

## *Human Stigmergic Problem Solving*

### **6.1 What Is Stigmergic Problem Solving?**

#### *6.1.1 Background*

What is stigmergy? The French entomologist Grassé coined the term “stigmergy” in the 1950s. The term is formed from the Greek words stigma “sign” and ergon “action,” referring to individual actions that leave signs in the environment, and determine subsequent actions by others. Stigmergy usually describes how many individual agents are able to coordinate collective action only by leaving information in a shared environment (Parunak, 2005). The basic principle of stigmergy is extremely simple; traces left by agents in the environment provide feedback information to new agents (Theraulaz & Bonabeau, 1999). When one agent leaves a trace in the environment, this trace will even stimulate or motivate others to do subsequent work. The aggregated collective work serves the purpose of being externalized information that ensures that new tasks are executed in the right order. The complete solution will gradually emerge when different individuals interact with the “evolving information” in the environment at different points of time (Rezgui & Crowston, 2018).

Stigmergy can also be explained as a feedback loop that does not require any direct communication between the individuals because all coordination is done through the traces of information left in the medium. When information remains available, it can guide new agents at any later point of time, and there is no need to be present at the same time. Nor is mutual awareness a requirement since every individual works independently of each other. The individuals do not even need to know that other agents are participating in the work. The collective actions are materialized in the environment and function like a shared external collective memory (Heylighen, 2011, 2015). For example, an ant colony will record its collective activity as traces in the physical environment, and this helps

them organize their collective behavior. Information can be stored in the environment in several different ways, as gradients of pheromones, material structures impregnated by chemical compounds, or by spatial distribution of colony elements. These traces of the collective activity function as “stimulants,” both by directing and constraining the individual behavior of the ants. New actions are triggered by the perceived recent changes in the trace (Theraulaz & Bonabeau, 1999).

The notion of stigmergy allowed Grassé to solve the “coordination paradox,” the question of how social insects could collectively tackle complex projects like building a nest. The notion of stigmergy highlights that it is possible to generate robust, intelligent behavior at a system level by following very simple behavioral rules. Compared to traditional methods of organization, stigmergy makes minimal demands on the agents. There is no need for a plan or overall goal, individuals only need to know the present state of the activity (Heylighen, 2015; Parunak, 2005; Rezgui & Crowston, 2018). The two basic requirements in stigmergy are that the agents can recognize the right conditions to start their work, and that they can access the medium in which these conditions are registered. These agents are goal-orientated in their attempt to maximize “fitness,” “utility,” or “preference.” The underlying mechanism is local trial and error or variation and selection, where two interacting agents mutually adapt their actions, until they reach an acceptable “coordinated” pattern. This local pattern is then adopted by neighboring agents until it includes the whole system. A global order will spontaneously emerge out of local actions, illustrating that intelligence does not reside in each individual agent, but in the interactions among the agents and the shared dynamical environment (Heylighen, 2015; Parunak, 2005; Rezgui & Crowston, 2018). It has even been suggested that stigmergy is the only way a large distributed population can solve collective problems if it has a limited amount of computational resources (Parunak, 2005).

If one observes each insect separately in a colony, they do not seem to be involved in a coordinated, collective behavior. However, they interact indirectly through medium, and both physical and geometrical constraints will influence the choices of the colony. Social insects use a large variety of olfactory, tactile, visual, and vibrational messages, as well as multi-modal combinations of these in their communication. In general, these messages can be divided into three groups. First, some messages require direct contact between individuals, being local in both space and time. Second, some signals are local in time but not space, typically alarm signals. Third, some messages, building on stigmergy, are local in space but not in time (Feinerman & Korman, 2017).

### 6.1.2 Quantitative Stigmergy

Moreover, in stigmergy it is common to distinguish between quantitative and qualitative stigmergy. Quantitative stigmergy refers to perceived conditions that differ in strength or degree, and where stronger traces typically elicit more intense or frequent actions. Two or more actions are performed on the same object or task, and the stimulus-response sequence comprises stimuli that do not differ qualitatively from each other like gradients in pheromone fields. It is only the probability of others performing the same action that will change. A stronger trace will, over time, lead to more frequent actions by increasing the number of individual contributions, resulting in a more intense overall activity. Both ant trail laying and termite nest building use quantitative stigmergy (Heylighen, 2011, 2015; Theraulaz & Bonabeau, 1999).

In ant trail laying, ants discover a range of different food sources independently, and stigmergic mechanisms effectively select the closest source to the nest. If the path to the food source is short, the traffic will be sufficiently intense for the pheromonal trace to remain. However, if the distance between a food source and the nest is long, the time interval between the trips of two foragers may exceed the evaporation latency of the pheromone and the trail disappears. The time scale of pheromonal communication will depend on chemical evaporation times and vary between species and tasks. Depending on the distribution of food, the same behavioral rules may produce very different collective behavior. Individual ants will refine and amplify complete ant trail structures that other ants have made (Feinerman & Korman, 2017; Theraulaz & Bonabeau, 1999) (Figure 6.1).

Another example is termite nest building, which is performed without a plan (Figure 6.2). The individuals will locally interact with features of the structure by adding building material to them. Termites may add mud to the same pillar, or several individuals will push the same load of mud. Soil pellets impregnated with pheromone are used to build the pillars in two phases. First, pellets are deposited randomly until one of the deposits reaches a critical size. If the group of builders is sufficiently large, the coordination phase starts, and pillars or strips begin to emerge. The higher the emerging heap of mud is, the stronger the trace will be. This makes it even more attractive, strengthening the probability of more mud being added, thus creating an amplifying effect. The workers are stimulated to continue the building process through a positive feedback mechanism (the snowball effect), since the increasing amount and accumulation of material



Figure 6.1 Leafcutter ants following the same trail when carrying leaves back to the nest, photo Ricardo Riechelmann/EyeEm/Getty Images ©



Figure 6.2 Cathedral termite mounds near Adelaide River, Northern Territory, Australia. The termite mound structures are approximately 100 years old and can stand up to seven meters in height. The mounds are made with a combination of soil, mud, chewed wood, and saliva. The life of the termite is a constant race against rain because a heavy downpour can ruin part of the mounds. Therefore, the termites will always be rebuilding their mounds, photo Yvonne van der Horst/Getty Images ©

reinforces the attractiveness of deposits through the pheromone on the pellets. A spatiotemporal structure emerges from a random spatial distribution of soil pellets (Feinerman & Korman, 2017; Heylighen, 2011, 2015; Theraulaz & Bonabeau, 1999).

Once the structures are created, they are stabilized through negative feedback, mainly pheromone decay and competition among neighboring pillars. When pheromones are deposited on the building materials, it adds a temporal dimension to the physical structure because pheromones evaporate over time. If the number of builders become too small, the pheromone disappears and the amplification mechanism stops. This self-organizing system can result in several different stable states. In the ant trail laying, this depends on the initial conditions (path dependency). The collective behavior can also change completely at a critical density, like at some point in the nest building; no pillars will emerge below the construction, only above it. When building and extending a termite hill, no final goal will be reached as the maintenance will be part of an ongoing long-term project (Feinerman & Korman, 2017; Heylighen, 2011, 2015; Theraulaz & Bonabeau, 1999).

### 6.1.3 Qualitative Stigmergy

Qualitative stigmergy refers to conditions and actions that differ in kind rather than in degree. It differs from quantitative stigmergy in that individuals respond to qualitative differences in the type of stimuli; a different trace will stimulate a different type of action. Actions are performed automatically in the right order, since an action will not be started until the right condition is in place.

In nest construction in the solitary wasp *Paralastor sp*, it is the completion of one stage that provides the stimulus that begins the next stage (Figure 6.3). The wasp begins with the excavation of a narrow hole, and only when the nest hole has been completely lined with mud will the wasp begin the construction of a mud funnel above its entrance. This funnel is built in five distinct and highly stereotyped stages from a series of mud pellets. Stage 1 involves the building and application of a series of mud pellets until it reaches a length of 3 cm. At Stage 2, the wasp ceases to build upwards, and by adding more mud to one side, it begins the construction of a uniform curve. The end of each stage of building will provide the stimuli that initiates the next building stage. Coordination of the collective behavior will always build on the previous consequences of building actions (Theraulaz & Bonabeau, 1999).



Figure 6.3 Potter wasp building mud nest, São Paulo, Brazil, photo Kassá/Getty Images ©

A second example is nest building in the wasp *Polistes*. Building decisions are based on perceived configurations of previous construction that direct the collective work. Here, several building actions may happen in parallel, and actions will be performed on separate, independent parts of the medium. One advantage is that different wasps can then do the same work (Figure 6.4). Nor does a wasp distinguish between its own or others' work, making it possible for a new wasp to continue with the ongoing nest construction work at any stage. This can potentially result in conflicting actions when they are performed simultaneously. However, cells are added according to specific simple buildings rules, such as the rule in which the wasps tend to finish a row of cells before initiating a new row. For example, the probability of adding a cell to a three-wall site is about ten times higher than in the case of a two-wall site.

Experimental studies have shown that if a stimulus is presented at the “wrong stage,” it automatically leads to a redundant structure. It is the architecture itself that constrains the building activity and prevents its disorganization. Since the sequence of the tasks is imposed, a single individual does not need an overview of the complete of task. A new cell is added based on several templates that best characterize the current local shape of the nest (Theraulaz & Bonabeau, 1999).



Figure 6.4 Two hornets building a nest together by making hexagonal cells, Mana, HI, United States, photo Craig Damlo/Getty images ©

#### 6.1.4 Sematectonic Stigmergy

Furthermore, it is possible to distinguish between sematectonic and marker-based stigmergy. In sematectonic stigmergy, the structure of the domain itself will provide sufficient signals for coordinating the collective behavior without any need for special markers. A new action is triggered by “the current state to the solution” which is the accumulated activities of prior agents. This mechanism is present in the aforementioned example of nest building in the wasps *Paralstor sp* and *Polistes*. Another example is ants who cluster corpses in their cemeteries, directed by variations in the density of the corpse distribution. This is quantitative sematectonic stigmergy, where collective decisions are made by individuals who follow gradients in this field, avoiding repellers and approaching components (Parunak, 2005). A classic example of quantitative sematectonic stigmergy is the creation of human trails. Humans wear down vegetation on routes that are used frequently, while the vegetation will grow again if an old path is not used. This example illustrates that all actions count whether you choose to walk along the path or not (Parunak, 2005).

*6.1.5 Marker-Based Stigmergy*

In marker-based stigmergy, coordination is built around a signaling mechanism where a mark is explicitly left with the intention of being a signal. Unlike sematectonic stigmergy, which is a response to an environmental modification, marker-based stigmergy does not make any direct contribution to the work in it itself. Here, ant trail laying is a relevant example, in which ants leave a trace of pheromones, a chemical signal, as marks when they return to the nest after they have found food. Each ant has the disposition to move towards the scent left by other ants. When there are few ants, the scent has little effect on the collective behavior of the group. However, when there are many ants, and each of them moves towards the strongest scent and at the same time lays down their own scent, trails of scent to the food begin to emerge. So the stronger the pheromone trail, the larger the probability of a response. Gradually, an extensive network of pheromone trails will connect the nest to the surrounding food sources. Shorter paths collect more pheromone, so the network and the collective “external memory” will continuously become more effective (Heylighen, 2013; Marsh & Onof, 2008).

Another example is how animals leave marks in the terrain, often communicating their presence to other conspecifics. The most basic message is a sign that informs others that, “I’m here.” This mark will typically give information about the identity of the animal and its relative dominance (Giuggioli, Potts, Rubenstein, & Levin, 2013). A number of mammals will be marking the environment with glandular secretions, urine, or feces (Figure 6.5). Besides the function of informing other conspecifics, marks may help individuals orient themselves in their area.

In animal groups, marking is also done by the whole group in their area. They leave more marks at much-visited sites such as junctions, dens, or zones where individuals from other groups may be encountered (Theraulaz & Bonabeau, 1999). Likewise, human marker-based stigmergy will leave some kind of signal that informs others about their actions. For example, humans will tend to improve trails with markers such as direction signs, illustrating how sematectonic stigmergy is often combined with marker-based stigmergy that reinforces the preferred actions (Parunak, 2005).

*6.1.6 Human Stigmergic Problem Solving Is Solution-Centered*

Obviously, stigmergic mechanisms will be present in many types of human problem solving. As with insects, human individuals do not need to have a





Figure 6.5 Cheetahs scent marking their territory together, Masai Mara in Kenya, photo Mike Powles/Getty images ©

complete overview of the work, the problem, or the solution. According to Rezgui and Crowsten (2018), information on how to improve human collective work can also be communicated through the current state of the work itself. Coordination signals can be elicited from the ongoing shared work. Tasks depend on each other and build on previous work when it has been stored. The implication is that problems depend on the current status of the solution.

Parunak (2005) suggests the binary distinctions in quantitative vs. qualitative stigmergy, and sematectonic vs. marker-based stigmergy, can also be used in the analysis of different types of human stigmergy. In quantitative stigmergy, signals follow a single scalar, whereas in qualitative stigmergy, the signals form a set of discrete options. In marker-based stigmergy, the signs consist of special markers that agents deposit in the environment, while in sematectonic stigmergy, individual actions respond to the current state of the solution (Parunak, 2005).

In this chapter, the analysis is inspired by these stigmergic dimensions. For example, in human qualitative stigmergy, a preliminary part of a solution will be stored in the system or medium, and individuals will then respond to the unfinishedness in the solution in different ways. New

actions will be triggered in an attempt to come closer to a solution. If many versions of a solution already exist, human quantitative stigmergy can be used to rate the most optimal solutions. In the online setting, solutions will continuously be compared with each other. Here, the aggregation of a large number of individual reviews will help direct attention toward the best solutions.

Furthermore, the chapter suggests that human stigmergic problem solving is primarily solution-centered and emerges through four distinct problem-solving mechanisms:

1. Rating complete solutions
2. Reestimating the solution
3. Completing solutions
4. Adapting complete solutions

These four subtypes of stigmergic problem solving are further explained, both through the introduction of new examples and by analyzing previous examples from Chapter 3. In the final section, the four problem-solving mechanisms are related to the different stigmergic dimensions (quantitative vs. qualitative, sematectonic vs. marker-based).

## 6.2 Rating Complete Solutions

### 6.2.1 Search Engines and Collaborative Filtering

Because of digitization, more information is available than ever before. In Chapter 3, several examples showed how knowledge products or “complete solutions” are being published openly, including research articles, open textbooks, and videos (e.g., YouTube). These solutions are typically complete in the sense that they are of direct value without any need for further modification. Quite a lot of knowledge products, like videos published on the online platform YouTube, are automatically included in a larger collection of similar types of work. Most knowledge products will remain available for a very long time. For example, when an instructional video is published to help a specific individual in an online community, it can also be relevant to other new viewers at any later point of time. However, because new knowledge products are being published constantly and everything is stored, this increase in published solutions need to be filtered. The sheer size is overwhelming and requires new ways of sorting and finding relevant information. Key questions are: How do we decide what we need to know, how do we find solutions that will

address our need, and how do we evaluate what solutions are most relevant?

Today, it is the search engines that help us find and reuse the best solutions others have made. They select and rank solutions, and ultimately define what knowledge is of value today. The web became search-centric in the mid-2000s. Although people still use the phrase “surfing the Internet,” they no longer move from site to site through hyperlinks like in the 1990s (Halavais, 2018: 30–31). Search engines are built around three components – the crawler, the indexer, and the front end. First, information about webpages is gathered from around the web. Second, all this collected data is evaluated according to their relevance to a particular set of keywords. Almost every modern search engine extracts key terms to create a keyword index of the web by an “indexer” (Halavais, 2018: 19). In the index of a printed book, one gets an overview of all the page numbers where all the keywords appear in the book. On the web, this strategy is not effective because much of the material on the web may be spam and intentionally misleading. Search results need to rank content according to how relevant it is. Algorithms will typically rely on hyperlinks to assess how relevant specific web pages are according to a search term. This is used to create an index, the “secret recipe” of a search engine, which crawls the web and creates a database of indexed material, which every individual search will be built upon. Third, in the final step, the search engine will present the processed results of a query at the “front end.” The results are usually ranked in a simple list that display the most significant hits, providing a clear prioritization of the most relevant solutions. The results are designed to reveal possibilities without overwhelming the user (Halavais, 2018: 19–25).

Google became the most popular search engine because they were the first to recognize that the number of links could be used as an estimation of how good a web page was. When one page linked to another page, it indicated that the content was worth reading. Hyperlinks were not just connections, but rather votes on the relevance of a web page. A large number of links indicated that a page was particularly relevant. For example, according to one estimation, Google claims the median number of backlinks to the first three sites to be 1,210, while the hundredth site receives 61 backlinks, and the last site receives none. However, not all backlinks are necessarily of equal value. Therefore, Google assigned a PageRank to each page in the search engine’s index, calculating the number of sites pointing to the page as well as estimating the “popularity” of these sites by also calculating the number of links these sites had. More

weight is given to links from popular sites. PageRank provides a way for the web community to tacitly vote on the quality of a page. Google also extracts keyword information from the text on the backlinks, so that relevant content on these pages have an impact too. For example, if a web page is linked to from a site with many visitors; this popular page is “weightier” than other backlinks (Halavais, 2018: 108–116). If somebody publishes a solution, it will often be located by a search engine, and when more people link to an answer because they find it relevant, it is ranked higher in the search results.

In large part, the design of search engines is determined by our limited span of attention. Although there has been an enormous increase in the production and access to information and knowledge, the human capacity to consume information is still the same (Halavais, 2018: 112–114). Because of the existence of billions of webpages, it is not possible to present them all equally. Therefore, search is as much a process of ignoring as it is of presenting. Gaining attention has always been important, and the search engine is changing the ways in which attention is concentrated and distributed (Halavais, 2018: 99). Over a period of two centuries, mass media like radio, newspapers, and television sent out the same message to the public, setting the agenda and directing people’s attention to what was considered worthwhile information. However, on the web, the attention becomes distributed because of all the options that are available. Because of the increased use of participatory media, the fight for people’s attention is now harder than ever before. Attention is increasingly regarded as a valuable commodity. The visitor of a website is someone who gives you attention in return for information. In addition, this ability to consistently attract attention is something advertisers want because it can produce revenue. Here, the search engine is the tool that select some of these winners who get a lot of attention. In one example, an online diamond retailer called Skyfacet.com was selling jewelry worth about 3 million each year, but a change in Google’s ranking algorithms removed the site off the first few pages of results for popular queries about jewelry. The result was that sales dropped by 17 percent in three months. While sales in an offline setting is about location, in the online setting, it is all about getting attention through search engines (Halavais, 2018: 112–114).

In the recent decade, data from social media platforms have also increasingly been used to inform and rank search results. In “social search” or search personalization, the search engine results are re-ranked according to that person’s search history and the interests of the person’s larger social network. In addition, people increasingly search for information via social

media platforms instead of using open search engines. Platforms like Facebook and LinkedIn filter information from the rest of the web. For example, in a study in 2014, equally many found and read news articles on Facebook as on open-ended search engines. On social media, users not only find information, but they can identify who shared the information and connect with this person if they want to. The platforms even offer their own search engines, and in 2016, Facebook handled more than 2 billion searches per day from their search box (Halavais, 2018: 76–78, 89–90).

On one hand, most information is ranked automatically through our collective behavior on the Internet, but in addition, many systems today let individuals actively rate complete solutions. For example, there are dedicated Community Question Answering (CQA) sites, markets for answering questions that involve people with particular expertise. Some examples are Stack Overflow, Quora, and Yahoo! Answers. Generally, these sites allow individuals to ask a question, and anyone can rate the various answers that are provided. Points can also be awarded, and the market is largely reputation-based. All of these systems attempt to gain expert opinion without the costs of finding and hiring an expert. For a single question, that expense is frequently overwhelming. In many cases, the questions are practical and quite simple to answer (Halavais, 2018: 92–93). These sites build on the assumption that many individuals pose the same questions at different locations, but with the help of search engines, one can save time by effectively finding and reuse solutions that others already have provided. However, if a problem is difficult to understand or conceptualize, it may be much easier to ask an expert directly,

Another type of “sociable search” centers on the aggregation of explicit user-centered evaluations of content. They index ratings or measure how a community evaluates a web page instead of indexing the web pages. For instance, collaborative filtering harness information from explicit reviews performed by a large group of users who visit a site. Those who share similar interests can discover relevant material with very little individual effort. A “front page” of options is presented based on criteria provided by the user. Preferences in one area gives information about other topics that may be of interest, and the aggregated opinions of similar peers guide the user to new relevant resources. Homophily is one label that is used to describe the tendency of like-minded people to “flock together.” If two people have similar backgrounds and preferences on some matters, they are likely to have similar tastes in other areas, and this can easily lead to a self-reinforcing congruence. By explicitly taking into account the searchers’

social networks, results can be even more effectively ranked (Halavais, 2018: 82–88).

Reddit is one example of a collaborative filter site that explicitly let user votes decide what web pages get the most attention. This site covers a large variety of topics that allow the community to vote on the most interesting links and comments. Although Reddit does not actively seek to establish an online community, it is possible to publish comments. However, a major challenge at the site is how to reduce the “groupthink” of group filters, and limit unintended or intended bias. The algorithms that determine what appears on the front page need to be regularly changed because of attempts to manipulate the feed. One alternative strategy to cope with spam and bad, off-topic comments is to let users vote on comments. The website Slashdot even created a meta-moderation process that permitted voting on the moderators’ votes. Every search engine depends on trust because the searchers will always perceive a threat of being either intentionally or unintentionally misled. In traditional search engines, this will always be a challenge because the user will not know how the search engine works. However, in a system like Reddit, the transparency of all the activities creates trust. Being aware of how others use the system increases the overall trust in the system (Halavais, 2018: 84–88).

### 6.2.2 *Different Rating Methods*

The quality of complete solutions published online will be closely connected to how others value or rate the work. When a knowledge product is reused by others, the digital traces will always aggregate simple statistical meta-information, such as the number of visits. Many online sites also let users actively rate solutions in different ways, but typically with simple voting methods. For example, features such as subscription counts, ranking, likes, and dislike counts in YouTube provide an indication of the popularity of posted content and how well it is received (Lee et al., 2017). Systems will usually allow any viewer to rate a video by registering a “like” or “dislike” rating by using the “thumbs up” approval icon or the “thumbs down” approval icon. Viewers can also favorite videos as a way of “spreading the word” on videos (Postigo, 2016). The simplicity of this assessment system increases the likelihood of receiving feedback from more users. However, it may be difficult to interpret why people like or dislike a video. Some may also press “like” on everything they read. Although one study finds that some of the features in YouTube (e.g., recommendations, searching) allow learners to find

relevant educational videos in an effective way, it is important to be aware that many views does not automatically imply high quality or credibility (Lee et al., 2017).

In online platforms like Reddit, the quantitative results of simple voting, like upvoting or downvoting, will have a direct impact on what information gets the front page of attention. These voting systems allow transparent voting, by letting new individuals see how others have voted before they vote themselves. This transparent quantitative rating will be part of the assessment when deciding whether one will use time reading an article or viewing a video. If a large number of votes are given, such rating systems aim to reduce bias and provide a fair and precise assessment of the quality of a knowledge product.

Another important way of understanding popularity on YouTube is through the number of subscribers each channel attracts. The “most viewed videos” and “most subscribed channels” represent different types of engagement, the most viewed has the greatest outreach, but the most subscribed has most engagement. The increasing number of channels with more than a million subscribers is indicative of the growth of YouTube. In 2010, there were only five such channels, while in 2016 there were 2000. The most subscribed channels are dominated by “YouTube stars.” These stars aim to be authentic and create a community by interacting with the fans. Today, the subscriber count is considered to be a measure of audience engagement and return visits. It is considered the key metric that can generate revenue from sponsorship and merchandising (Burgess & Green, 2018: 87–92).

The subscriber will have an overview of their favorite channels and automatically be notified whenever a new video is added to these channels. This feature promotes more viewing and is very important in getting consistent views. If the viewers are not happy with the channel, they can stop subscribing. For instance, video game commentators, will use a lot of effort to both get subscribers and not lose them. They both compete against each other in getting subscribers, but at the same time they share them with each other. Top commentators will often favorite each other’s videos, promoting their channels to each other’s subscriber bases. The most popular commentators also have each other on their respective channels as guest commentators, or provide links to other commentators on their own channel. This reciprocity strengthens the relationship between gameplay commentators and help both channels grow (Postigo, 2016).

Typically, most sites will provide a collection of many comparable solutions that are ranked against each other. The aggregated quantitative rating will often be used as a part of ranking system that provide a lot of attention to a few persons or solutions. Today, these network gatekeepers are often labeled as “influencers,” and they have a huge impact on the information flow in the system and what gets attention. They choose what information is valuable by connecting networks or clusters of individuals and information to one another (Halavais, 2018: 66–68, 87–90, 106).

Furthermore, most rating systems include qualitative data, in the form of comments or reviews of knowledge products. Although fewer persons will write comments compared with giving likes, they still play an important role in providing additional meta-information about the quality of a solution. When the comments are shared openly like in YouTube, viewers can read all the comments that appear at the bottom of the web page. The number of comments will increase over time, and make the aggregation of comments increasingly valuable and relevant. Still social participation will usually be minimal with a majority of comments being very short (Klobas et al., 2019). At some sites, the comments invite to more lengthy reviews of books (e.g., Amazon), movies (e.g., Internet Movie Database) or home-stays (e.g., Airbnb). In general, comments can provide more detailed information than a quantitative scale, including both strengths and weaknesses. In addition, reviews of open textbook are important (see Section 3.2). Since the online version of the book is free, there is more skepticism regarding the quality. Open reviews make the quality of the textbook transparent and add supplementary information, in contrast to traditional textbooks (Pitt et al., 2019). Many amateur online communities are also built around informal peer review, like for instance, fan fiction communities (Black, 2008). In other contexts, such as science, peer review is still not open and done anonymously. This intends to ensure honest feedback. Open public comments could be valuable, but scientists appear to be reluctant to publicly criticize someone else’s work (Nielsen, 2011: 179–180).

In relation to stigmergy, the rating of complete solutions can be seen upon as a human marker-based stigmergy, which today has become essential in organizing the abundance of human knowledge in the online setting.

### 6.3 Reestimating Solutions

In the US presidential election in 2016, all the polls got it wrong when Trump surprisingly won the election. This spurred an interest in



Table 6.1. *A comparison between the election results in the swing states in the US 2020 presidential election and the results from a prediction market*

Swing states	Actual result of the US presidential election	Predictions made at PredictIt the day before the election (November 2, 2020)
Georgia	Biden winner. Margin 0.2% (49.5% vs. 49.3%)	Trump winner (61¢)
Arizona	Biden winner. Margin 0.3% (49.4% vs. 49.1%)	Biden winner (52¢)
Wisconsin	Biden winner. Margin 0.7% (49.6% vs. 48.9%)	Biden winner (70¢)
Pennsylvania	Biden winner. Margin 1.2% (50% vs. 48.8%)	Biden winner (61¢)
Nevada	Biden winner. Margin 2.4% (50.1% vs. 47.7%)	Biden winner (74¢)
Michigan	Biden winner. Margin 2.8% (50.6% vs. 47.8%)	Biden winner (69¢)
North Carolina	Trump winner. Margin 1.4% (50.1% vs. 48.7%)	Trump winner (56¢)
Florida	Trump winner. Margin 3.3% (51.2% vs. 47.9%)	Trump winner (60¢)
Texas	Trump winner. Margin 5.6% (52.1% vs. 46.5%)	Trump winner (74¢)
Ohio	Trump winner. Margin 8.0% (53.3% vs. 45.3%)	Trump winner (72¢)
Iowa	Trump winner. Margin 8.2% (53.2% vs. 45%)	Trump winner (76¢)

alternative ways of predicting how the US 2020 election will end. Many looked to prediction markets, which allow individuals to buy and sell “shares” on whether a future event will happen or not. In the recent 2020 US election, PredictIt, a well-known prediction market, set Biden at 61¢ and Trump at 44¢ (“Who will win the 2020 U.S. presidential election?,” 2020).

The market suggested that Biden was most likely to win, which he eventually did. Since Biden was the favorite among the bookmakers, this may not seem so impressive. However, what is remarkable is that the prediction market also picked the correct winner in all swing states, with the exception of Georgia (Table 6.1).

Biden won a number of swing states, where PredictIt had also priced him as the most likely winner, like Pennsylvania (61¢), Wisconsin (70¢), Michigan (69¢), Nevada (74¢), Arizona (52¢). It is interesting how the

market estimated a close race in Arizona with a Biden win at 52¢, and it did also actually end with a very narrow winning margin to Biden at 0.3 percent. The market also correctly predicted a most likely Trump win in swing states North Carolina (56¢), Florida (60¢) Texas (74¢), Ohio (72¢), and Iowa (76¢). However, the prediction market is not perfect. In Georgia, Trump was predicted as a winner at 61¢ with Biden at 40¢. Here, the extremely tight race ended with Biden winning with a margin of only 0.2 percent. In sum, the prediction market showed an impressive accuracy that few experts would have been able to match.

Although political betting is illegal in US, PredictIt is an exception because it is primarily used by academic institutions for research purposes. The website has between 110,000 to 120,000 “funded accounts,” and around 1,000 of them are highly active users who trade as a part-time to full-time job in terms of their trading volume. The site is highly regulated, individuals cannot use more than \$850 on an individual bet, and only 5,000 traders can be in one specific “market” (Mashayekhi, 2020).

In this section, I will look closer at the stigmergic mechanisms that make this performance possible. Prediction markets are basically a crowd decision system, which use a market mechanism, often real money, to aggregate information from a large numbers of individuals. Individuals are invited to buy and sell contracts of predictions on the outcome of a future event. These virtual shares will typically pay one dollar if an event happens or a candidate wins the election and nothing if otherwise. The initial price of the contract is 50 cents. The aggregated effect of all the individuals who buy and sell the contracts will be the equilibrium market price of the contract. This price will change during the period up towards the time of the event, and the latest price is presumed to represent the current best guess about the probability of the event occurring. A contract trading at 60 cents implies the crowd believes there is a 60 percent chance that the event will occur at that given point of time. If circumstances change, and individuals change their mind and think that a political candidate will not win after all, they can sell (or “short”) the contract, and take the profit or loss before the event has happened. When the trading in the market changes, the market price will be adjusted accordingly (Buckley & O’Brien, 2017). Evidence suggests that prediction markets can outperform sales projections, journalists’ forecasts, and expert economic forecasters, and often it even performs at least as well as opinion polls (Atanasov et al., 2017).

As a type of stigmergic problem solving, it is the fluctuating market price that provides the crowd wisdom. It represents the collective opinion

of the group, which is constantly being *reestimated*. During the entire betting period, the price can be regarded as an estimated solution that gives the best information about the likelihood of an event happening. At a collective level, the forecast is built on letting everyone participate on equal terms (Buckley & O'Brien, 2017). Traders are motivated by profit, so if they obtain new and relevant information, they act quickly in the market. The probability of the solution will therefore be continuously reestimated. Both the price availability and the price history promotes transparency during the process, illustrating that knowledge and opinions can be shared openly in this system (Atanasov et al., 2017). In prediction markets, individuals exchange information by placing orders, while in polls, the design let individuals make solo predictions. Since the value placed on the assets is set in an open market of buyers and sellers, participants are informed and socially influenced by each other through various market indicators (e.g., movements in prices, trading volume, volatility). The process is therefore different from the focus on independent contributions in the classical wisdom of crowd approach (Surowiecki, 2005). The collective problem-solving process is much more dynamic because of continuous aggregations made by multiple parties (Tindale & Winget, 2019). Economists have highlighted that the latest price reflects all information available to market participants. Prediction markets are designed to produce continuously updated forecasts about uncertain events (Atanasov et al., 2017). They can also cover multiple options of mutually exclusive alternatives, like all the candidates in a presidential election race (Buckley & O'Brien, 2017).

The first modern prediction market was the Iowa Electronic Market (IEM), which opened in 1988. Historically, prediction markets have been used to forecast elections, but many organizations are now also interested in using such methods to gather information from employees, both in risk management, product sales, project completion, or idea generation (Buckley & O'Brien, 2017). Today there are numerous markets, with PredictIt being among the most popular. Here, people can also discuss predictions and provide arguments about the current prices on specific events.

A number of studies find that prediction markets can provide more accurate forecasting than other aggregation methods like polls. It appears that several crowd wisdom mechanisms are present. *First*, participants are encouraged to search for relevant information and actively use that information to their benefit. It is therefore likely that more individuals will make informed contributions. The market incorporates differences in

forecaster knowledge and skill. *Second*, the market price provide an automatic and continuous aggregate of the collective opinion. This is more effective than repeated polling and shows how prediction markets can provide updated information over extended periods. *Third*, prediction markets can easily scale to very large groups, and it will still work well when a majority have little relevant information because the few well-informed participants will be more motivated to increase the trading. If some participants are very confident, they will just buy more contracts. *Fourth*, participants can be anonymous which is important since social interaction can have a negative effect on the betting. *Finally*, bias is less of problem since unbiased participants can profit by exploiting others' biases. If somebody wants to manipulate the price, the prices will be corrected within a relatively short amount of time because participants wanted to profit from contract mispricing (Buckley & O'Brien, 2017).

However, one should be aware that prediction markets also fail. In the 2016 US presidential election, Trump was priced at 22 cent at PredictIt, but he still won the election. One explanation can be systematic bias among the participants using prediction markets. They are typically well educated and among the upper income groups. Another possible explanation is that traders sometimes move into a bubble, they became convinced of the inevitability of Clinton winning (Graefe, 2017).

Another type of dynamic forecasting is prediction polls. Here, individuals place predictions on future events, but they make probabilistic forecasts, either independently or as a team. The participants can change opinions during the betting period, and as in a prediction market, will receive feedback after the event is over.

In one study, the prediction poll addressed geopolitical issues like "Will any country officially announce its intention to withdraw from the Eurozone before April 1, 2013?" More than 2,400 participants made forecasts on 261 events over two seasons. Here, forecasters in teams outperformed the individual forecasters. The online teams comprised 15 persons and these teams both shared and discussed different issues, but they did not need have to reach consensus. Instead, they made individual predictions, and the team score was based on the median forecaster. The teams would have information about the forecasts of their teammates. Every team also had an overview of accuracy scores for each member and a separate leaderboard, which compared performance across teams. Both the use of individual leaderboards (within teams and across independent forecasters) and team leaderboards (across teams) created a significant level of gamification in the design. The competitive features between the

forecasters are assumed to strengthen motivation, and there are many learning opportunities in the process. However, teams were not allowed to talk with persons in other teams. The statistically aggregated forecasts also gave more weight to the most recent contributions and the best performers (Atanasov et al., 2017).

In this study, the team prediction poll performed better than prediction markets, especially in the first phase of a prediction period. It is likely that both motivation and learning processes have played a role. The design offers a complex mix of intrateam cooperation and interteam competition. Teams would share information, learn of each other, and motivate each other to update their forecasts more regularly. This strategy is especially effective when the number of active forecasters is small (Atanasov et al., 2017).

Dynamic forecasting can also be used to filter irrelevant solutions. In online innovation contests, the best solution must often be identified among a large number of proposals. There may be hundreds of irrelevant ideas that are time consuming to review (Klein, 2017). One example is Google's charitable 10 to the 100th project. More than 150,000 suggestions were submitted and Google had to deploy 3,000 employees to filter all the ideas, putting the process nine months behind schedule (Lykourantzou, Ahmed, Papastathis, Sadien, & Papangelis, 2018).

The most typical crowd-filtering strategy is majority voting, which is used by Threadless, for instance. Anyone can vote on the best T-shirt designs and the most popular ones win prize money. Another alternative is multiple voting, which gives each individual in the crowd a certain number of votes to allocate to the ideas they prefer (Garcia & Klein, 2017). However, recent studies have shown that the crowd can be effective in filtering proposals by removing irrelevant ideas. In one study, the crowd was assigned to assist a review panel in a contest by removing irrelevant ideas and keeping the best solutions. In the Diverse Bag-of-Lemons (DBLemons) strategy, each participant is given ten lemons (corresponding to 19 percent of the idea corpus) and asked to identify the worst ideas instead of the best ones. By using a dynamic ranking system, participants can vote at different points of time. It is also possible to view how previous users have ranked the ideas. In addition, algorithms "force" reviewers to compare more diverse ideas with each other. The results show that each participant only had to look at approximately 50 percent of all the ideas to include all the best proposals, illustrating the time efficiency of this strategy. (Lykourantzou et al., 2018).

DBLemons is partly inspired by studies of the portfolio effect, which claims that if you add more diversity to your investment, you take on less risk. People who are exposed to just a few ideas initially appear to get fixated on them, which can impede successful problem solving. The finding underlines the importance of designing diversity procedures across domains in crowd decision-making (Lykourantzou et al., 2018). DBLemons resembles prediction markets in that each voter can view how previous users have ranked the idea, but it is different in that one does not have to choose the winner, but only remove the losers. This approach represents an innovative way of involving the crowd because it makes the individual task much easier and still highly effective. When the number of ideas or the amount of shared knowledge is huge, effective removal of noise becomes increasingly important.

The accuracy of crowd forecasting shows signs of a new type of CI that builds on decentralization of expertise. The stigmergic mechanism in dynamic forecasting provide a transparent environment that let individuals learn from each other and build on this knowledge through different types of gamification. The result is that a diverse group of people is able to quickly synthesize all their knowledge, represented by a quantitative indicator in a scalar measurement, like a price or score. This gives everyone an accessible overview of the aggregated collective opinion of the best solution, which is being constantly reestimated.

#### 6.4 Completing Solutions

In this section, it is proposed that a third type of stigmergic problem solving is directed towards “completing solutions.” It is characterized by new work, which build on the current unfinished version of a solution. In Chapters 2 and 3, several examples illustrate how a large amount of people can develop such complex knowledge products through asynchronous contributions in an online setting, such as the Wikipedia project, argument mapping, and open databases.

Wikipedia is considered to one of the best examples of this type of human stigmergic problem solving in an online setting (Heylighen, 2011). Most of the collective work is done on separate articles where new modifications build on the current state of the specific article. All the articles can be regarded as a collection of separate attempts to complete a solution. All contributions leave “traces” in a shared medium that enable

new actions to build on previous ones. This stigmergic mechanism is present when edits in Wikipedia made by one individual trigger new edits made by another person. Other users can then again continue to work with the aggregated traces of others' work at any later point of time. This introducing a high degree of flexibility into the work.

Moreover, some online environments also let participants discuss how the work should be done. For example, in Wikipedia, every article is attached to a separate talk page that allows anyone to discuss the coordination of the work. Some of the most active authors might discuss specific issues with each other, while others may just briefly leave a comment or make suggestions on how to improve the article. The discussion is asynchronous and there is no guarantee that anyone will respond to a request, but all posts are archived and made easily accessible for anyone at a later point of time. However, a significant amount of the work in a Wikipedia article will be done without explicit coordination. For example, in one study of a Wikipedia article, about four fifths of the edits are done by editing the content in the article; while the remaining part were discussions on the talk page (Rezgui & Crowston, 2018).

Even mistakes can be valuable because they “trigger” others to make corrections, both minor spelling, but also misinterpretations of the content. Different types of errors, vagueness, or lack of information will stimulate different types of improvement (Heylighen, 2011; Heylighen, 2015 #483).

Corrections in Wikipedia can also be done without knowing how the complete article should end up looking. Each individual will only perform actions according to that person's interests, skills, or background knowledge. Normally, the more “confident” an individual is about the correct answer, the more it will be stimulated by the condition, and the quicker it will begin working (Heylighen, 2011, 2015). For instance, a spelling error triggers a proofreader in Wikipedia, while an imprecise article about Viking ships motivates an archaeologist to add new content. Because of the wide access in the global online setting, a diverse group of people with different expertise can easily make contributions based on their skills. Any contribution counts, everything from minor proofreading to major revisions. Over time, the quality will emerge through continuous attempts to improve the collective work.

Part of the coordination success in Wikipedia may be due to how discussions in specific areas are linked to separate articles. The structure is easy to grasp, minimizes the chances of redundant discussions, and makes it time efficient to join the discussion. By using a watch list feature

in Wikipedia, editors can also keep an overview of all recent changes (Rezgui & Crowston, 2018). This modularized structure enables millions of article projects to move forward in parallel where all simultaneously aim to complete a solution.

Furthermore, completing solutions can be about “filling in the missing part” of a solution. This can be a bird observation in the eBird database or adding a new argument to an argument map (see Chapter 3). Here, many persons contribute with different pieces to complete map that aims to provide a fair and accurate overview of a solution. When a new contributor needs to position a new argument into an existing map structure, it is necessary to read some of the arguments already published. These argument maps can be used to support and organize complex political discussions, but usually a moderator will need to help organize the map and approve comments.

According to Bullen and Price (2015), the abundance of online information challenges us to develop a new form of literacy that make us better able to grasp the interconnectedness between many different problems and potential solutions. Collective argument mapping may be a help in describing complex problems. In one example, College of Contemporary Health’s (CCH) Obesity used DebateGraph, an argument map in an attempt to create a more comprehensive and coherent visual representation of the obesity policy space. The obesity problem is rapidly increasing, and by 2030, it is expected that approximately 40 percent of the world’s population will be overweight or obese. Obesity is a complex issue and needs to be analyzed from a range of different perspectives. At the same time, healthcare professionals and policy makers face an ever-expanding amount of data that needs to be synthesized and understood. The argument map includes causes, impacts, interventions, evidence, and barriers to change.

The goal of the Obesity DebateGraph is to help all stakeholders better understand the complexity of the problem in a larger societal context, and facilitate dialogue and critical thinking across the community. Here, the mapping system aims to integrate all kinds of information resources within the same map. It can also be used as a dynamic tool to update new information. The aim is to provide an overview of the most important resources and the current debates in the field (Bullen & Price, 2015). However, when the structure in the map becomes more complex, it is a challenge to sustain a complete overview of the collective work.

Another similar example is how the Climate CoLab organized innovation contest within the framework of a contest web that includes a family of related contests. This web aims to covers a broad space of possible



solutions, by organizing a large number of proposals into a complete taxonomy. The taxonomy covers different parts of the problem, and every category is mutually exclusive from each other. This strategy centers on dividing a complex problem into many small parts so individuals can more easily make different contributions. Because all the work is transparent and accessible, the community of participants can also be challenged to combine these partial solutions and create a more comprehensive solution (Malone, 2018: 182; Malone et al., 2017).

## 6.5 Adapting Complete Solutions

### 6.5.1 *Background*

It is suggested that a fourth type of stigmergic problem solving is directed towards “adapting complete solutions.” This type of problem solving aims to reuse and modify existing solutions. In Chapters 2 and 3, there are several such examples. One example is the integration contests hosted by the Climate CoLab, which require that new solutions must build on previous winner solutions from the innovation contests. All winners can participate in a new contest and the contestants are challenged to combine their own and others’ work in developing an integrated proposal (Malone, 2018: 173; Malone et al., 2017). A prominent example of this type of problem solving is open source software projects. According to Nielsen (2011: 57), the most basic characteristic with open source is that programmers don’t have to start from scratch, but can build on and incrementally improve what others have developed. The open distribution of code to anyone stimulates programmers to build a publicly shared information commons. Originally, the great programmers would write their programs largely from scratch within a very short period. In stark contrast, the best programmers today will instead know how to quickly reuse code from the commons, and assess what additional code they need to write from scratch. One could claim that every new solution indirectly builds on the work of thousands of other programmers. The advantage is that problems can be solved faster and more reliably compared with working from scratch. As the size of the information commons grow, the quality of the collective work will in general also improve (Nielsen, 2011: 57–59).

### 6.5.2 *Open Textbooks*

A more recent relevant example is the production of open textbooks, building on Open Access policies and the OER movement (see

Chapter 2). Today, there is an increased interest in the production of textbooks that anyone can access both because they can support education for all and possibly also maintain and improve the quality of these books (Al Abri & Dabbagh, 2018). Many open textbook projects build on a complete version that has already been published and which only needs minor modifications to be adjusted to a new context. Because open textbooks are usually published with a Creative Commons license, other textbook authors or educators can modify the original textbook so it better fits the local educational context. It is possible to both remix, adapt, combine, and add content. Because the open textbooks are digital, they can also easily be accessed by anyone. In a recent UK report, the “American” version of the open textbooks was not seen as barrier to their usage despite issues around language and other contextual issues. The license permits both minor and major changes. Among UK academics, there is considerable interest, not only because of cost savings, but even more because of the freedom to adapt and develop textbooks (Pitt & B., 2019: 1164).

Furthermore, new open textbooks movements are being established in other parts of the world. For example, in March 2016 the first Open Textbook Summit in Africa was hosted in Cape Town (Wiens, 2016). In the pilot project Open Textbooks for Africa (OT4A), the objective is to support the adaption of currently available open textbooks and the development of new textbooks that display African knowledge to the world. The cost of textbooks represents an even larger economic cost for students in the Global South (Wiens, 2016). In one example, a group of physics teachers at the University of Cape Town revised an open textbook in physics originally written by American authors and published through OpenStax (Wiens, 2016). In the topic, history of astronomy, the authors have replaced images of an archaeoastronomical site like Stonehenge in England with other similar sites in Egypt and Kenya (Merkley, 2016). In addition, this shift to an open textbook will save 150 first-year South African students 12,000 dollars at one institution over one academic year. Some of the textbook content has been changed to better fit with an African context (Wiens, 2016).

Today, there are more textbooks being published that build on adaptation of other textbooks. A core textbook is adapted to new language editions, and several different types of curriculum. Additional content about local, regional, or national preferences is included. The adaptation of material can also make more content accessible for people with disabilities. Students can even contribute. The book format is more flexible, and it is easy to reduce the size of the book and integrate it with other media

resources. These books are often published with a Creative Commons license that allows for easy modification of the original version without needing to ask for permission. In this way, new solutions can be adapted to different contexts as a part of long-term collective knowledge advancement.

### 6.5.3 *Internet Memes*

Internet memes comprise another example of adaptation of solutions, but with a purpose that is very different from textbooks. An internet meme can be defined as a group of digital items that share common characteristics of either content, form, and/or stance. In the popular culture, it usually describes the propagation of items such as jokes, rumors, videos, and websites from person to person via the Internet. A central attribute is that they spark user-created derivatives articulated as parodies, remixes, or mashups. One example is the video “Gangnam Style” performed by South Korean singer named PSY, which became the first YouTube clip to be viewed more than one billion times in 2012. In addition to watching the clip, thousands of people also created and posted their own versions of the video imitating the horse-riding dance from the original video, with videos such as “Mitt Romney Style” and “Arab Style” (Shifman, 2014).

The memetic content is simple, typically conveying one uncomplicated idea that imitates the original video in some way. Usually, some degree of repetitiveness complements this simplicity. This can be highly repetitive lyrics and melody (e.g., “Leave Britney Alone”). The repetitiveness may trigger active user involvement and make it easier to remake video memes. In other memes, repetition will be about imitating a well-known person, and others may again imitate these “imitations” (Shifman, 2014: 78–88). Humor is often important, but it can be “quirky and situational,” including bizarre translations and wacky teenagers. Some memes belong to specific subcultures that share their own language and symbols (e.g., LOLcats, rage comics). This culture flourishes on specific sites such as 4chan, Tumblr, and Reddit (Shifman, 2014: 118).

The memetic form is the concrete manifestation of the message in the meme, including the visual/audible dimension and the genre-related patterns (e.g., animation). The video will often be filmed in one single shot, making it time efficient to create. The visual effects are simple with little or no editing work. For example, the meme “Leave Britney Alone” only required a white piece of cloth, a camera, and a modicum, making it possible to make a new version with limited resources (Shifman, 2014: 78–88).

Previously, mass-mediated content would often be transmitted simultaneously from a single institutional source to many people. Memes are different because they spread gradually through many interpersonal contacts. The producers of memes are aware of each other because the content is circulated, imitated, and transformed by many different users. Memes should be understood not just as single entities that propagate well, but also as groups of content units with common characteristics and shared values in a digital culture (Shifman, 2014).

The memetic content will usually invite others to reuse and share the original work. Because most of the content is user-generated and reflect the opinions of a layperson, people will tend to react more to the memetic video compared with a professionally made video. It will often be perceived as more meaningful to respond to a peer than a celebrity. It is also more likely that other peers will comment on the new video responses, and thus reinforce a stronger sense of community (Shifman, 2014: 75–76).

In addition, because amateurs produce the videos, they will usually be textually incomplete or flawed, compared with a professionally made video. Paradoxically, a “bad” video production can make a “good” meme because inconsistencies often stimulate further spread and dialogue, active user involvement, and recreation of content. The unpolished, amateur-looking videos motivate people to address the puzzles: what is missing, or how bad it is. The memes become part of a socially constructed public discourse that include diverse voices and perspectives (Shifman, 2014).

Another attribute with memes is that they easily become popular because they are interlinked with each other. A new version of a specific video will draw attention back to the initial memetic video in a reciprocal process. This increases the likelihood of appearing in YouTube’s suggestions bar as a highly relevant search result when viewers search for the initial meme. This is particularly important with user-generated content and amateur videos that do not necessarily receive many views. The metadata with viewer statistics and comments will constantly be aggregated and displayed to all users. This information is increasingly becoming an influential part of the process itself – with people considering it before they decide whether to remake a video, (Shifman, 2014: 32–33).

In the political domain, internet memes open new types of democratic discussions in the online setting. It is an accessible, cheap, and “enjoyable” way of voicing one’s political opinions. It allows for the creation of multiple and diverse opinions. Major political events often spur a large

number of commentary memes, which are used for political advocacy in different ways. Social media allows for new types of political participation, especially among younger citizens. This was first demonstrated in the 2008 US presidential election campaign, when massive amounts of politically oriented user-generated audiovisual content was created. Clips such as “Obama Girl,” “Wassup,” and “Yes, we can” attracted millions of viewers. The political campaigners produced only a fifth of the most viral videos, while interest groups and other nontraditional actors produced the rest. Popular videos were also transformed into memes, addressing and advocating issues with both humor and seriousness. (Shifman, 2014: 51–52, 120–126).

Another example is the “Pepper-Spraying Cop,” a meme originating from November 2011, when students gathered as part of the Occupy Wall Street protest. When they refused to move, two police officers reacted by pepper-spraying a row of still-sitting students directly in their faces. Shortly after, videos documenting the incident were uploaded to YouTube, generating negative reactions in the public opinion. A photograph in which one of the officers was spraying the students quickly evolved into an internet meme. In the aftermath of the protest, the image was photo edited into a large range of contexts, spanning historical, artistic, and pop-cultural-oriented backgrounds (Shifman, 2014: 51–52).

Two main groups of memes were produced. The first group of user-generated images focused on political contexts. For instance, the police officer is shown pepper-spraying iconic American symbols such as George Washington, the Constitution itself, and other freedom fighters across the globe. All these new meme versions tell the same story, that the police brutally violated the basic values of justice and freedom (Shifman, 2014: 51–52).

The second group of memes is pop-culture oriented and shows the police officer pepper-spraying icons such as Snoopy and Marilyn Monroe. Some of these versions show an entirely different use. In one case, the pepper-spraying cop is used to criticize Rebecca Black – a widely scorned teen singer and internet phenomenon. The memetic content can also lead to stance alternations. While the politically oriented versions are mainly sardonic, the tone in the pop-culture-oriented ones is more playful and humorous. The “Pepper-Spraying Cop” meme shows a diverse type of diffusion and evolution (Shifman, 2014: 51–52).

Furthermore, memes can be part of grassroots action that links personal stories and political issues in an attempt to empower and mobilize citizens. A protest like the American Occupy Wall Street was not backed by a

strong formal organization, but rather by digital networks. The slogan in the Occupy Wall Street's "We are the 99 percent" meme refers to the argument that 1 percent of the American population controls the country's financial wealth. The memes showed a person holding a handwritten text depicting a gloomy personal story. People would show their agony with a serious facial expression and by holding "I am the 99 percent." The stories could be about not affording medication or struggling to provide for children. The combination of repetition and variation turns these personal stories into a larger political issue. The misery is not just a personal problem because the collective network of memes show that the system has failed (Shifman, 2014).

On the one hand, a meme is unique, and on the other hand, the new version will signal membership in a large community that use similar messages. The multitude of new versions also helped promote the topic on the mass media agenda, drawing more attention to the movement. The popularity of the "99 percent" meme even generated a countermeme: the 53 percent meme. Conservative activists introduced a rhetoric with an opposite stance, bringing in conflicting information, underlining that only 53 percent of American people pay income tax. It illustrates how a stance in the meme can either imitate a certain position or introduce an alternative perspective based on the same idea (Shifman, 2014).

Memes play an important role because shared slogans communicate easily across large and diverse populations. Personalized content is also shared in large-scale, fluid social networks across the globe by ordinary internet users. The power in the memes lies in the message not just being standardized content distributed to everyone, but instead being personalized and adapted so individuals can tell their own stories. The community is simultaneously both local and global (Shifman, 2014: 127–129).

## 6.6 What Is Human Stigmergic Problem Solving?

### 6.6.1 *Solution-Centered Collective Problem Solving*

In stark contrast to swarm problem solving, human stigmergic problem solving is centered on the improvement of solutions, and not predefined problems. These solutions can be improved by rating, reestimating, completing, or adapting them. These solutions change continuously, whether it is a new individual rating or a new edit of an unfinished draft. Problems are relevant to these solutions, but they do not constitute the premise for this type of collective problem solving.

In stigmergic problem solving, it is assumed that solutions exist independently of the problems. As pointed out by Von Hippel and Von Krogh (2015), one should be aware that there already exists a wide range of solutions in the environment. They use a story to illustrate their point: An employee visits a trade show “just to see what is new.” There, the employee discovers a new payroll-processing software, and thinks it might be relevant to use in his firm: “I wasn’t thinking that we had any payroll-processing deficiencies, but now I recognize that we do, and that this technology might make this work more effective.” It is the new solution that creates a need. The formulation of a problem, that the payroll system in his firm is not as effective as it should be, is formulated after the solution has been discovered. In the solution–need pair there is no initial independent problem identification. One could describe the inferior previous arrangement as “a problem,” but this is only possible to do post hoc after the discovery of a new solution (Von Hippel & Von Krogh, 2015). With the online setting, these “trade shows” are everywhere because so much information is stored on the Internet.

Many more solutions are available than ever before. They are stored as a collective memory in the online setting, which makes it possible to easily find previous solutions, reuse and modify them at any point of time. A solution can be used to solve many different problems for different persons at different points of time. The solutions range from complete independent solutions, estimated solutions, to very incomplete solutions. The environment functions like an external memory that registers and stores a proposed solution or part of a solution. When stored in a digital format, solutions can be reused in many intended and unintended ways. As the examples show, it opens up for the production of many different versions of an existing knowledge product. This happens as a distinctly asynchronous problem-solving process, contrasting human swarm problem solving that stresses a rapid and synchronous response to a specific problem. If the traces of the work (solutions) are shared openly, they can be reused in many different ways in the future.

The problem of free riding is almost removed because it requires little additional effort or cost to leave traces of your work in the online setting. “Free riders” are here defined as individuals who benefit from others’ efforts without doing anything in return. An answer to a problem stored openly in an online setting will not only solve a specific time-restricted issue in a local context, but the solution will automatically become part of a huge collective “map” of interconnected knowledge. Open sharing often demands very little extra work. For example, if someone posts an answer to

Table 6.2. *An overview of the four different types of stigmergic problem solving*

	Sematectonic vs. marker- based stigmergy	Quantitative vs. qualitative stigmergy	Examples	Similarities with other animals
<i>Rating complete solutions</i>	Marker-based	Quantitative (actions performed on the same action)	– Reddit (collaborative filtering) – YouTube	Ants: Gradient following in a single pheromone field
<i>Reestimating solutions</i>	Sematectonic (marker- based)	Quantitative (actions performed on the same action)	PredictIt	Ant trails
<i>Completing solutions</i>	Sematectonic (marker- based)	Qualitative (actions performed on separate parts of the medium)	– Wikipedia – Argument maps	Wasp nest construction
<i>Adapting complete solutions</i>	Sematectonic (marker- based)	Qualitative (actions performed on separate parts of the medium)	– Open textbooks (revising) – Memes (remixing)	Termite nest building (decisions based on combinations of pheromones and the construction work as it is)

a problem in a discussion forum or makes a video response that demonstrates how to solve a practical issue, many other people can at a later point access this “frozen” solution and reuse it. On the Internet, mutual interaction mechanisms like tit-for-tat (e.g., “prisoner’s dilemma”) do not function in the same way as in the offline setting. In this way, the web itself can be regarded as a shared memory with solutions that are left as permanent traces in the online setting (Heylighen, 2011, 2015).

Table 6.2 gives an overview of the four different types of stigmergic problem solving that have been discussed in this chapter and what type of stigmergic mechanisms they build upon.



The table shows that “completing solutions” and “adapting complete solutions” (remixing and revising) primarily build on qualitative stigmergy, while “rating solutions” and “reestimating solutions” build on quantitative stigmergy. Human stigmergic problem solving will also often combine aspects of sematectonic stigmergy with supplementary marker-based stigmergy that produce relevant metainformation.

#### 6.6.2 “Rating Complete Solutions” as Marker-Based, Quantitative Stigmergy

Because an enormous amount of knowledge products is being stored in an online setting, it is essential to develop mechanisms that can rate the quality of all these contributions. On one hand, the rating of these complete solutions can be regarded as marker-based stigmergy. Users will actively “mark” knowledge products, by providing different types of metainformation, like qualitative comments and quantitative ratings. Videos published on YouTube automatically map user behavior by generating reading statistics, but users can also actively choose to like, subscribe, or comment on a video.

On the other hand, these assessments and reviews will change over time. The aggregated metainformation about the solution will often be displayed as a rating result, whether this is number of likes, subscribers, or the frequency of knowledge sharing (e.g., retweets). In science, number of citations is one such rating mechanism. All of these build on quantitative stigmergy.

At a macro level, this metainformation compares and ranks a collection of solutions according to their quality and relevance. When there are a huge number of alternatives, a rating system provides a collective assessment and filtering of the available solutions. Algorithms in search engines also give much weight to these user ratings when they index and rank knowledge products. Like with the ant trail-laying systems, the rating systems attempt to provide the shortest route to the best solution through the enormous amount of digital information that constitute our human collective memory.

However, because so many solutions are produced online, only a few will be visible in the search engines. One challenge with top-rated solutions is that the most popular ones are not necessarily the best one. The structure of the web follows a winner-take-all distribution that amplifies attention towards a few solutions, while many remain unnoticed. Because of positive feedback, the initially most promising solutions will grow very

quickly, while the rest will be lost. Without already being on the ranking list, it may be difficult to initiate interest in a solution because few persons will know about it. Search algorithms tend to increase the current imbalance and reinforce existing networks of popularity. Their lack of transparency also makes it more difficult to understand why some solutions get more attention (Halavais, 2018: 102–115).

### 6.6.3 *“Reestimating Solutions” as Quantitative, Sematectonic Stigmergy*

“Reestimating solutions” is a type of stigmergic problem solving, which centers on different types of dynamic forecasting. A market mechanism typically gives continuously updated information about the probability of a collective outcome. While “rating complete solutions” builds on marker-based quantitative stigmergy, “reestimating solutions” utilize quantitative, sematectonic stigmergy. The market price or the “voting leaderboard” is considered the most accurate indicator of the solution. Individuals can access the aggregated prediction of the entire group, which is continuously updated (e.g., PredictIt, Kickstarter). Based on their individual background knowledge, they can choose whether to engage or not. At an individual level, single individuals make estimates or bets, which vary in performance, but at the collective level, the aggregated fluctuating market price is considered to cancel errors and provide the best predicted estimate of a solution. Both “rating complete solutions” and “reestimating solutions” are similar in giving individuals access to transparent updated information about the crowd opinion at any time.

### 6.6.4 *“Completing Solutions” as Qualitative, Sematectonic Stigmergy*

“Completing solutions” is another type of stigmergic problem-solving which manifests itself in the urge to fill in the “missing part,” like a piece of the puzzle in an argumentation map, or missing data from a geographical area in an eBird database. The perceived incompleteness of a solution triggers the motivation to make new contributions. However, the individual does not have to know what the final solution should look like. This mechanism builds on qualitative, sematectonic stigmergy.

Stigmergic actions stimulate their own continued execution via the intermediary of the marks they make. The completion of one task triggers a new task. The motivation will be to fix something that is missing, by “filling something out,” “fixing an error,” or “creating new order” (Heylighen, 2015). The “shared material” is itself regarded as a type of

communication that allows for coordination independent of prior planning, norms, or explicit discussions (Rezgui & Crowston, 2018).

It echoes the construction work of bees, who are also automatically triggered by a special “configuration of incompleteness” in the solution. In their construction work, a type-1 stimulus triggers action A by individual 1. Action A will then transform the type-1 stimulus into a type-2 stimulus that triggers action B by individual 2. This mechanism allows for effective indirect cooperation between individuals (Theraulaz & Bonabeau, 1999). Like the bees, a contributor on Wikipedia will first begin to contribute when recognizing the right start conditions that correspond to that individual’s background knowledge. The solution is mediated through a draft version that is changing over time. Further work is stimulated by the current state of the incomplete document with its limitations and errors that stimulate further modification. Each author is stimulated by what previous authors have written and use this information to either add, revise, or remove content (Parunak, 2005).

The contributors build on others’ work by removing, adding, or correcting existing content. Anybody can change almost anything, and there is no editor who divides the tasks. On a macro level, this type of stigmergic problem solving can also involve large complex, self-organizing system with contributions from individuals distributed all over the globe. One example is the eBird database or the global network of Wikipedia articles. In these systems, thousands of people make independent contributions according to their interests and competence without any centralized control. Large projects will often be modularized, making it easier to participate in a smaller separate part. Based on their expertise, individuals make relevant contributions of different size, only coordinated indirectly through an online environment.

This type of collective work is also reliant on appropriate digital technologies like a wiki or an argument map technology. Several systems combine sematectonic stigmergy with marker-based stigmergy that permit asynchronous reflective communication about the ongoing collective work. In Wikipedia, this includes discussions of content, wording, and structure on the talk page of the different articles. The transparency of this metadiscourse provides future contributors with an informal review of the quality of the article and information on how it can be improved. These discussions may emerge over a very long time, involving many persons who are unknown to each other and not even aware of the new contributions being made.

In addition, the active contributors in the community use much time discussing and developing policies and procedures on separate wiki pages. This explicit coordination constitutes a growing proportion of all the work being done in Wikipedia (Kittur & Kraut, 2008). On one hand, sematectonic stigmergy will be present in the editing of the Wikipedia article. On the other hand, marker-based stigmergy will be present on the talk pages that are separate but attached to each article.

#### 6.6.5 “Adapting Complete Solutions” as Sematectonic, Qualitative Stigmergy

In adapting a complete solution, a part of the original solution is reused and some type of revision is done, which can both involve minor or major changes. Because information is digitized, it has become much easier to make multiple new versions of a knowledge product, whether this is a textbook or a meme. This adaptation builds on sematectonic, qualitative stigmergy, since the new solution is reliant on direct modification of the original content. This repackaging can be done through either revising or remixing a solution.

In revising the knowledge product, the entire content is modified. The original open textbook will be transformed into a new complete textbook, like when an open textbook is translated into a new language to make it appropriate for a local context. This customization process can involve both major and minor changes, parts of the original content can be removed, modified, and new content can be added. Adaptations of open textbooks will usually require expert contributions because the goal is to end up with a “polished” knowledge product that can be used in an educational institution. Because of the Creative Commons license, a new author can modify the text without notifying the original author. This license exemplifies marker-based stigmergy that permits flexible reuse of the original work.

When remixing a work, a new interpretation is created from an individual part of it. In memes, some part is retained, while other parts are substituted with a new version of the same content through editing an image or adding a new soundtrack. In an offline setting, memes will usually change both their form and content because it is almost impossible to retell something in exactly the same way. However, in the online setting, one can easily retain some part in its complete original form. For example, in the “Occupy Wall Street” memes, everyone made a meme with a written text saying, “I am the 99 percent” on the video clip. New versions followed the same “production or remix rules.” The content

will often aim to be unpolished, authentic, and emotionally laden, which increases the likelihood of getting response from others. This is an example of sematectonic qualitative stigmergy.

Furthermore, memes are interesting because the interlinking and sharing of content is important. The adaptation time is quick because many individuals post a video based on their first shooting. Because a large number of new versions are produced, one could claim that memes also utilize a type of sematectonic quantitative stigmergy. Since all memes are linked together in a network structure, they tend to get a lot of attention in search engines. If the memes are made as personal stories about a political issue, they can together represent a powerful political statement (Shifman, 2014).

#### *6.6.6 Improvement of Solutions as the Basis for Human Stigmergic Problem Solving*

This chapter has intended to show how human stigmergic problem solving process is solution-centered. It builds on a version of a solution that already exists, either partially or complete, and aims to improve it by rating, reestimating, adapting, or completing it. These solutions may be relevant in many different ways, like a video or a Wikipedia article, which can be used to help solve problems for many different persons at different points of time. Consequently, a specific problem is not defined in advance, but can instead be regarded as an offspring of a solution. The solution is offered to anyone as part of a shared collective memory that emerges over a longer period.