Predicting the spread of AI-applications

Purpose and aims

To understand the impact of AI we must first understand what AI applications will be used and by whom. The goal of the project is therefore to develop a framework with which we can predict the spread of AI applications.

Specifically, we aim to:

- develop and evaluate a methodological framework for predicting the spread of AI-applications in different demographics within the coming decade;
- apply this framework to compare the spread of a set of AI technologies and then identify which applications that are likely to spread the fastest;
- identify applications where the social impact is likely to be large, and thus highlight areas in which proactive action is the most urgent - both in terms of limiting and enhancing spread.

Better predictions about the spread of AI applications would enable society to take full advantage of AI’s opportunities and to act proactively to mitigate its potential damages.

This is not only very important but there are several reasons why this is feasible, at least to the level of precision needed to significantly improve on current assessments. Primarily, while general predictions have turned out to be beyond the capabilities of experts, research in psychology shows that experts are good at judging underlying features on which predictions can be built. Secondly, previous research within methods for analyzing technology acceptance and market dynamics provide compelling evidence that it is feasible to predict weather a specific technology will be used based on assessments of its characteristics. Furthermore despite their ubiquity, almost all current AI applications rely on training deep neural networks and therefore share some notable characteristics: a reliance on data that are often gathered through sharing of user information; predictions using a large number of input variables; an ability to improve over time, given a well-defined goal function; functional perception and control across different types of input, and an automation of tasks and consequent submission of (human) user control, e.g., Autonomous Emergency Braking (AEB) systems in passenger cars. The possibility to build a common framework for evaluating AI is motivated by the fact that current applications typically share these attributes.

Finally, the components and the assumptions of the framework would be transparent and can therefore be improved as more data becomes available. Transparency will enable us to improve the framework within this project, but it can also be improved further by future research as well as be extended to other contexts or time frames. An important point is that the project would begin at a time when AI usage is likely to set off. In this sense, the coming five years provide a window of opportunity for developing and evaluating predictive models of AI spread.

State-of-the-art

Artificial Intelligence (AI) here refers to computational practices that allow machines to perform cognitive functions, such as perception, learning, and language use. Applications of AI typically depend on algorithms that are based on large amounts of data to enable automated functions, predictions and decision support. Technologies based on AI have recently advanced fast thanks to drastic improvements in statistical methods and hardware. For instance, during the last 20 years there has been an eightfold increase in the number of research papers, and venture capital investments in AI-based businesses have grown by 450% during the last five years [1].

It is therefore reasonable to expect a corresponding increase in the use of new applications based on AI. Evidence of this development is already apparent in retail, transport, finance, the public sector, etc. Some AI applications are likely to have a transformative impact on society in the coming decades [2]. A prerequisite for social impact of a technique is that it will actually be used. To systematically investigate the social impacts of AI we need to be able to compare, quantify and predict the spread of different types of AI applications in various demographics.

Previous studies have attempted to predict the spread of new technology using top down approaches [3, 4, 5]. However, such efforts have often failed [6]. There is also a substantial body of research on foresight methods in general [7, 8, 9, 10]. However, these methods have yet to be applied systematically to the domain of AI.

Importantly, none of the frameworks that we are aware of capture the dynamics needed to predict whether a new AI application will be used widely or not. More specifically, there is no framework that considers 1) the specific characteristics of AI systems (discussed above); 2) both market dynamics and usage when the product has been purchased; and 3) the need for prediction of time-frames rather than qualitative assessments. We have concluded that there seem to be an important gap in the scientific literature regarding methods to quantitatively predict the spread of AI applications in different demographics in the coming decade.

Our focus on the near-term spread of AI among consumers is distinct from two other questions about future. AI that have received more attention. One question is what impact AI will have on the labor market [11], focusing on how future AI applications may come to replace human labor. In this literature the focus is on decisions made by the industry which depend almost exclusively on cost benefit analysis in contrast to our project where the decisions are made by individuals and therefore rely on a large host of other factors such as norms, perceived usefulness etc. Another question is how to avoid the worst-case scenarios of how future AI may impact on humanity, such as the threat of a misaligned superintelligence, wars resulting from AI arms races, or AI-induced authoritarianism [12, 13]. The near-term development and the long-term development of AI are both important but here our focus is on the former.

Project description

The goal of the project is to develop a framework for prediction. By a framework we mean an algorithmic description of how to predict the spread of a specified application within a specified demographic. This includes a detailed description of the data needed, typically feature scores generated by assessors, as well as a guide for how to collect that data in terms assessors of different features. It also consists of a mathematical model to turn that data into predictions. It will account for delimitations that concern under which circumstances it is appropriate to use the different methods. Finally, it will include a set of case studies that will exemplify the use of the framework and work as benchmarks for the time assessments.

The ambition is that the framework will be applicable to many different types of AI applications. To begin with, however, a set of representative applications will we chosen, on which the framework will be developed. Particular focus will be put on features selection and model development, as addressed below. Subsequently, the framework will be applied to other AI applications - both established systems, for evaluation, and new ones, in order to make predictions.

Theoretical background

Experts are often good at estimating the features that are relevant to the spread of a technology [14], even though they are generally not capable of predicting the speed of the spread itself. From this fact, we will design a structure that allows us to weigh together expert and user judgments of specific features in order to produce concrete predictions about the speed at which a particular AI technology will spread in a particular demographic. To build the framework, we draw on four scientific discourses in particular: near-future scenario generation (e.g., used in [15]), the Technology Acceptance Model (TAM) [16], Discrete choice theory [17], and cultural evolution models [18], as described below.

Near-future scenario generation entails consulting experts and non-experts and have them articulate situations that are likely to occur in the near future (10-15 years ahead). One implementation is the scanning-method, in which scenarios are devised by individuals who read news articles relating to, say, new technology or emerging social problems [15].

TAM uses surveys to measure perceived usefulness and perceived ease-of-use, which in turn predicts adoption rates of Information Technology (IT) [19]. Applying TAM in 1989, Davis et al. [20] concluded that computer usage could be predicted quite well from users' intentions, and, in turn, that their intentions may be predicted by perceived usefulness and perceived ease of use. An important conclusion was that simple, yet powerful, models could predict user acceptance relatively efficiently and thus guide managerial interventions [19]. Since then, TAM has been applied by researchers and practitioners to a range of different IT systems and populations, and its validity and reliability have been confirmed [20].

Discrete choice theory [17] has proven to be a powerful tool for predictive analysis of market dynamics. In particular, it accounts for two assumptions that we draw on in the proposed project: that choice depends on a finite set of measurable attributes of the available alternatives, and that related statistical predictions can be made based on a linear combinations of these data. This is a relevant approach given that a choice between products has to be made, and that the market size is exogenous; for example, it might be applicable to predicting the choice between a traditional vacuum and a robot vacuum, conditional on that one of them will be purchased.

Cultural evolution modelling is focused on the spread of cultural traits through a population [22]. These entail mathematical formulation of dynamical systems based on assumptions about the frequency of interactions between different individuals and the probabilistic effects of those interactions. The dynamical systems then describe how the traits spread through the population. This enables the transformation of estimates of specific attributes such as perceived usefulness to be turned into predictions about the rate by which the cultural trait will spread [23]. While it has been used successfully to predict the spread of other cultural traits, for instance how moral opinions change [24, 25] it has not yet been adopted to understand the spread of AI applications.

The above set of existing approaches to prediction constitute a starting point and a benchmark for our project. Our aim is to develop a framework that is more precise than the methods for scenario generation, as it will be developed and evaluated based on quantitative data. Further, it will account for market dynamics in addition to technology acceptance, in contrast with TAM, which concerns the use of technology that is available. Moreover, it will be more generic than the discrete choice models that have been applied to predict market dynamics in a certain product category, such as passenger cars [26]. Finally, it will draw on cultural evolution modelling to make distinct predictions about the actual, and not just relative, speed of the spread of different AI applications.

To evaluate the social impact of different AI applications in various demographics, we will use a well-established research method for systematic forecasting called Delphi, or Estimate, Talk, Estimate (ETE) [35, 36]. In ETE, a group of experts make assessments that are structured in a particular way. The experts answer questions independently in the first round, and facilitators will then provide them with anonymized summaries of these assessments in later rounds. The experts will reveal the way they have reasoned and they will take part in each other's reasoning. There are several rounds, and it is expected that the assessments will converge with time, with the mean score typically providing the final prediction. In this project, the experts on potential impact will be researchers within psychology and social science, which is motivated considering the purpose of evaluating the social impact of different AI technologies.

Methodology for developing the framework

- **Select key features**

We aim to identify a set of features that differ between AI applications and which are likely to impact their spread. The features will be selected based on several criteria. First, they will include characteristics that previous literature has shown to impact on various determinants of technological spread, such as perceived usefulness [19], norms [27] and advertisement investment [28]. Second, they will be features that experts, users and potential users are capable of giving good assessments of. For this we will be guided by psychological research of what experts can and cannot accurately assess [14]. Third, the selected features will be inspired by models within cultural evolution [29, 30]. Such models simulate decisions and interactions of all individuals in a population, relying on assessments of how often a demographic evaluates the choice of application, the extent to which choices depend on the people they interact with, and the pattern of those interactions. When evaluating the features experts and users will also be able to provide features they think are relevant but missing so that they can be included in the framework at later stages

Experts, users and potential users will be selected depending on their expertise relating to these features. They will include AI developers at Google DeepMind, industry experts for the particular cases we are interested in, and representative samples of potential users from the demographics for which the application is intended. As we aim to develop a tool that will be readily available for each feature we will make clear the reason behind the choice of assessor.

- **Design the model**

We will evaluate different ways of combining the feature scores, given by the assessors, into predictions. This includes the definition of a set of models that weight the assessors' evaluations to derive predictive models. It means developing several models which differ in generality as well as predictive precision.
We will divide the AI-applications into two categories: applications that are readily available at almost no monetary cost (e.g., Google Maps) and applications that depend on purchasing power (e.g., robot vacuums). These two categories will be assessed separately. The first will mainly depend on technology acceptance, while the second will depend on market dynamics and willingness-to-pay for the AI feature.

The more general models mainly consist of linear combinations of the feature scores. Previous research has shown that linear unit weights perform very well for this size of dataset [33, 34]. Potentially, regression will be used, assuming that we are able to assess many AI applications providing sufficient data to enable statistical modeling. For each application this will give us a combined measure of how fast it will spread which can be tested against actual data.

Further, we will develop more specific and advanced mathematical models of the type developed within cultural evolution [31, 32]. Cultural evolution models are dynamical systems built on micro scale events. Examples of events would be someone telling their friend about the robot vacuum they just bought or someone seeing an ad for a new service. The model then lets us aggregate the microscale events into macroscale predictions about the spread of the applications. These models let us predict the speed of the spread relative to how many individuals currently use the application. This is more information than can be obtained through linear estimates, which only provide a combined score for each application. However, as these models rely on more assumptions than the linear combinations, they are likely to require more data testing to ensure that the right assumptions are being used before they give accurate predictions.

All these models constitute drastic simplifications of the real world, but they are expected to capture the big picture, as such models have successfully done in other domains [34]. They will yield specific time predictions concerning when a particular AI application will be used by a certain proportion of potential users in a particular demographic.

Illustrative example

As an example, consider an AI based Virtual Home Assistant and the demographic of young adults in Stockholm County: First, we intend to specify a number of concrete attributes, such as price (e.g., 500 SEK); software/hardware (software); type of interface (voice); advertisement investments (e.g., 5 million SEK/year); frequency of replacement (e.g., every three years); type of user data that will be logged (intimate, as it concerns personal behavior and movie preferences); and, is it sensitive to failure? (no). Subsequently, we would identify and target two representative samples within this demographic, potential and actual consumers. Then, we would ask them to assess the following features from a scale from 1-10: Is using this system free from effort (perceived ease-of-use)? (e.g., 5). Does the system facilitate everyday tasks (perceived usefulness)? (e.g., 9). Does it strengthen social relationships? (e.g., 3, due to sharing of movie tips, for example). Then we will use a set of mathematical models that will weigh together the listed attributes to predict spread. In its simplest form this means turning all features into scores between 1-10, where higher means increasing speed, and then adding them up. The most advanced form will be cultural evolution models detailing how decision-making processes depend on the feature scores and social interactions. The predictions will then be evaluated using variables representing spread, for instance sales data or mentions in social media.

- **Predict the spread of AI applications**

Once the framework is developed it will be used to predict the future spread of 20 recently introduced AI applications. We will make these predictions as early on as possible to be able to evaluate them within the timeframe of the project. A set of AI applications with different characteristics and potential user groups will be selected in order to test the general framework. The final list will be developed with input from AI developers and Industry experts. However, the preliminary list includes the following AI applications: deep fakes, i.e., fake videos that seem real and which may be used as political misinformation; AI for habit formation, AI applications for learning, scam baiting; AI therapists; voice assistants in smartphones, such as Apple Siri; and voice assistants in the home, such as Amazon Alexa or Google Assistant. The spread of these applications will be evaluated using the following measures: sales statistics, mentions in social media, download frequency of applications, as well as surveys measuring actual usage among current users, when necessary. We will evaluate these early predictions at the end of the project.

- **Improving the framework using existing data**

When the early predictions are in place the framework will be improved by applying it to older AI applications for which the spread can already be measured using existing data. These tests will allow us to select the best performing models and exclude the features that do not sufficiently improve the predictions. Examples of AI technologies that have been available on the market for some time includes robot vacuums, language translation pens, and matching algorithms for dating.

- **Evaluate the spread and potential social impact of AI-applications**

We will then use the improved framework to predict the spread of a larger set of AI applications for the coming decade. To better understand which applications policymakers need to focus on we will use the Delphi method [35, 36], discussed above, to specify their impact. In this project, the experts assessing impact will be researchers within psychology and social science, who have the necessary background for those judgments. We will also consult with philosophers who work on relevant ethical dilemmas to understand the ethical implications. Combined, these assessments will provide us with information regarding which AI applications are likely to have the most considerable impact on society and whether those effects are wanted, unwanted or ethically indeterminable.

- **Evaluate the early predictions empirically**

In the last part of the project, we expect to gather sales data and other statistics regarding usage regarding the applications that we made predictions for at the beginning of the project. In this way, the strengths and limitations of the framework will be identified.

Time plan

In summary, the project can be divided into five distinct parts:

1. A detailed description of a framework for assessing the spread of AI technologies.
2. Application of the framework to predict the spread of a set of relatively new AI applications.
3. Application, evaluation and continued development of this framework using historical data.
4. Application of the improved framework to predict the spread of a larger set of AI applications, as well as analyses of the social impacts of both fast-spreading and other applications in order to identify the need to enhance or limit spread.
5. An evaluation of the predictions made in point two based on empirical data.

We plan to develop the predictive framework during the first year, to make testable predictions for the future spread of AI applications in the year after that, and to test and refine it using historical cases in the year after that. In the fourth year we will use the improved framework to predict the spread of a larger set of AI applications and the type and size of their social impact. This will enable evaluations of whether there is a need to increase or limit the spread. As a few years will have passed at this point, the project will conclude with an evaluation of the predictions made in the second year.

Publications and outreach

The project will lead to several traditional peer reviewed publications detailing the framework, the historical case studies as well as the predictions of the future. However, as these findings are time sensitive we will present the findings directly to the general public as well as to policy makers. This will be done through seminars as well as through the publications of open reports both of which will be greatly helped by the Institute for Futures Studies' strong infrastructure for outreach.

Scope and delimitations

An important feature of AI is its generalizability. The spectra of potential applications and user groups are therefore extremely wide. Thus, it is essential to clearly define the scope of the proposed project, which will have these delimitations: First, the category of users that we are focusing on are consumers, i.e., individuals who may/may not use a specific technology in their private lives, rather than organizations or companies (private or public). The AI applications that we will study will be both in the form of software (e.g., voice assistants) and hardware (e.g., robot vacuums). Second, we are focusing on the spread within different demographics in Sweden and the US and we will assume that the current legal framework in the respective countries will be valid. However, we plan to consider AI applications that may diffuse in society despite legal restrictions, such as scam baiting. Third, the time frame is one decade, so that the framework will be designed to make predictions up to 2030. Fourth, we will study AI applications that are available in the market today or are likely to become available within a few years. Lastly, we will consider applications and not specific companies' versions of applications.

Significance and scientific novelty

Current predictions about the spread of AI are poor and improving them would provide a large contribution to the field, even if these predictions would end up being far from perfect. To ensure feasibility, we have limited the scope of the project in several ways (see the methods section). Even with these limitations, the proposed framework would constitute a large improvement of current prediction techniques, not least because it will be constructed so that it can be used by non-researchers and because it includes several modelling paradigms that can be selected dependent on the data available. Finally, the framework is constructed so that it can be improved as more data becomes available. Such a framework would not only be of great importance to society but it would constitute a new paradigm in the research on the impact of AI as well as contribute to future methodological research on predictive tools for technological foresight and social change.

Understanding the spread of AI applications is necessary to consider the impact of those applications, for example concerning the ethical implications of AI. Improving predictions can therefore be seen as a necessary precursor to all research trying to assess the impact of AI, including its ethical implications. We therefore expect that the framework will be used in future research that evaluates the social, societal and environmental impacts of different types of new AI. We are ourselves very engaged in such applications, and this helps to ensure that the framework will be relevant and usable.

However, while the significance for other research is large, the significance for society is likely to be even larger. Considering the magnitude of the potential impact of AI, a greater understanding of which effects are more imminent - and therefore in need of immediate attention - would allow society to act in a timely manner. This is likely to be key to ensure that the benefits of AI are maximized, and the harms minimized.

Project organization

The team consists of researchers who will study the development of AI in general, model the spread of selected AI applications, and assess the social and cultural impact of these applications. The project will be situated at the Institute of Futures Studies (IF), where it will be incorporated in the existing team for research on the societal impact of new technology, led by Strimling. The project will capitalize on IF's strong connections to the industry and policymakers, as well as the institute's extensive experience with research dissemination beyond academia.

PI Pontus Strimling is deputy director at the Centre for Cultural Evolution at Stockholm University as well as a research director at the Institute for Futures Studies. His expertise is interdisciplinary, including a PhD in mathematics, a docent in economics, and postdocs in Anthropology (UC Davis), Economics (IAS Princeton) and Political Science (Indiana University). Strimling currently holds the prestigious Wallenberg Academy Fellowship. His research has explored how cultural traits change over time including making predictive models for how values change. He has published extensively on this in journals such as PNAS and Nature Human Behaviour. Strimling will be involved in all parts of the project.

Emma Engström is a post-doctoral researcher at the Institute for Futures Studies. She has experience from a range of research environments, as a Fulbright Visiting Student Researcher at the Lawrence Berkeley National Laboratory at the University of California, a research assistant at the University of Chicago, and an honorary PhD student (Excellens-doktorand) at the KTH School of Architecture and the Built Environment. Engström holds a MSc in Engineering Physics, with a specialization in machine learning and statistics, and a PhD within Predictive modelling and environmental risk assessment. She will develop the models that allow us to make predictions out of the expert and user scores.

We will also recruit a postdoctoral researcher from psychology, with expertise on the quality of expert judgments or prediction methods. The position will be internationally recruited with good terms of employment to ensure a high-quality candidate.

International and national collaboration

The project relies on several already developed international collaborations. Throughout the project, we will consult with our contacts in two of the world leading research institutes for understanding AI impact: the Future of Humanity Institute in Oxford (connected through professor Anders Sandberg and director of AI governance professor Allan Dafoe) and the Centre for the Future of Intelligence at Cambridge University (connected through Dr. Jess Whittlestone). Finally we have strong connections with the world's leading research center for AI development, DeepMind (connected through staff researchers Timothy Lillicrap, Gregg Wayne and Neil Rabinowitz). This will ensure that we take the latest AI technology into account as well as have access to AI developers for expert judgments when needed.

Other applications or grants

Pontus Strimling already receives funding as a Wallenberg Academy Fellow (How do social norms change? 2017.0257) and will only need funding for 40% of his time in the project. . If we receive funding for this project IF has agreed to contribute with 50% of Emma Engström's salary. We have applied for funding for smaller versions of this project from Riksbankens Jubileumsfond and the Swedish Research Council but no decision has been made yet.

References

1. Yoav Shoham, Raymond Perrault, Erik Brynjolfsson, Jack Clark, James Manyika, Juan Carlos Niebles, Terah Lyons, John Etchemendy, Barbara Grosz and Zoe Bauer, "The AI Index 2018 Annual Report", AI Index Steering Committee, Human-Centered AI Initiative, Stanford University, Stanford, CA, December 2018.
2. Economist Intelligence Unit. 2018. Intelligent economies: AI's transformation of industries and society. White paper.

3. Rogers, E. M. (2010). Diffusion of innovations. Simon and Schuster. Richerson, P. J., Christiansen, M. H., & Lupp, J. (Eds.).
4. Eveland, J. (1986). Diffusion, Technology Transfer, and Implementation. *Knowledge*, *8*(2), 303-322. doi: 10.1177/107554708600800214
5. Grubler, A. (1990). The Rise and Fall of Infrastructures: Dynamics of Evolution and Technological Change in Transport. Springer Verlag, Heidelberg.
6. Curran, J. (2012). Reinterpreting the internet: James Curran. In Misunderstanding the internet (pp. 9-39). Routledge.
7. Avin, S., Wintle, B., Weitzdörfer, J., Ó hÉigeartaigh, S., Sutherland, W., & Rees, M. (2018). Classifying global catastrophic risks. *Futures*, *102*, 20-26.
8. Mellers, B. A., Baker, J. D., Chen, E., Mandel, D. R., & Tetlock, P. E. (2017). How generalizable is good judgment? A multi-task, multi-benchmark study. *Judgment and Decision making*,
9. Rowe, T., Beard, S. (2018). *Probabilities, methodologies and the evidence base in existential risk assessments.* Working paper. Centre for the Study of Existential Risk, Cambridge, UK.
10. Voros, J. (2003). A generic foresight process framework. *Foresight*, *5*(3), 10-21.
11. Frey, C., & Osborne, M. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting And Social Change*, *114*, 254-280.
12. Bostrom, N., 2014. Superintelligence: Paths, Dangers, Strategies. OUP Oxford, Oxford.
13. Dafoe, A. (2018). AI Governance: A Research Agenda. Unpublished report, University of Oxford. Available at: https://www.fhi.ox.ac.uk/wp-content/uploads/AI-Governance_-A-Research-Agenda.pdf
14. Kahneman, D., & Klein, G. 2009. Conditions for intuitive expertise: a failure to disagree. *American psychologist* 64(6), 515-526.
15. Honda H., Washida Y., Sudo A., Wajima Y., Awata K., Ueda K. (2017) The difference in foresight using the scanning method between experts and non-experts. *Technological Forecasting and Social Change* 119: 18-26.
16. Davis, F. (1989). Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology. *MIS Quarterly*, *13*(3), 319.
17. McFadden, Daniel L., 1984. (1984) Econometric analysis of qualitative response models. Handbook of Econometrics, in: Z. Griliches & M. D. Intriligator (ed.), Handbook of Econometrics, edition 1, volume 2, chapter 24, pages 1395-1457 Elsevier.
18. Eriksson, K., & Strimling, P. (2015). Group differences in broadness of values may drive dynamics of public opinion on moral issues. Mathematical Social Sciences, 77, 1-8.
19. Legris, P., Ingham, J., & Collerette, P. (2003). Why do people use information technology? A critical review of the technology acceptance model. *Information & management*, 40(3), 191-204.
20. Davis, Fred D., Bagozzi, Richard P. and Warshaw, Paul R. (1989) User Acceptance of Computer Technology: A Comparison of Two Theoretical Models. *Management Science*, Vol. 35, No. 8, pp. 982-1003.
21. Davis F.D., Venkatesh V. (1996) A critical assessment of potential measurement biases in the technology acceptance model: three experiments. *International Journal of Human-Computer Studies* 45: pp. 19-45.
22. Richerson, P. J., & Boyd, R. (2008). Not by genes alone: How culture transformed human evolution. University of Chicago press.
23. Strimling, P., Enquist, M., & Eriksson, K. (2009). Repeated learning makes cultural evolution unique. Proceedings of the National Academy of Sciences, 106(33), 13870-13874.
24. Eriksson, K., & Strimling, P. (2015). Group differences in broadness of values may drive dynamics of public opinion on moral issues. Mathematical Social Sciences, 77, 1-8

25. Pontus Strimling, Irina Vartanova, Fredrik Jansson, & Kimmo Eriksson (in press). The connection between moral positions and moral arguments drives opinion change. To appear in Nature Human Behaviour.
26. Train, K.; Winston, C. (2007). Vehicle Choice Behavior and the Declining Market Share of US Automakers. *International Economic Review*. 48 (4): pp. 1469–1496.
27. Hsu, C. L., & Lu, H. P. (2004). Why do people play on-line games? An extended TAM with social influences and flow experience. *Information & management*, 41(7), 853-868.
28. Assmus, G., Farley, J. U., & Lehmann, D. R. (1984). How advertising affects sales: Meta-analysis of econometric results. *Journal of Marketing Research*, 65-74.
29. Boyd, R., & Richerson, P. J. (1988). Culture and the evolutionary process. University of Chicago press.
30. Eriksson, K., & Strimling, P. (2015). Group differences in broadness of values may drive dynamics of public opinion on moral issues. *Mathematical Social Sciences*, *77*, 1-8.
31. Wainer, H. (1976). Estimating coefficients in linear models: It don't make no nevermind. *Psychological Bulletin*, *83*(2), 213-217.
32. Dawes, R. (1979). The robust beauty of improper linear models in decision making. *American Psychologist*, 34(7), 571-582.
33. Cohen, J. (1990). Things I have learned (so far). *American Psychologist*, *45*(12), 1304-1312.
34. Strimling, P., de Barra, M., & Eriksson, K. (2018). Asymmetries in punishment propensity may drive the civilizing process. *Nature Human Behaviour*, *2*(2), 148-155.
35. Helmer-Hirschberg, O. (1967). A systematic use of expert opinions. RAND Paper. Document Number: P-3721.
36. Nicholas Rescher (1998). Predicting the Future: An Introduction to the Theory of Forecasting. State University of New York Press, Albany, 1998.