Stefano De Paoli:

The Automated Production of Reputation: Musing on bots and the future of reputation in the cyberworld

Abstract:

Reputation is considered as the summary of a person's relevant past actions in the context of a specific community and is a concept which has gained huge relevance in the cyberworld as a way of building trust. Increasingly, however, reputation is awarded to users after they have carried out repetitive, mechanical or trivial actions. This opens the space to a phenomenon which we can define as the automated production of reputation: reputation produced by the means of software technologies known as bots that can easily automate repetitive online actions. In this paper the phenomenon of automated production of reputation is preliminarily defined and presented using three different empirical examples: Massively Multiplayer Online Games, the social network twitter and the reputational hub Klout. The paper also discusses some of the foreseeable negative consequences of the automated production of reputation and in particular the risks related to the loss of trust in online communities.

Agenda:

Author:

Dr. Stefano De Paoli:

- Fondazione <ahref, Vicolo Dallapiccola 12, 38122, Trento, Italy
- ☎ + 39 - 0461 - 235794 , ✉ stefano@ahref.eu stefano.depaoli@gmail.com
- Relevant publications:
  - De Paoli Stefano and Kerr Aphra. On Crimes and Punishments in Virtual Worlds: bots, the failure of punishment and players as moral entrepreneurs. Ethics & Information Technology, 14(2): 73-87, 2012.

ISSN 1614-1687

Introduction

There is a phenomenon which will have an increasing relevance for the future of reputation in the cyberworld: *the automated production of reputation*. In this position paper I will define it, reflect on it using empirical examples and make some preliminary observations on some problematic issues related to it.

Reputation is considered as the summary of a person's relevant past actions in the context of a specific community and is a concept that has gained huge relevance in the cyberworld as a way of building trust. Some authors even see online reputation as the central aspect of contemporary digital society and talk about "The Reputation Society" (Masum and Tovey, 2012). In the cyberworld, reputation is "created" and "disseminated" by the means of technological systems known as reputation systems (Dellarocas, 2012). Increasingly, however, in these systems, reputation is awarded to users after they have carried out repetitive, mechanical or trivial actions. An immediate example would be awarding simple "likes" to a fan page on Facebook[15] (an action that simply requires the repetitive clicking of a button on the interface).

Generally speaking, repetitive and mechanical actions are often automated with technologies. An example we can think about is repetitive actions carried out by workers in manufacturing — a process that can be automated using assembly lines. The same is often true for digital repetitive actions as well (e.g. a "like" on Facebook). In this case the automation is often achieved by means of software known as *bots or socialbots* (when they are used on social network sites). Bots are software agents that can replace users in carrying out repetitive tasks and can easily automate several online actions. In many contexts bots are legitimate technologies as they support the user in conducting repetitive actions. For instance Wikipedia bots support the Wikipedia community in carrying out repetitive tasks to maintain the English language Wikipedia. Many bots however can be used deceptively and for illegal purposes. For instance in online games, bots can be used to cheat, causing direct damage to fair players (De Paoli & Kerr, 2010) and game businesses. Bots can also be used to "produce" reputation values on behalf of the user — by an automation of repetitive actions awarding reputation: this is, in the first place, what I call the automated production of reputation. The automated production of reputation is a form of cheating and also a deceptive use of bots that could have serious negative consequences, first of all undermining the role of trust as a social regulatory feature of interplay in the cyberworld. The goal of this paper is to explore these problems. I will do so by introducing some empirical examples, following an approach that uses empirical material to introduce what is essentially the beginning of a conceptual exploration.

In the remainder of the paper I will firstly introduce the concept of reputation and its relevance for the cyberworld. Secondly, I will substantiate the concept of the automated production of reputation and augment it with three short empirical examples (Massively Multiplayer Online Games, the social network, twitter, and the reputational hub, Klout) that justify my claims about the increasing relevance of this phenomenon. I will finally discuss the main risks that the automated production of reputation can have for the cyberworld and finally trace a perspective for further research into this subject.

Reputation in Context

At an individual level, reputation is the summary of a person's relevant past actions in the context of a specific community. It is a collective value of trust that a community awards to a person. In other words people prefer to interact with reputable persons, whose trustworthiness has been assessed by the social group to which they belong (Dasgupta, 1988). As a form of trust, reputation allows actors to reduce the complexity of action and take decisions in situations of risk when otherwise they would possess insufficient knowledge (Luhmann, 1979). As a form of trust, reputation can be seen as a functional alternative to rational prediction for the reduction of the complexity of social action. Hence, to a certain extent, reputation can be considered a form of what Taddeo (2009) calls referential trust: "the kind of trust that one develops in an unknown agent by considering only the

[15] Further examples of mechanical actions awarding reputation are described in the section, Automated Production of Reputation.

recommendations about that agent provided by other agents or by other information sources, such as newspapers or television". Referential trust therefore enacts an array of expectations that people have of each other based on cross-references related to past actions. Taddeo goes on by saying that "Referential trust is one of the main kinds of trust developed in digital environments in which communication processes are easily performed". Because of this particular feature, reputation (i.e. a form of referential trust) is a concept that has been largely adopted as a way to build trust in the cyberworld (Jøsang et al. 2007).

According to Capurro (2006), in its broader sense information ethics deals with questions of digitization: the reconstruction of all possible phenomena in the world as digital information and the problems caused by their exchange, combination and utilization. This is a useful perspective for framing the phenomenon of online reputation. Indeed, we can argue that reputation in the cyberworld is a relevant example of digitization of referential trust: it is, as Dellarocas (2003) has clearly put it, a sort of digitization of the word-of-mouth existing in traditional face-to-face networks. For Dellarocas (2003, p. 1409) "Word-of-mouth networks constitute an ancient solution to a timeless problem of social organization: the elicitation of good conduct in communities of self-interested individuals who have short-term incentives to cheat one another". Even if Dellarocas adopts an atomistic and rationalistic perspective that does not capture the whole complexity of this phenomenon, he is right in saying that traditional word-of-mouth networks (what Taddeo considers trust based on communication exchanges) can be considered as an effective solution for building social order. This is possible, in particular, because word-of-mouth networks present two relevant aspects: they can support good and stable reciprocal forms of conduct among unknown participants in social interactions and they can be used for preventing deception and cheating in such interactions. These are also the issues that I consider relevant in terms of a discussion for information ethics when we have a digitization of reputation.

The creation of digitized and internet based word-of-mouth reputation networks are attempts to rebuild the key aspects of traditional networks (good conduct and cheating-prevention). Clearly, however, there are some contextual differences that must be acknowledged between traditional and digitized reputations making this a complex challenge. Indeed, it is sufficiently evident that online reputation partly differs from offline, face-to-face reputation. The problem is that, in contrast to face-to-face interactions, online interactions are dis-embedded from any specific social context (Lash, 2002). It is quite different buying a book on Amazon from a seller whose shop is in another country or instead, buying the same book from a store located in the neighbourhood where you live. In the second case you can see and touch what you are buying, you can interact directly with the seller and ask for advice. The reputation of the local seller is known in the community where you live and you can decide whether to buy also depending on the reputation awarded to the seller by the community. In the case of online interactions, many of the features of face-to-face interactions are missing. Indeed, online you are interacting with the e-commerce portal interface and not directly with the seller or the goods you are purchasing. Furthermore, you will need to place a great deal of trust in the seller and the product you are purchasing as they are described on such an interface. Online Reputation systems have been identified as a solution to bring social order and structure (Farmer & Glass, 2010) in these dis-embedded social interactions: reputation systems collect, aggregate and display ratings, votes, comments and other informational, reputational values (i.e. references) on several aspects of the online behaviour of entities (e.g. a seller, a user, a product). These reputation values are then represented[16] in a variety of ways at the interface level to support online interactions (e.g. online purchases in e-commerce). Users (e.g. customers) will then base their actions (e.g. purchase from an online seller) on the values of reputation displayed on the interface (see figure 1). These informational reputations, like more traditional reputations, are also communitarian values because they are often produced by community of users (e.g. the Amazon or TripAdvisor user communities) by the mediation of reputation systems.

Because of the necessary use of informational technologies such as reputation systems in the cyberworld, the reputation of a user (or other entities) is increasingly a matter of numerical values and aggregation of these values. To capitalize on the terms used in the Call for Papers, user reputation is in many cases a matter of "bit-strings": single numerical digital values and their aggregation in meaningful numerical wholes. Indeed on many

[16] See chapter 7 of Farmer and Glass, 2010 for an overview on reputation display.

web platforms user reputation is measured with points which have been awarded, the number of likes or thumbs-up received, the number of views of an informational content (e.g. a video) and so forth. Reputation systems are then systems that collect these numerical values and aggregate them into synthetic scores and finally disseminate them to other users via the interface (see figure 1). It is this process of collection, aggregation and dissemination of informational reputational items, via reputation systems and their interfaces, that is supposed to sustain users in their good conduct during their interaction with other unknown users. In the same way reputation systems are used as ways to prevent cheating by creating an informational governance mechanism based on referential and distributed trust.

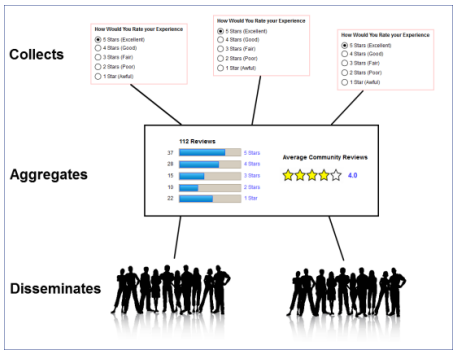

Figure 1: Concept of reputation system: collects, aggregates and disseminates reputation values

The Automated Production of Reputation

Very often, however, numerical reputation values are awarded to the user after the completion of rather mechanical and repetitive actions. As for many other contexts, repetitive and mechanical actions open the space to automation and replacement of human tasks and skills with machines[17]. This is clearly evident, for example, in the case of industrial labour where workers' tasks and skills are often recomposed in large industrial machineries (Marx, 1976). This is a process currently taking place also in other productive sectors, with artificial intelligence replacing skilled workers (Brynjolfsson & McAfee, 2011). What is relevant for this discussion is that the principle also works in digital contexts. In these contexts repetitive actions can be automated by means of autonomous agents, also known as bots: computer programs whose goal is to automate digital relations, replacing and supporting humans in carrying out repetitive and/or complex tasks. In many contexts, bots are legitimate technologies. For instance Wikipedia bots support the Wikipedia community in carrying out repetitive and mundane tasks to maintain the English language Wikipedia. Crawling bots, such as those used by search engines to provide users with up-to-date data about web content, are also legitimate bots. It is this consideration which opens the space for the idea of an automated production of reputation by means of machines (i.e. bots): because in many cases the actions that award reputation to a user are mechanical and repetitive and simply lead to awarding numerical values, these actions can be easily automated with bots. The automated production of reputation is then the production of reputational values with bots. The automated production of reputation is also largely a deceptive process and a violation of the shared rules of online services. Many social network sites for example explicitly forbid the use of bots and other forms of automation. This inevitably leads to a number of problems which I will discuss later in more depth. Firstly, however, it is crucial to better focus on what I mean by automated production of reputation and its deceptive nature. I will introduce three simple examples.

[17] Properly what we have is a replacement of human-labour with machine labour.

Example 1: Massively Multiplayer Online Games

Firstly, I will introduce an empirical case I have studied widely over the last 3 years (De Paoli & Kerr, 2009, 2010, 2012), that of Massively Multiplayer Online Games (MMOGs) and their reputation systems: game player rankings. It was indeed the in-depth study of MMOGs that led me in the first place to observe the existence of an automated production of reputation. This example is therefore of paramount importance for building my case and I will describe it at length.

MMOGs are a sub-sector of the digital games industry (Kerr, 2007). There are hundreds of MMOGs around, with World of Warcraft often cited as the prototypical example. In MMOGs millions of players interact in a persistent Virtual World through their avatars (the in-game persona controlled by the player). A key task of MMOG game-play is that of levelling one's avatar. Avatar levelling is pursued by killing monsters inside the game: killed monsters award so called "experience points" (i.e. simple numerical reputation values) whose accumulation leads to enhancing the player's ranking inside the game. Increases in levels means that the avatar can usually perform better in the game. Rankings are an important type of reputation system in these "competitive communities" (Farmer and Glass, 2010), whose goal is to make users compete with each other[18]. Being at the top of the game ranking makes the user/player the most reputable in the MMOG community as this means that she has performed very well and likely better than her opponents inside the game. In this way MMOG rankings work in exactly the same manner (figure 2) as any other reputation system, as described in figure 1. Comparison between players is based on the accumulated experience points. Ranking at the top makes a player highly reputable for the MMOG community.

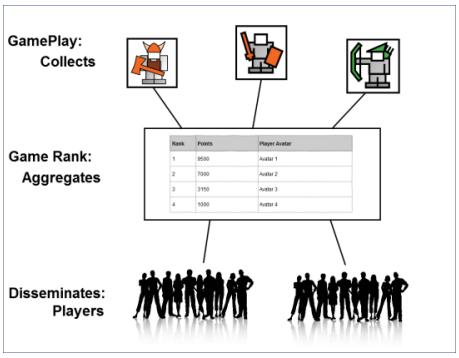

Figure 2: The concept of an MMOG rank/reputation system

Killing monsters inside an MMOG is a very repetitive and mechanical activity that requires little intellectual ability in a situation in which the player is forced to repeat the same actions over and over hundreds of times. This activity is referred to as "grinding" by players and has often being compared in academic literature to industrial labour and Taylorism (Ruggil et al, 2004). Because of this repetitiveness, many MMOGs suffer from the diffusion of bots that can replace players that can be used to automate the "grinding" MMOG levelling[19]. Bots can fully replace the player in killing monsters, in the subsequent accumulation of experience points and in climbing game rankings. With bots, MMOG experience points (i.e. reputation values) are clearly machine-made. This creates an unfair situation between players who play fairly (and need to manually repeat the same actions endlessly) and those who instead fully automate the levelling, since using a bot is a form of cheating. However there are further and more relevant negative consequences.

In the industrial context, automation of work is often seen as a solution to increase productivity: a reduction of the labour-time needed for producing goods. This holds true also for experience points: a direct consequence

[18] Less competitive communities do not use ranks as a reputation system, as their goal is to promote collaboration rather than competition.

[19] Bots in MMOGs are in any case a form of cheating and a violation of the legal documents of the games.

of automation of MMOG levelling is the increase in ”productivity” as greater amounts of experience points can be easily produced in less time with bots compared to what human players can do (De Paoli, 2013). In some games, players estimate that bots can produce, in a few months, an amount of experience points that would take years for a fair player to produce (De Paoli, 2013). Bots produce more points in less time than human players (figure 3: simply provides a qualitative idea of productivity increases, it is not based on real data). This is indeed the main reason why bots are a form of cheating. In this way machine-made experience points flood the game rankings and the rankings themselves can easily become a false representation of the community reputation.

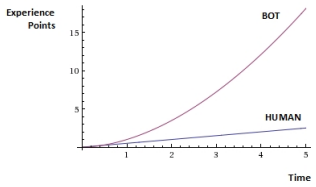

Figure 3: Machine-made versus human made experience points (time on X-Axis, accumulated points on Y-axis)

Example 2: Fake followers on twitter

Another example which displays pretty much dynamics similar to MMOGs is that of false or fake followers on twitter. Because of the wide dissemination of twitter, the phenomenon of fake followers has recently caught the attention of the media[20]. Fake followers are machine-made followers that resemble human twitter users: they have a photo, a bio and, if well-crafted, these fake followers look like real people. In some cases they can also be backed with a bot that can entertain interactions with other users. Basically bots can fabricate these fake followers by creating real-looking twitter profiles, with the aggregation of photos and bios. Bots can produce thousands of these fake followers that can then be later sold over the internet. A twitter user can buy these followers for a few dollars and add them to her public profile (whilst violating the platform’s legal documents). As I will now show, this is a further clear example of automated production of reputation, with distinctive and deceptive outcomes.

The key aspect is that very often the number of followers that a user has on twitter is largely understood as a score or mark of the reputation and social influence of that user. The most influential people or other entities (e.g. companies) in the social media ecosystem are those that are followed by masses of followers. Justin Bieber or Lady Gaga are often quoted as examples, with the latter being the user with largest follower base (around 30 Million). Counting a single follower is again a numerical value of reputation and the total follower base can be seen as an aggregated reputation score. The equation is rather simple then, if the total number of followers — and not the quality of these followers — equals the level of reputation, then adding an increased number — if not masses — of (fake) followers can boost reputation.

As anticipated, the wide dissemination of twitter made the phenomenon of fake followers mainstream. For instance, an event that indicated the problem was the sudden increase in followers of the official twitter account of the 2012 US presidential candidate, Mitt Romney[21], which in a span of about 24 hours had an increase of more than 100 thousand new followers (more than 10% of his total number of followers). The following graph which circulated widely in online newspapers and blogs shows some of the dynamics of this particular case

[20] See also a recent study by Barracuda Labs (http://barracudalabs.com/?p=2989) has shown the depth of this phenomenon.

[21] This was discovered by Barracuda Labs (http://www.barracudalabs.com/).

and, if we focus on it, we can see that it displays pretty similar dynamics to that of figure 3[22]: when automated production of reputation (machine-made fake followers added to one's account) enter the stage we have an exponential increase of values, compared to legitimate manual reputation which has a linear progression. Automated production of reputation allows an increase in productivity and floods the web with machine-made reputational values. This largely undermines the number of followers on twitter as an indicator of reputation.

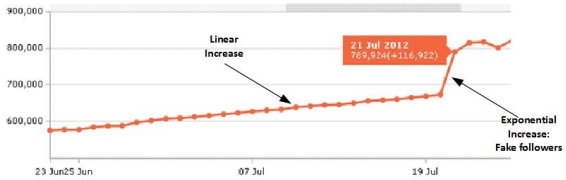

Figure 4: Fake followers and exponential increase in the Mitt Romney case[23]

Recent research conducted by Camisani (2012) on Italian twitter showed that several well-known companies (both national and international companies using Italian twitter with a follower base > 10k) have in their follower base "A very high number of users with "bot" behaviour [...], with percentages in excess of 45%". The same author found that a pretty similar dynamic is displayed by the twitter accounts of Italian politicians. The author of the research concludes (in an interview) that "the number of followers is no longer a valid indicator of the popularity of a twitter user" and this, I would add, is a consequence of a larger process of automated production of reputation. In other words, the falling of the numerical model of reputation is a consequence of automation and deceptive production of reputation.

Example 3: Klout manipulation

An interesting and growing phenomenon of recent years has been the attempt to establish central reputational scores across the web, with Klout (http://klout.com/home) being the most successful so far. A Klout score "is a single number that represents the aggregation of multiple pieces of data about your social media activity" (http://klout.com/corp/how-it-works). Many signals coming from social network sites' (e.g. twitter, google+, Facebook etc.) on a user's activities are aggregated to compose the Klout score. Among them (but the list is more extensive than this), the number of likes or mentions on Facebook, the number of followers or re-tweets on twitter, the connections graph on linkedin and so forth.

An interesting post by Jeff Turner (http://www.jeffturner.info/game-klout/) describes an experiment that the author did to manipulate Klout largely without human intervention. By using automated software he was able to take a "Klout score of 1 to 35 in 30 days, and from 27 followers to 141". He used, in particular, a bot called rep.licants that is able "to simulate the activity of the user, to improve it by feeding his account and creating new contacts with other users". In this case, a Klout score has been entirely produced by bots. Turner reaches an interesting conclusion, namely, that despite prompting the idea that Klout serves as a quality indicator, in fact "Klout doesn't really care about the quality of the 'conversations' it is measuring. Klout can only care about the quantity". Reputation in centralized hubs (like Klout) that mainly leverage quantity and mechanical actions can therefore be easily produced by bots: a further clear example of the automated production of reputation. But there is more.

Having discussed the case of fake twitter followers, we can easily see a preliminary consequence: automated production of reputation in the case of fake twitter followers could also easily lead to increases (if not a major boost) in the Klout score. This is a kind of second-order effect of the automated production of reputation, which

[22] Figure 3 was just a qualitative example, whereas instead Figure 4 is based on real data.

[23] Image from http://www.nbcnews.com/technology/technolog/romney-twitter-account-gets-upsurge-fake-followers-where-928605

does not relate just to fake followers but also re-tweets, Facebook likes and many other signals on social network sites that can be produced easily with bots. The second-order effect here is the situation in which automated production of reputation affecting a service X (e.g. twitter fake followers) also leads to increases of reputation on other services Y (e.g. boost of Klout score) that use social signals to provide reputation scores. In other words, with the central hubs that aggregate reputation values from various services, the negative consequences of the automated production of reputation could become viral for the whole social web.

Discussion: Automated Production of Reputation and its Consequences

The automated production of reputation is an emerging phenomenon touching several aspects of trust in the cyberworld. Very often this process is deceptive in nature — it is a form of cheating in social interplay and could lead to essentially negative consequences. Clearly, there are many different threats that can undermine reputation systems (Carrara & Hogben, 2007), but the automated production of reputation remains a whole new phenomenon whose direct consequences are yet to be explored.

I will now concentrate on a short focused discussion of the implications of automated production of reputation in terms of information ethics. In this regard, I second the approach of taking a critical and emancipatory perspective developing a criticism of possible consequences of the automated production of reputation in the information field, with a particular focus on the collective level. As reputation is indeed a communitarian and collective form of referential trust, this is particularly relevant. How can we then discuss the issue of the automated production of reputation in this frame? And especially its possible negative consequences on the enforcement of rules against violations and the stability of online conduct? These are relevant questions which I will now consider.

Reputation systems are practical, distributed means for internet users to support their actions and decisions. They play a relevant role in the creation of social order in the cyberworld, by engendering trust among unknown participants in online interactions of many sorts (e.g. games, commerce, plain social intercourses). They are based on what Taddeo (2009) calls referential trust: the references about an agent provided by other reliable information sources. The automated production of reputation is a problem that directly attacks this referential process by creating unreliable and fake machine-made references that could not be considered representative of an authentic collective level of trust.

Given that reputation could be easily produced by automated software, indeed a massive amount of machine-made reputation values or references could inundate the web. It takes time and effort to build a legitimate reputation. For instance, it takes quite a long time to climb the game rankings of an MMOG or to build a healthy twitter following. However, if the action that awards reputation can be replaced easily by machines and if the 'productivity' of reputation increases dramatically as a consequence, then the outcome would be that reputation (which can be considered as a form of social capital) will inevitably lose value. This is basic political economy. The value of a product is largely determined by its scarcity on the market. If scarcity is no longer an issue, then the value of the product will fall. The real problem here is for those who legitimately work hard on building their reputation (e.g. fair players, twitter users who personally manage their accounts and so forth) and then see the value of their social capital falling. This is a clear case of unfair competition and a form of cheating in social interactions.

The automated production of reputation could therefore easily lead to 'breakdowns' of reputation systems: direct consequences could be social disorder and inability to represent collective trust within active communities of users. Automated reputation-generation could, in particular, easily undermine users' ability to orient their conduct according to the level of trust being represented by reputations systems. Indeed, if the automated production of reputation becomes a mass phenomenon, then the reputation represented on reputation system interfaces will no longer be representative of the level of trust that a community has placed in a person. In other words, reputation systems will no longer be a distributed social regulatory feature of interplay upon which the user can rely when deciding with whom or what to interact in the cyberworld. This is a second and much larger negative consequence that the automated production of reputation could have on the collective level. In

brief, the automated production of reputation could easily undermine the fact that reputation in the cyberworld is meant to inherit the positive aspects of traditional Word-of-Mouth Networks (Dellarocas, 2003): stability of conduct in social interactions could fall and cheating could proliferate.

When it comes to the enforcement of rules against violations (i.e. cheating), online services have their rights as well as their responsibilities. Certainly the automatic production of reputation in most cases violates the legal documents of online services: Facebook, for instance, prohibits the use of automatic software to "like" information items, online games prohibit the use of bots to play, twitter prohibits the use of bots and the adding of fake followers and so forth. But the violation of legal documents (which I will not discuss here) is definitely not the key relevant negative consequence of automated reputation-production.

Companies (e.g. game companies, social network providers) invest heavily in information security technologies for preventing bots operating within their services or for detecting them with the goal of banning those who use bots. There are many concerns for user privacy and control over the use of these monitoring technologies. Some of these monitoring technologies act ubiquitously in the background, collecting user information and have often been criticized for being too invasive of user privacy. However, technical security solutions are not necessarily the only direction for achieving better services. The mechanical nature of reputation-generation could also be modified, and this would probably reduce the amount of privacy monitoring needed to detect bots.

Conclusion: What's next?

This manuscript is a position paper whose goal is to raise awareness of the problems emanating from the automated production of reputation and to describe some of the immediate foreseeable consequences of this phenomenon. Clearly this is not sufficient, however. Indeed, I largely believe that more needs to be done if we are to understand and tackle the problem. In this conclusion I will briefly touch upon this aspect.

In the first place, a much more solid theoretical definition of the concept of the automated production of reputation will be necessary. The description of the concept provided in this paper merely points to some possible directions of investigation, but clearly it does not have sufficient depth for theorizing about the implications that automated reputation-production has for the 'reputation society' at large. Possible directions for building a more solid theoretical approach have been briefly touched on in this work: the problem of productivity and the replacement of human work by technologies, the issue of the automation of work and the link with current processes of automation, the relations between reputation and the enormous internet-governance problem. Exploring these aspects more deeply, and suitably linking them with the problem of the automated production of reputation will be of paramount importance for research into reputation in the cyberworld.

Secondly, because my working approach is based on empirical research and developing theory as part of empirical data analysis (i.e. a ground-theory-driven approach), it is clear that further empirical research will be required to fully understand the boundaries, implications and evolution of the automated production of reputation. Some research fields have been described in this paper: online multiplayer games, a very dynamic and emerging field; twitter and other social network platforms as places where automated production of reputation acquires the most social form; centralized web reputational hubs that are prone to second-order negative effects depending on the automated production of reputation.

A better theory and more extended empirical fieldwork necessarily constitute the next steps in research into the automated production of reputation.

References

Brynjolfsson, E. & McAfee, A. Race Against The Machine: How the Digital Revolution is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy. [kindle edition]. Lexington MA: Digital Frontier Press, 2011.

Camisani Calzolari, Marco. Analysis of Twitter followers of leading international companies. Quantitative and qualitative study of behaviours demonstrated by humans (users which are presumably real) or by bots (users which are presumably fake)", 2012. URL: http://www.camisanicalzolari.com/MCC-Twitter-ENG.pdf

Capurro, Rafael, "Towards an Ontological Foundation of Information Ethics," Ethics and Information Technology 8, 2, (2006): 175–86Dasgupta, Partha. Trust as a commodity. In D. G. Gambetta (Ed.), Trust: 49-72. New York: Basil Blackwell, 1998.

Carrara, Elisabetta and Hogben Gilles. Reputation-based Systems: A Security Analysis. ENISA Position Paper No. 2, October 2007.

Dellarocas, Chrysantos. The Digitization of Word of Mouth: Promise and Challenges of Online Feedback Mechanisms. Management Science, 49(10):1407-1424, 2003

Dellarocas, Chrysantos. "Designing Reputation Systems for the Social Web". In Building web reputation systems, edited by Farmer, F.R. & Glass, B. 3-12. Cambridge, MA: O'Reilly, 2012.

De Paoli, Stefano and Kerr Aphra (2009). "We Will Always Be One Step Ahead of Them": A Case Study on the Economy of Cheating in MMORPGs . Journal of Virtual Worlds Research, 2(4). 2009 URL: http://journals.tdl.org/jvwr/index.php/jvwr/article/view/865/630

De Paoli, Stefano and Kerr Aphra. (2010). The Assemblage of Cheating: How to Study Cheating as Imbroglio in MMORPGs. FibreCulture Journal, Issue 16. 2010. URL: http://sixteen.fibreculturejournal.org/the-assemblage-of-cheating-how-to-study-cheating-as-imbroglio-in-mmorpgs/

De Paoli, Stefano and Kerr Aphra. On Crimes and Punishments in Virtual Worlds: bots, the failure of punishment and players as moral entrep. Ethics & Information Technology, 14(2): 73-87, 2012.

De Paoli, Stefano. (2013). Bots and the Speed of Levelling: How bots change the player perception of MMORPGs. Full paper unpublished.

Farmer, F. Randy. and Glass, Bryce. Building web reputation systems. Cambridge, MA: O'Reilly, 2010.

Jøsang, Adun, Ismail Roslan and Boyd Colin. "A survey of trust and reputation systems for online service provision." Decision Support Systems, vol. 43 (March 2007): 618-644, March 2007.

Kerr, Aphra. The Business and Culture of Digital Games: Gamework / Gameplay. London: Sage, 2007.

Lash, Scott. Critique of Information. London: Sage, 2002.

Luhmann, Niklas. Trust and Power. New York: John Wiley, 1979.

Marx, Karl. Capital: Critique of Political Economy, Volume 1. trans. B. Fowkes. London: Penguin, 1976.

Masum Hassan and Tovey Mark. (eds.). The Reputation Society: How Online Opinions are Reshaping the Offline World. The MIT Press, 2012.

Ruggill, Judd Ethan., McAllister, Ken S. & Menchaca, David. (2004). The gamework. Communication and Critical/Cultural Studies 1(4), 2004: 297-312.

Taddeo, Mariarosa. Defining Trust and E-trust: From Old Theories to New Problems. International Journal of Technology and Human Interaction (IJTHI) 5(2), 2009: 23-35.