatives where investors support a worthwhile project in return for a merely symbolic token of appreciation. Another source of surplus comes from the *non-rival* nature of information: giving information to others does not take anything away from the one who gives, because information can be infinitely duplicated without decreasing in value (Heylighen, 2007b). Thus, the immensely valuable knowledge repository of Wikipedia is available for free to anyone in the world, while the relatively small infrastructure behind it is maintained by volunteer efforts and donations. As the economy becomes increasingly dematerialized, the proportion of value carried by such informational goods can only grow (Heylighen, 2007b).

A final source of potential surplus comes from the productivity gains achieved through on-going technological innovation (Heylighen, 2008, 2014a)—especially through ICT-driven automation of production, logistics and services. As physical goods are efficiently produced by robots or 3-D printers controlled by software freely downloaded from the web, their cost plummets. This reduction of costs will be further boosted by the *Internet of Things*, an emerging suite of protocols and wireless controllers that promises to greatly enhance the efficiency of physical, logistic and industrial processes (Atzori, Iera, & Morabito, 2010). Moreover, the on-going shift to solar and other renewable energy sources that do not consume any resources will eventually make energy nearly free.

According to Jeremy Rifkin's (2014) analysis, these trends will culminate in a *zero-marginal cost society*, in which the cost of producing goods and services—and with it the value of money—essentially vanishes, once the investments needed to build the necessary infrastructure (robots, printers, solar panels, etc.) are recouped. That means that in a couple of decades there will be an abundance of offer relative to demand (Heylighen, 2014a), so that demands can be trivially satisfied. This will relegate the accumulation of money, which is the cornerstone of our capitalist economy, to the margins, because it becomes increasingly difficult and meaningless to make financial profit in a society where things are nearly free (Rifkin, 2014)

Tackling inequality through redistribution

The problem that remains at present is that the benefits of productivity increase are very unevenly spread, with the bulk going to the people that were rich and lucky enough to invest in successful innovations, while the majority may effectively become poorer because they lose their jobs, and thus their sources of income, to automation (Brynjolfsson & McAfee, 2014). Assuming that the first sources of surplus (superfluous goods, altruism, and free information) are not sufficient to fill the gap, this means that society will need to develop a better way to redistribute wealth.

Commonly proposed solutions are to impose taxes: on the one hand, on capital, capital gains, and/or financial speculation in order to dampen the "rich getting richer" feedback; on the other hand, on pollution (e.g. a carbon tax), consumption of non-renewable resources, and unhealthy lifestyles (e.g. tobacco, alcohol and sugar) in order to counter negative side effects of consumption. The income generated by such taxes would then be used to provide either welfare payments to the poor, or, perhaps more simply, a *universal basic income*, i.e. a guaranteed, unconditional payment made to all citizens so that they can afford the basic necessities they need for life, independently of any work they may or may not be doing (Van Parijs, 2004).

This scenario for combating inequality via redistribution is still based on the traditional notion of value as quantified by money. This means that it suffers from the shortcomings of money as a universal, one-dimensional exchange medium. For example, providing welfare or a basic income in the form of money in principle allows the recipient to exchange that money for drugs, alcohol, guns or other things whose overall value to society is negative. An offer network may avoid these problems by instead providing a "universal basic offer" that includes all and only the resources that people need to live a decent life—such as food, shelter, medical care and education. Non-essential resources, such as alcohol, fashion or jewelry, on the other hand, would only be available under more restrictive conditions.

Redistribution in this case would require that agents who have more resources than they need would somehow be driven to offer at least some of the "basic ones" for free. This could initially still be achieved via a monetary tax, which would be used by the government to buy the required surplus of basic resources. This surplus would be offered for free via the offer network, so as to make sure that they are distributed in a balanced matter. But eventually the system may shift from monetary taxation to taxation in the form of a "community service" that would be required from those who have more than they need. This service would consist in providing a basic resource that is in demand, and that the agent in question is skilled in providing. For example, a building firm could fulfill the service requirement by producing some amount of social housing for free—next to producing commercial property for sale (or in exchange for non-monetary offers).

Ideally, such a service should be voluntary rather than imposed by some central authority. This could be achieved by organizing the offer network in such a way that it would *mobilize* people to contribute of their own accord to the community. We already saw how reputation systems can motivate people to "give" things or services to others, so as to increase their status—as happened in the pre-monetary gift economies. But there are a range of other motives that can stimulate people to act towards the common good, including altruism, curiosity, desire for feedback, and sheer enthusiasm for a particular cause (Heylighen, 2007b). With the help of a user-friendly, "gamified" interface (Deterding, Dixon, Khaled, & Nacke, 2011) and the social-psychological mechanisms of reputation, flow (Nakamura & Csikszentmihalyi, 2002) and stigmergy (Heylighen, 2007b), a smart ICT system can stimulate and coordinate such motivations so as to produce effective collective action. This would turn the offer network into a *mobilization system* (Heylighen, Kostov, & Kiemen, 2013). That means that it would not only match existing offers and demands, but also *elicit* new offers for anything for which there is an unfulfilled need.

Dealing with externalities

Next to poverty and inequality, there is another fundamental shortcoming of a market economy that an intelligent offer network may be able to tackle: the problem of externalities. A market transaction will normally take place if both parties expect to derive some benefit out of it. However, a transaction often has benefits or costs for third parties that are not involved in the transaction. These consequences that are external to the transaction itself are called *externalities* (Cornes & Sandler, 1996). For example, if I sell my farm to a company that intends to build a factory on that land, both the company and I may consider this a good deal. However, the building of a factory will have repercussions for the people living around my former property. Some of these are negative (e.g. noise, pollution, additional traffic), some positive (e.g. employment opportunities, more business for local shops). The negative externalities amount to costs for the wider community, the positive externalities to benefits.

Everything taken together, it may well be that the balance is negative, and that the wider community would have been better off if the transaction had not taken place. But in a pure market economy, externalities are not taken into account, and therefore there is nothing to prevent transactions whose overall effect is negative. There is also no mechanism to promote transactions whose overall effect is positive, unless they bring benefit to the parties directly involved in the transaction. For example, I might have sold my land to a school that needs space to expand, but if their offer is lower than the offer of the factory, the transaction will not take place, even though school expansion would have been better for the community than factory building.

The traditional solution to this problem is for the government to tax negative externalities, such as pollution, and subsidize positive externalities, such as education—thus making globally beneficial transactions more attractive (Brynjolfsson & McAfee, 2014). The difficulty with this approach is that estimating the overall external impact of anything is extremely complex and situation-dependent, as it involves an unlimited number of known and potential effects and side effects. Taxation, on the other hand, needs to follow a set of simple rules, formalized as laws, that are set by a central authority. This only allows a coarse and rigid correction that is unable to adapt to the subtleties of a complex situation. Another difficulty is that the most important, long-term externalities, such as climate change or loss of biodiversity, are global, whereas the authority of governments is national and their outlook short-term. Therefore, governments typically have neither the power nor the political will to efficiently regulate such global externalities.

How could a global offer network deal with externalities? One advantage of an offer network is that it does not reduce demand to an aggregate of discrete transactions in which one party is willing to pay money to one other party in order to obtain some specific good or service. Demand should rather be seen as a collective preference for certain outcomes rather than other ones. This includes demands for non-monetizable values such as clean air, biodiversity, peace, or freedom. An intelligent, "holoptic" offer network should in principle be able to track to what degree individual or collective agents contribute to satisfying such demands. Agents that add to the most important values would be rewarded by a general increase in reputation, and therefore in their ability to satisfy their own demands. Agents that contribute negatively to values would be penalized by a loss of reputation, and therefore a loss in their ability to produce further harm. Thus, the network as a whole would motivate agents to seek positive externalities and avoid negative ones.

This mechanism already plays to some degree at the global level. For example, firms that are seen to participate in rain forest destruction or exploitation of low-wage laborers lose reputation and therefore potential income from clients who would rather buy from more conscientious competitors. Vice-versa, organizations that contribute to the global good, e.g. by building schools in poor African villages, gain in reputation and are rewarded with donations and volunteer work. At the moment, this propagation of reputation still happens via the intermediary of NGOs, government institutes or specialized groups, such as Amnesty International, Greenpeace or Unesco. These monitor otherwise obscure external effects, and use the mass media and the Internet to broadcast their observations so as to expand people's awareness of the global situation.

Because of the general disintermediation dynamics, it is likely that this monitoring and propagation of "external" reputation will become increasingly integrated with the universal protocols of the Global Brain. This means in particular that there would no longer be any independent authority, such as a government, NGO, or UN organization, to decide what is good or bad for society. Instead, the different "authorities" would themselves be subjected to the same kind of reputation mechanisms as individual agents. This makes it possible to compute a global reputation as an aggregate of evaluations by individual and organizational agents, weighted by their respective reputation scores. Such an aggregate reputation can be determined in an intelligent manner through the propagation of trust weights in a social network (Rodriguez et al., 2007; Rodriguez & Steinbock, 2006), using an equivalent of the PageRank algorithm (Page, Brin, Motwani, & Winograd, 1999; Rodriguez & Bollen, 2006). The result would be a distributed measure of trust or authority that would be more balanced, more diverse in its outlook, and less easily corruptible than any central authority (such as a government) in estimating to what degree the overall effects of an action would be positive or negative.

Promoting sustainability and resilience

A long-term effect of externalities is that they can make the system unsustainable, either by the accumulation of waste or pollution, or the exhaustion of non-renewable resources. A money-based economy can to some degree deal with the exploitation of scarce resources via the price mechanism: as a resource becomes scarcer, its price increases; therefore more effort will be done to save it, or to replace it by a more abundant resource (Heylighen, 1997). However, since waste does not have a price, it is ignored by the "invisible hand", and therefore the market has no incentive to reduce it. Taxing waste or pollution is a traditional way of reducing this problem.

However, an offer network can do better, by considering waste as an unconditional offer entered into the system. The network's coordination algorithms would not only try to maximize the satisfaction of demands ("backward chaining"), but also the exploitation of offers ("forward chaining"). That means that they should search to match this offer of "waste" with some sequence of production rules that may convert it into something for which there is a demand. For example, the grape seeds that are a waste product from wine making can be transformed into grape seed oil, which in turn can be used for the production of food, while plastic from packaging can be reconverted to produce inexpensive park benches and toys.

Thus, a well-balanced offer network should maximally recycle any "waste" that is produced as a side effect of satisfying demand. The realizability of full recycling is demonstrated by Chemical Organization Theory, a mathematical formalism based on "reactions" similar to the ones we described as production rules (Dittrich & Fenizio, 2007; Heylighen, Beigi, & Veloz, 2015). The theory shows that a network of reactions tends to self-organize into a self-sustaining subnetwork, where everything that is consumed is also produced, and vice versa. The algorithms supporting an intelligent offer network should help us to find such self-sustaining system of cycles more quickly and efficiently.

Another global challenge that offer networks may be able to tackle is *resilience*. Complex networks of interdependencies are vulnerable to *cascading failures* (Dueñas-Osorio & Vemuru, 2009; Helbing, 2013): whenever part of the network is damaged by some unforeseen perturbation (e.g. an earthquake, a bankruptcy, an electricity overload, a computer virus) this problem may start to propagate through the network more quickly than it can be contained, until the network as a whole breaks down. This is for example what happened with several large-scale electricity blackouts and with the 2008 financial collapse (Helbing, 2013). A system is resilient if it can recover from perturbations before it is irreversibly damaged. Strategies to achieve resilience include:

- preventing the positive feedbacks that make local disturbances cascade into global problems;
- increasing the diversity of the components of the system, so that a problem affecting one is unlikely to affect others;
- increasing redundancy, so that there are sufficient components to take over the function of those that were damaged; and
- accurately monitoring (potential) problems so that they can be remedied before they get out of hand.

All of these strategies are supported by offer networks. The envisaged surplus of offers creates a buffer against potential scarcity. The great diversity of available offers and methods to connect offers and demands provides sufficient redundancy so that demands can be satisfied even when the traditional supply routes are disrupted. This is similar to how the TCP/IP protocol made the Internet more resilient than rival networks, by providing a wide variety of alternative routes for retransmitting information packets that get lost. Moreover, the offer network would create a fine-grained map of the different needs the moment they appear (Gao et al., 2011), while providing intelligent algorithms that mobilize and route offers immediately to the place where the need is highest. Finally, because it does not rely on the accumulation of money, it prevents the positive feedbacks of speculation and the Matthew effect that can lead to financial collapses and growing inequality.

Implementing an intelligent offer network

We have noted several times that while the technologies for implementing an offer network in principle already exist, in practice such implementation will be a very complex issue. Implementation is likely to take several decades before it would reach the scale and efficiency necessary to eliminate the global problems of scarcity, inequality and sustainability. Let us here make the problem a little more concrete by noting some specific requirements and methods to tackle them.

An offer network functions as a *cognitive system*: it avoids searching in the physical world by performing these activities in the nearly frictionless, virtual world of knowledge and information. The essence of cognition is that the external world is represented internally, in the mind, so that problems can be solved by manipulating abstract concepts and rules rather than concrete objects and processes. This simplified representation functions as a model or "vicarious selector" (Campbell, 1987), which can be explored much more easily than the real environment. Thus, it leads to better solutions more quickly and with less effort and risk of error.

The network of condition-action rules, when properly represented, plays the role of such an abstract knowledge system. This defines the first requirement for a universal offer network: all users and supporting programs should use the same, unambiguous language for representing conditions and actions. If I need something specific from the offer network, I should be able to formulate that need so that the network can recognize that it matches the offer that someone else has formulated. Natural language tends to be too ambiguous for the purpose, because words can have many meanings or just be too vague to express what I need. Let us illustrate this and other requirements by means of an example of a situation in which offers and demands need to be matched.

Suppose that I am organizing a party (which is my offer) and that I would like it to be accompanied by live piano music (which is my demand). Before that demand can be satisfied, I will need to specify what I mean by music. This could for example lead the offer network to propose me a choice from different categories, such as jazz, blues, classical, or "any genre". Such a formal system of categories defines what is known in AI as an *ontology* (Hepp, 2007). Building consensual, comprehensive and accurate ontologies is a difficult and time-consuming task, which has up to now only been done for a relatively small number of domains. A universal offer network will in principle require an ontology that covers anything that could potentially be proposed or requested anywhere on Earth.

Assume that I chose the category "classical". Now we come to the second basic requirement, which is that the offer network should contain a very broad variety of real-world knowledge, i.e. production rules that tell it how the different categories are related, and in particular how the one in demand can be derived from the ones on offer. This knowledge will be mostly implemented in AI knowledge bases, but the less common rules may also still reside in the heads of the people that participate in the network. In this example, one rule that the network needs to know is the following: pianist + piano + playing → live piano music.

By backward chaining from this rule, the offer network can infer that my demand could be satisfied by locating (1) a classically trained pianist, (2) who is willing to play at my party, and (3) a piano. Thus, it decomposes a simple and general request into a combination of several more specific requests. The network now needs to search through its database for people that are listed in the category "classically trained pianist", and people or organizations that own pianos. Belonging to these ontologically defined categories are *hard constraints* that any potential solution to my demand must satisfy: you cannot produce live piano music without a piano or a piano player.

However, the decision about which of the potential offers to investigate will further depend on *soft constraints*—i.e. requirements where there is some leeway in determining what is acceptable and what is not. For a pianist, these "soft" criteria may include how good the pianist's play is and how easy it is for that pianist to come to my party. The quality of the play may be estimated from the pianist's reputation built up from playing at earlier events, while ease of travel may be inferred from the distance between the pianist's home and the location of the party.

Another constraint is whether the pianist would be willing to play at my party. The offer network may try to satisfy that constraint by *mobilizing* the pianist, i.e. by creating sufficient motivation for the pianist to do the effort. Possible motives that could convince the pianist to play may be an offer of free food and drink at the party and the fact that several friends of the pianist will be at the party. If this is combined with the fact that the party takes place two blocks away from the pianist's residence, at a free time, this may be sufficient to elicit a positive answer to the demand. The offer network can then make a list of potential pianists ordered according to the soft constraints, such as reputation for quality of play or degree of friendship with other party participants. It then asks each of them in turn whether they accept the invitation, while pointing out the benefits to them, until one says "yes".

In the case of the piano, the offer network may have located someone coming to the party who has a piano at home and is willing to lend it out for the party. However, the condition for realizing this offer is that a truck can be found to move the piano to and from the party location. This defines another attempt to match demand and offer, which now includes the additional hard constraint that the truck should be large enough to contain the piano, and the soft constraint that the truck should be located as close as possible to the location of the piano and/or the party. The offer network may find another partygoer who has a truck and moreover is willing to drive it, provided someone helps to load the piano in and out. Once a volunteer mover is found, the requirements for producing live piano music are fully met, and the party can go ahead—without any money needing to be exchanged in the process.

While the demand in this example is rather straightforward, in other cases the most difficult problem may be that there are no clear categories or criteria that distinguish offers that match well from offers that match less well. For example, if you are looking for an interesting book to read, you normally will not be able to formulate "interestingness" in terms of the typical categories used in a bookseller's ontology—such as genre, author, year or publisher. However, this problem is pretty well tackled by *recommender systems* (Ricci, Rokach, & Shapira, 2011). These typically use the popularity of items among people that in the past have expressed preferences similar to yours in order to recommend something that you too are likely to appreciate—a method known as *collaborative filtering* (Heylighen, 1999). Another powerful new technology for classifying and associating ill-structured data is "deep learning" and related AI techniques built on the induction of hierarchies of patterns (Hinton, Osindero, & Teh, 2006; Ray Kurzweil, 2012). These can create fuzzy ontologies, in which an item belongs to a category to a larger or smaller degree.

While there are plenty of theoretical and technological arguments for concluding that an intelligent offer network can be built over the coming years, the question remains whether such an network will work in practice, i.e. in an actual social system with people formulating concrete demands. Next to the technological challenge of formulating offers and demands in a way sufficiently precise so that computer programs can match them, there are the social, political and psychological challenges of getting people to abandon ingrained habits, powerful institutions and centuries-old traditions of economic exchange, while investing time and effort in a radically new system of sharing. The success of the many local sharing communities seems to indicate that this is possible, but it is likely that new difficulties will emerge when such systems try to scale up to a more universal and global level.

Therefore, a first step will be to set up smaller-scale experiments, which are still limited in the number of participants, but unlimited in the type of offers and demands that these can make, and more intelligent in their matching algorithms than existing sharing sites. Our research group, in collaboration with Goertzel's group and others, is planning to start such experiments over the coming years, while continuing to investigate the theoretical possibilities and limitations of the concept of offer network. A potential platform for implementing an offer network is the "web of needs" (Kleedorfer, Busch, Pichler, & Huemer, 2014), a proposed ICT system in which offers and demands are represented as "user proxies", i.e. formal descriptions of what a particular user needs or proposes.

Conclusion

The Global Brain concept denotes the idea that the Internet is connecting all people and machines on this planet into a distributed intelligence able to tackle the major problems that confront humanity. This emergent mind has an implicit system of values, a perception of its situation, a store of knowledge, and a capability for action. Its intelligence helps it to find the right sequence of actions that would allow it to maximize value in the present situation. Distribution means that this capability for perception, cognition, valuation and action is not centralized in a particular subsystem or agent, but spread across all human and technological components of society. Effective intelligence therefore requires strong *coordination* between the perceptions, inferences, values and actions of all these agents, so that they function in a harmonious, synergetic manner. In other words, their activities should be performed so that they efficiently complement each other, producing a maximum of value for a minimum of resources.

We have argued that the knowledge underlying such intelligence can be represented as an immense network of condition-action rules. The issues it addresses can be conceived as *challenges*, i.e. situations or aspects of situations that elicit action, because acting in this condition would produce a higher value than not acting. Challenges can be positive (opportunities, resources, offers) or negative (problems, needs, demands). The application of an appropriate condition-action rule will typically transform a challenge into a weaker, "relaxed" challenge (Heylighen, 2014b). The goal of the distributed intelligence is to find the combination of condition-action rules that would fully resolve a challenge, i.e. extract the value inherent in an offer or satisfy an unmet demand.

In a traditional society, challenges propagate in a rather slow and haphazard way from individual to individual along the links in a social network. In an ICT-supported society, a challenge can in principle be routed much more efficiently from the agents who experience it to the ones that are best positioned to resolve it. In recent years, a variety of platforms have sprung up across the Internet to implement this principle. Such platforms for example point people with a question towards an answer, people needing a ride towards a driver willing to take them along, or people having grown more fruit than they can eat to neighbors willing to collect that fruit. However, as yet these platforms are limited in the kind of challenges they cover, in the number of people they serve, and in the intelligence they exhibit. Moreover, the platforms are disparate—some commercial, some community-based, some open, some closed—while using a variety of rules, protocols, and interfaces. This strongly limits their effectiveness, because challenges cannot propagate across platforms in order to find an optimal resolution.

A similar situation arose in the early days of computer networking, when a variety of networks and standards made it nearly impossible to transmit information from one network to another. The TCP/IP protocol solved this problem by providing a robust, reliable and universal way to interconnect all local networks—thus creating the global Internet. The concept of an offer network (Goertzel, 2015) promises to do the same for the online economy: creating a universal protocol for sharing or exchanging any kind of goods, services or information.

The idea is that anyone with an offer or a demand can list this challenge on the network, including any conditions that may need to be satisfied before it can be realized. The system then uses AI algorithms to analyze the corresponding network of condition-action rules, taking into account both hard and soft constraints, in order to find an optimal match between offers and demands. Thus, it fulfills the function of distributed intelligence that defines the Global Brain. As more knowledge, offers and demands are expressed in such a network, while its memory and computing capabilities grow exponentially according to Moore's law (Kurzweil, 2005), the matching function should become ever more effective, finding highly efficient, synergetic and balanced solutions to complex and ill-structured problems involving numerous stakeholders.

Such an offer network would potentially revolutionize society, by eliminating the main shortcomings of our capitalist, money-based economy. Since offer networks can in principle allocate resources without the intermediate of money, they can promote values such as well-being, democracy, or biodiversity that cannot be expressed in monetary terms, or that function as mere "externalities" to a monetary transaction. This can be achieved by formulating these values as concrete demands and then mobilizing people or organizations to come up with offers that satisfy these demands. Agents offering effective solutions can be rewarded either directly by satisfying some of *their* demands in return for their investment, or indirectly, by increasing their reputation, so that their future demands are more likely to be met.

By promoting gifts, altruism, synergy, and efficiency, an offer network can elicit a surplus of resources, thus turning scarcity into abundance. The surplus of basic offers can be used to combat poverty and inequality, while the more efficient use and recycling of resources promotes sustainability. The diversity and redundancy of pathways connecting offers to demands makes the system more resilient against shocks and failures. Offer networks furthermore discourage the accumulation of the "universal resource" of money, thus averting the positive feedbacks that produce inequality and instability.

An offer network can in principle be implemented using existing ICT and AI algorithms, although its full deployment is likely to require many years, if not decades, of research and development. Given its many benefits, it seems likely that such a universal network will eventually take over and replace the highly inefficient patchwork of sharing platforms that presently exists. However, it will probably take many more years before the resulting system of exchange would be ready to replace our capitalist economy. In the meantime, money is likely to co-exist with different alternative currencies and other concrete resources within the network as one way among many to satisfy conditions or demands. While an offer network does not *need* money as an exchange medium, it can still profit from the flexibility of purely monetary offers for the time being—while the full power of the intelligent sharing network is still developing. Thus, as Rifkin (2014) argues, traditional capitalism will probably not disappear in the foreseeable future, but it is likely to become a niche phenomenon, on the margins of what Rifkin calls the *collaborative commons* and what we have called an offer network.

Once a universal offer network is fully deployed, it will form the foundation for what we have called the Global Brain: a distributed intelligence able to solve all problems confronting humanity—including poverty, inequality, scarcity, pollution, lack of sustainability, and poor resilience. The resulting state of peacefulness, freedom and abundance could be seen as a true return to "Eden" (Heylighen, 2014a), the idyllic state that humans putatively enjoyed as small communities of hunter-gatherers before the advent of civilization.

Acknowledgements

This research was supported by the Yuri Milner Foundation as part of the Global Brain Institute. Several of its ideas were inspired by discussions with Ben Goertzel, Jean-François Noubel, and the members of the Global Brain Institute.

References

Agatz, N., Erera, A., Savelsbergh, M., & Wang, X. (2012). Optimization for dynamic ride-sharing: A review. *European Journal of Operational Research*, *223*(2), 295–303. doi:10.1016/j.ejor.2012.05.028

Agrawal, A. K., Catalini, C., & Goldfarb, A. (2013). *Some simple economics of crowdfunding*. National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w19133

Anderson, J. R. (2014). *Rules of the Mind*. Psychology Press.

Atzori, L., Iera, A., & Morabito, G. (2010). The internet of things: A survey. *Computer Networks*, *54*(15), 2787–2805. Retrieved from http://www.sciencedirect.com/science/article/pii/S1389128610001568

Berners-Lee, T., & Fischetti, M. (1999). *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by Its Inventor*. Harper San Francisco. Retrieved from http://portal.acm.org/citation.cfm?id=554813

Brynjolfsson, E., & McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton & Company.

Campbell, D. T. (1987). Evolutionary epistemology. *Evolutionary epistemology, rationality, and the sociology of knowledge*, 47–89.

Cheal, D. J. (1988). *The Gift Economy*. Routledge.

Cornes, R., & Sandler, T. (1996). *The Theory of Externalities, Public Goods, and Club Goods*. Cambridge University Press.

De Alfaro, L., Kulshreshtha, A., Pye, I., & Adler, B. T. (2011). Reputation systems for open collaboration. *Communications of the ACM*, *54*(8), 81–87. Retrieved from http://dl.acm.org/citation.cfm?id=1978560

Deterding, S., Dixon, D., Khaled, R., & Nacke, L. (2011). From game design elements to gamefulness: defining "gamification." *Proceedings of the 15th International Academic MindTrek Conference: Envisioning Future Media Environments*, MindTrek ’11 (pp. 9–15). New York, NY, USA: ACM. doi:10.1145/2181037.2181040

Diamandis, P. H., & Kotler, S. (2012). *Abundance: The Future Is Better Than You Think*. Free Press.

Dittrich, P., & Fenizio, P. S. di. (2007). Chemical Organisation Theory. *Bulletin of Mathematical Biology*, *69*(4), 1199–1231. doi:10.1007/s11538-006-9130-8

Doan, A., Ramakrishnan, R., & Halevy, A. Y. (2011). Crowdsourcing systems on the world-wide web. *Communications of the ACM*, *54*(4), 86–96. Retrieved from http://dl.acm.org/citation.cfm?id=1924442

Dueñas-Osorio, L., & Vemuru, S. M. (2009). Cascading failures in complex infrastructure systems. *Structural Safety*, *31*(2), 157–167. doi:10.1016/j.strusafe.2008.06.007

Dugger, W. M., & Peach, J. T. (2015). *Economic Abundance: An Introduction*. Routledge.

Furuhata, M., Dessouky, M., Ordóñez, F., Brunet, M.-E., Wang, X., & Koenig, S. (2013). Ridesharing: The state-of-the-art and future directions. *Transportation Research Part B: Methodological*, *57*, 28–46. doi:10.1016/j.trb.2013.08.012

Gao, H., Barbier, G., & Goolsby, R. (2011). Harnessing the crowdsourcing power of social media for disaster relief. *Intelligent Systems, IEEE*, *26*(3), 10–14.

Gellman, R. (1996). Disintermediation and the internet. *Government Information Quarterly*, *13*(1), 1–8. doi:10.1016/S0740-624X(96)90002-7

Goertzel, B. (2002). *Creating internet intelligence: Wild computing, distributed digital consciousness, and the emerging global brain*. Kluwer Academic/Plenum Publishers.

Goertzel, B. (2015). Beyond Money: Offer networks, a potential infrastructure for a post-money economy. In B. Goertzel & T. Goertzel (Eds.), *The End of the Beginning: Life, Society and Economy on the Brink of the Singularity* (pp. 522–549). Humanity+ Press.

Graeber, D. (2014). *Debt: The First 5,000 Years*. Melville House. Retrieved from https://books.google.be/books?hl=en&lr=&id=lliTBQAAQBAJ&oi=fnd&pg=PT25&dq=david+graeber+gift+economy&ots=KKUYpdTYe7&sig=1rs9bfr5bIi6Gih566l8oIywML0

Helbing, D. (2013). Globally networked risks and how to respond. *Nature*, *497*(7447), 51–59. Retrieved from http://www.nature.com/nature/journal/v497/n7447/abs/nature12047.html

Helbing, D. (2014). *Qualified Money - A Better Financial System for the Future* (SSRN Scholarly Paper No. ID 2526022). Rochester, NY: Social Science Research Network. Retrieved from http://papers.ssrn.com/abstract=2526022

Hepp, M. (2007). Possible Ontologies—How Reality Constrains the Development of Relevant Ontologies. *IEEE Internet Computing*, 90–96.

Heylighen, F. (1997). The economy as a distributed, learning control system. *Communication & Cognition-AI*, *13*(2/3), 207–24. Retrieved from http://pcp.vub.ac.be/papers/MarketControl.pdf

Heylighen, F. (1999). Collective Intelligence and its Implementation on the Web: algorithms to develop a collective mental map. *Computational & Mathematical Organization Theory*, *5*(3), 253–280. doi:10.1023/A:1009690407292

Heylighen, F. (2002). The global brain as a new utopia. *Zukunftsfiguren*. Suhrkamp, Frankurt. Retrieved from http://pespmc1.vub.ac.be/papers/GB-Utopia.pdf

Heylighen, F. (2007a). The Global Superorganism: an evolutionary-cybernetic model of the emerging network society. *Social Evolution & History*, *6*(1), 58–119. Retrieved from http://pcp.vub.ac.be/papers/Superorganism.pdf

Heylighen, F. (2007b). Why is Open Access Development so Successful? Stigmergic organization and the economics of information. In B. Lutterbeck, M. Baerwolff & R. A. Gehring (Ed.), *Open Source Jahrbuch 2007* (pp. 165–180). Lehmanns Media. Retrieved from http://pespmc1.vub.ac.be/Papers/OpenSourceStigmergy.pdf

Heylighen, F. (2008). Accelerating socio-technological evolution: from ephemeralization and stigmergy to the global brain. *Globalization as evolutionary process: modeling global change*, Rethinking globalizations (p. 284). Routledge. Retrieved from http://pcp.vub.ac.be/papers/AcceleratingEvolution.pdf

Heylighen, F. (2012). *A Tale of Challenge, Adventure and Mystery: towards an agent-based unification of narrative and scientific models of behavior* (ECCO Working Papers No. 2012-06). ECCO. Brussels, Belgium. Retrieved from http://pcp.vub.ac.be/papers/TaleofAdventure.pdf

Heylighen, F. (2014b). Challenge Propagation: Towards a theory of distributed intelligence and the global brain. *Spanda Journal*, *V*(2), 51–63. Retrieved from http://cleamc11.vub.ac.be/papers/ChallengePropagation-Spanda.pdf

Heylighen, F. (2014c). *Cognitive Systems: a cybernetic approach on the new science of the mind* (Lecture Notes.). ECCO, VUB. Retrieved from http://pespmc1.vub.ac.be/Papers/CognitiveSystems.pdf

Heylighen, F. (2015a). Return to Eden? Promises and Perils on the Road to a Global Superintelligence. In B. Goertzel & T. Goertzel (Eds.), *The End of the Beginning: Life, Society and Economy on the Brink of the Singularity*. Humanity+ Press. Retrieved from http://pespmc1.vub.ac.be/Papers/BrinkofSingularity.pdf

Heylighen, F., Beigi, S., & Veloz, T. (2015). *Chemical Organization Theory as a universal modeling framework for interaction, self-organization, and autopoiesis* (ECCO Working Papers). Retrieved from http://pespmc1.vub.ac.be/Papers/COT-applicationsurvey.pdf

Heylighen, F., & Bollen, J. (1996). The World-Wide Web as a Super-Brain: from metaphor to model. *Cybernetics and Systems' 96. R. Trappl (Ed.).* Austrian Society For Cybernetics. Retrieved from http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.69.5617

Heylighen, F., Busseniers, E., Veitas, V., Vidal, C., & Weinbaum, D. R. (2012). *Foundations for a Mathematical Model of the Global Brain: architecture, components, and specifications* (GBI Working Papers No. 2012-05). GBI Working Papers. Retrieved from http://pespmc1.vub.ac.be/papers/TowardsGB-model.pdf

Heylighen, F., Kostov, I., & Kiemen, M. (2013). Mobilization Systems: technologies for motivating and coordinating human action. In Peters M. A., Besley T. and Araya D. (Ed.), *The New Development Paradigm: Education, Knowledge Economy and Digital Futures*. Routledge. Retrieved from http://pcp.vub.ac.be/Papers/MobilizationSystems.pdf

Hinton, G. E., Osindero, S., & Teh, Y.-W. (2006). A Fast Learning Algorithm for Deep Belief Nets. *Neural Computation*, *18*(7), 1527–1554. doi:10.1162/neco.2006.18.7.1527

Holland, J. H. (2012). *Signals and Boundaries: Building Blocks for Complex Adaptive Systems*. Cambridge, Mass: The MIT Press.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1989). *Induction: processes of inference, learning, and discovery*. MIT Press.

Hutchins, E. (2000). Distributed cognition. In N. J. Smelser & P. B. Baltes (Eds.), *International Encyclopedia of the Social and Behavioral Sciences*. Elsevier Science.

Jøsang, A., Ismail, R., & Boyd, C. (2007). A survey of trust and reputation systems for online service provision. *Decision Support Systems*, Emerging Issues in Collaborative Commerce, *43*(2), 618–644. doi:10.1016/j.dss.2005.05.019

Kleedorfer, F., Busch, C. M., Pichler, C., & Huemer, C. (2014). The Case for the Web of Needs. *2014 IEEE 16th Conference on Business Informatics (CBI)* (Vol. 1, pp. 94–101). Presented at the 2014 IEEE 16th Conference on Business Informatics (CBI). doi:10.1109/CBI.2014.55

Kurzweil, R. (2005). *The singularity is near*. Penguin books.

Kurzweil, Ray. (2012). *How to Create a Mind: The Secret of Human Thought Revealed*. Penguin.

Last, C. (2014). Global Brain and the Future of Human Society. *World Future Review*, 1946756714533207. doi:10.1177/1946756714533207

Last, C. (2016). Self-actualization in the Commons. *Technological Forecasting and Social Change*.

Mamykina, L., Manoim, B., Mittal, M., Hripcsak, G., & Hartmann, B. (2011). Design lessons from the fastest q&a site in the west. *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11 (pp. 2857–2866). New York, NY, USA: ACM. doi:10.1145/1978942.1979366

Mayer-Kress, G., & Barczys, C. (1995). The global brain as an emergent structure from the Worldwide Computing Network, and its implications for modeling. *The information society*, *11*(1), 1–27.

McLeod, P., Plunkett, K., & Rolls, E. T. (1998). *Introduction to connectionist modelling of cognitive processes*. Oxford University Press Oxford.

Minsky, M. (1988). *The society of mind*. Simon & Schuster.

Mollick, E. (2014). The dynamics of crowdfunding: An exploratory study. *Journal of Business Venturing*, *29*(1), 1–16. Retrieved from http://www.sciencedirect.com/science/article/pii/S088390261300058X

Mullainathan, S., & Shafir, E. (2013). *Scarcity: Why Having Too Little Means So Much*. Macmillan.

Nakamura, J., & Csikszentmihalyi, M. (2002). The concept of flow. In C. R. Snyder (Ed.), *Handbook of positive psychology* (pp. 89–105). New York, NY: Oxford University Press.

Page, L., Brin, S., Motwani, R., & Winograd, T. (1999). The PageRank citation ranking: bringing order to the Web. Retrieved from http://ilpubs.stanford.edu:8090/422

Piketty, T. (2014). *Capital in the Twenty-First Century*. (A. Goldhammer, Trans.). Cambridge Massachusetts: Belknap Press.

Ricci, F., Rokach, L., & Shapira, B. (2011). Introduction to Recommender Systems Handbook. In F. Ricci, L. Rokach, B. Shapira, & P. B. Kantor (Eds.), *Recommender Systems Handbook* (pp. 1–35). Springer US. Retrieved from http://link.springer.com/chapter/10.1007/978-0-387-85820-3_1

Richey, S. (2007). Manufacturing Trust: Community Currencies and the Creation of Social Capital. *Political Behavior*, *29*(1), 69–88. doi:10.1007/s11109-007-9028-7

Rifkin, J. (2014). *The Zero Marginal Cost Society: The Internet of Things, the Collaborative Commons, and the Eclipse of Capitalism*. Palgrave Macmillan. Retrieved from http://www.bookdepository.com/Zero-Marginal-Cost-Society-Jeremy-Rifkin/9781137278463

Rodriguez, M. A., & Bollen, J. (2006). Simulating network influence algorithms using particle-swarms: Pagerank and pagerank-priors. *Arxiv preprint cs/0602002*.

Rodriguez, M. A., & Steinbock, D. J. (2006). The anatomy of a large scale collective decision making system. *Los Alamos National Laboratory Technical Report [LA-UR-06-2139], Tech. Rep.*

Rodriguez, M. A., Steinbock, D. J., Watkins, J. H., Gershenson, C., Bollen, J., Grey, V., & deGraf, B. (2007). Smartocracy: Social Networks for Collective Decision Making (p. 90b). Los Alamitos, CA, USA: IEEE Computer Society. doi:http://doi.ieeecomputersociety.org/10.1109/HICSS.2007.484

Rousseaux, F., Saurel, P., & Petit, J. (2014). Knowledge Engineering or Digital Humanities? In C. Faucher & L. C. Jain (Eds.), *Innovations in Intelligent Machines-4* (Vol. 514, pp. 129–187). Cham: Springer International Publishing. Retrieved from http://link.springer.com/10.1007/978-3-319-01866-9_5

Seyfang, G. (2001). Money that makes a change: Community currencies, North and South. *Gender & Development*, *9*(1), 60–69. doi:10.1080/13552070127737

Van Parijs, P. (2004). Basic Income: A Simple and Powerful Idea for the Twenty-First Century. *Politics & Society*, *32*(1), 7–39. doi:10.1177/0032329203261095

Woergoetter, F., & Porr, B. (2008). Reinforcement learning. *Scholarpedia*, *3*(3), 1448. doi:10.4249/scholarpedia.1448