



Sensitivity analysis of agent-based models: a new protocol

Emanuele Borgonovo¹ · Marco Pangallo² · Jan Rivkin³ · Leonardo Rizzo⁴ · Nicolaj Siggelkow⁵

© The Author(s) 2022

Abstract

Agent-based models (ABMs) are increasingly used in the management sciences. Though useful, ABMs are often critiqued: it is hard to discern why they produce the results they do and whether other assumptions would yield similar results. To help researchers address such critiques, we propose a systematic approach to conducting sensitivity analyses of ABMs. Our approach deals with a feature that can complicate sensitivity analyses: most ABMs include important non-parametric elements, while most sensitivity analysis methods are designed for parametric elements only. The approach moves from charting out the elements of an ABM through identifying the goal of the sensitivity analysis to specifying a method for the analysis. We focus on four common goals of sensitivity analysis: determining whether results are robust, which elements have the greatest impact on outcomes, how elements interact to shape outcomes, and which direction outcomes move when elements change. For the first three goals, we suggest a combination of randomized finite change indices calculation through a factorial design. For direction of change, we propose a modification of individual conditional expectation (ICE) plots to account for the stochastic nature of the ABM response. We illustrate our approach using the Garbage Can Model, a classic ABM that examines how organizations make decisions.

Keywords Agent based modeling · Sensitivity analysis · Design of experiments · Total order sensitivity indices

1 Introduction

Simulation modeling has become increasingly important in studying organizational behavior (Carley 2002; Harrison et al. 2007). Among the several simulation approaches available, agent-based models (ABMs) play a special role (Prietula and Carley 1994). Since early classics such as the Garbage Can Model

✉ Emanuele Borgonovo
emanuele.borgonovo@unibocconi.it

Extended author information available on the last page of the article

(Cohen et al. 1972), researchers have increasingly used agent-based simulation techniques to address relevant organizational, strategic and operational questions (Prietula et al. 1998; Luo et al. 2018; Barnes et al. 2020). As Anderson (1999) emphasizes, ABMs allow researchers to examine open systems—common in management situations—whose behavior cannot be described analytically by equations derived from energy conservation principles or decision-theoretical axioms.

We can broadly define ABMs as computational models in which aggregate outcomes emerge from agents' properties, behaviors and interactions, without the imposition of any top-down constraint. This makes ABMs extremely flexible, as it is relatively easy to vary the building blocks of an ABM. Yet the flexibility of ABMs does not come for free: results are often hard to interpret (Rahmandad and Sterman 2008). Lacking a closed-form expression that links inputs to outputs, agent-based modelers often struggle to assess whether the results of their ABMs are robust and their conclusions are valid.

Addressing these issues is one purpose of sensitivity analysis. However, as our literature review shows, sensitivity analysis is often omitted or only performed in a partial, ad-hoc manner in work that employs ABMs. This is understandable as ABMs raise a number of challenges for sensitivity analysis. For instance, the flexibility of ABMs makes it possible to vary not only parameters but also agents' behavioral rules. While sensitivity analysis with respect to parameters is relatively straightforward, it is less clear how to perform sensitivity analysis on elements that do not belong to a well-defined numerical space. Yet varying non-parametric elements, not just parametric elements, is surely essential in any sensitivity analysis of an ABM. After all, ABMs are often regarded as axiomatic systems that generate outcomes which should be regarded as propositions or theorems (Gallegati and Richiardi 2009). We cannot test the robustness and stability of ABMs' theoretical findings without varying their non-parametric elements. (Henceforth, we use the term *non-parametric elements* to indicate elements of an ABM simulator other than parameters; thus, we do not use the term *non-parametric* in the sense of models that are free of assumptions about the frequency distributions of the variables being assessed, a usage found in statistics.)

To strengthen future work that employs ABMs, we develop a systematic process for conducting sensitivity analysis on ABMs. Our first contribution is to propose a general conceptual structure for ABMs that identifies their “moving parts”—the elements that can be subjected to a sensitivity analysis. This conceptual structure is helpful for researchers to catalogue all the assumptions that underlie their model and to assess the breadth of their sensitivity analysis.

It is a common misconception that sensitivity analysis is just about showing that the main conclusions of a paper are robust to a range of assumptions. Robustness is only one of several goals. Sensitivity analyses can reveal which elements of a model, or combinations of elements, have the greatest impact on the results, and how strongly various elements interact with each other to influence model outcomes. In our second, more technical, contribution, we propose a design that allows the simultaneous variation of parametric and non-parametric ABM elements to address simultaneously multiple aims of a sensitivity analysis.

A further challenge of sensitivity analysis, which usually involves running many variants of a core model, is the visualization of results. Consequently, we devote special attention to graphical representation. In this context, a third contribution of our paper is a modification of the well-known individual conditional expectation (ICE) plots, a modification that accounts for the stochastic nature of ABM's response, by adding a test on the difference of mean values to validate the statistical significance of the insights regarding direction of change. We call the modification "stochastic individual conditional expectation" plots, or S-ICE (see Appendix 3 for greater details).

To illustrate our approach, we apply our sensitivity analysis process to the Garbage Can Model (GCM) in the implementation of Fioretti and Lomi (2008, 2010). We choose this ABM because it is well-known in the management sciences and the software implementation is publicly available. We show that a non-parametric element is the most important driver of the results, demonstrating that a sensitivity analysis focused only on parameters misses important aspects of ABMs. We then show that interactions between model elements are relevant and offer new insights on the managerial interpretation of the findings. These insights would be missed by approaches to sensitivity analysis that vary "one parameter at a time." As we illustrate with our analysis of the GCM, sensitivity analysis can reveal new research insights and point to fruitful extensions for future modeling.

2 Related literature

Our paper builds on two related literatures, one on agent-based modeling and another one on sensitivity analysis. These subjects are vast and this review cannot claim to be exhaustive. In this section, we provide synthetic overviews of these fields in order to highlight the research gap that we address. Appendix 1 provides a more comprehensive review of the literature, with several additional details.

2.1 Agent-based models in the management sciences

ABMs have been used to address a number of topics in the management sciences. A first use of ABMs is for theory development. An early classic is the garbage can model (GCM). The model arose in association with the "garbage can" theory of organizational decision-making (Cohen et al. 1972). In the garbage can theory, organizations are viewed as "collections of choices looking for problems" (Cohen et al. 1972). Each opportunity for choice is like a garbage can into which problems, solutions, and decision makers have been dumped, and what emerges from the collection is "organized anarchy." To any scholar who has spent time in real organizations, the garbage can is a welcome complement to traditional models. Over the years, it was noted that the computer model presented in the original article did not reflect the corresponding verbal theory (Bendor et al. 2001). Several authors have reformulated or extended the initial simulation model, replicating and generalizing the original results (Masuch and LaPotin 1989; Fioretti and Lomi 2008, 2010; Lomi

and Harrison 2012; Troitzsch 2012). While a classic, the garbage can model (GCM) still remains an influential work in the management sciences (Glynn et al. 2020), with substantial spillovers into other disciplines (Simshauser 2018). Another classical group of ABMs with a similar focus is the family of NK models (Levinthal 1997; Rivkin and Siggelkow 2003; Baumann et al. 2018), which applies techniques first developed in evolutionary biology to study the interrelationship between organizational design and market selection forces. ABMs have also been used to support theoretical investigations in fields that range from innovation diffusion (Garcia and Jager 2011; Fibich and Gibori 2010), knowledge transfer (Levine and Prietula 2012), and organizational learning (e.g., Levinthal 1997) to management accounting (e.g., Wall 2016) and organizational design (e.g., Dosi et al. 2003; Clement and Puranam 2018).

A second use of ABMs is as simulation tools for reproducing the behavior of actual complex systems. Examples of this vast literature are works such as Amini et al. (2012) and Stummer et al. (2015), in which ABM is used to simulate product diffusion, Utomo et al. (2018), on modeling agri-food supply chains, Barnes et al. (2010), Ayer et al. (2019), Barnes et al. (2020) on modeling disease transmission. Often, agent-based models are also part of hybrid simulations; we refer to the reviews of Brailsford et al. (2019), Robertson (2019), Currie et al. (2020) for additional details. Further discussion can also be found in Appendix 1.

2.2 Sensitivity analysis

Broadly speaking, sensitivity analysis can be thought of as the exploration of a mathematical or numerical model. The model is typically regarded as a black box that processes a set of inputs and calculates one or more quantities of interest (outputs). Thus, the exploration is not performed by a direct inspection of the model. Instead, the properties of the model are obtained indirectly, by investigating how the output changes given variations in the inputs.

Despite much theoretical progress [see the handbook by Ghanem et al. (2016)], sensitivity analysis is often a neglected task (Saltelli et al. 2020). The recent investigation by Saltelli et al. (2019), who inspect how sensitivity analysis is performed to help scientific modeling across several disciplines, shows that either sensitivity analysis is not applied or it is applied unsystematically. Reviews focused on sensitivity analysis for ABMs reach similar conclusions: in a survey of papers published in *Ecological Modelling* and in the *Journal of Artificial Societies and Social Simulation*, Thiele et al. (2014) find that 88% and 76% of the papers published in the years 2009 and 2010, respectively, do not include a serious sensitivity analysis. Similarly, Utomo et al. (2018) survey agri-food supply chain ABMs, finding that 28% of papers do not incorporate any form of sensitivity analysis, 68% of them perform a basic analysis, and only 4% of them apply a more systematic approach.

We have also performed a closer investigation of the works that have applied some form of sensitivity analysis in agent-based modeling; it is reported in Appendix 1. The analysis shows that the majority of authors use sensitivity analysis to check if the main qualitative conclusions of their work are robust when

some pre-selected parameters are set to values different from those included in their baseline scenario. To assess robustness, modelers typically vary one-parameter-at-a-time, plot the output variable against different parameter values, and show qualitatively whether conclusions are robust. Only a minority of authors consider other goals beyond robustness, vary multiple parameters at the same time, and use more sophisticated quantitative sensitivity methods. However, realistic ABMs would benefit substantially from a systematic method to identify the importance of computationally expensive assumptions. For example, if the outcome of the model was not sensitive to replacing an expensive assumption with a computationally cheaper one, this would be really useful for model building. At the same time, research on ABMs with a theoretical focus could benefit from the systematic approach to sensitivity analysis that we propose, as it would allow a researcher to develop new insights about the importance and interaction of model elements. Yet, authors are not rigorous about mapping goals to methods, and do not vary models elements other than parameters (and usually only a subset of parameters is pre-selected without a formal procedure), possibly due to the lack of a systematic approach. We aim to fill these gaps in the remainder of our work.

In this respect, our work is related to Lorscheid et al. (2012). The authors propose a systematic approach to opening the black-box of simulation models. They are interested in the whole pipeline of model analysis, including the formulation of the research question. Our work focuses on the sensitivity analysis of ABMs. Given this difference in goals, the two papers focus on different aspects. Lorscheid et al. (2012) include a pre-experimental phase to decide the number of simulations that need to be run to achieve stable results, and they also consider an iterative approach to select interesting parameter ranges. This can be seen as a premise to our approach, which addresses the goals that can be answered by sensitivity analysis; we elaborate on the nature of ABMs creating a conceptual map of the elements of an ABM that can be subjected to sensitivity analysis. We also exploit recent results showing that new sensitivity measures can be extracted from the design; we consider new visualization methods that address different sensitivity goals.

3 The elements of agent-based models

In this section, we propose a conceptual structure for ABMs, with the goal of facilitating sensitivity analyses. Our structure classifies the elements of an ABM into a number of sets and subsets. Admittedly, it is not the only conceptual framework that can describe an ABM, and sometimes the distinction between elements might be blurred, subjective or dependent on the specific model under scrutiny. Nonetheless, our aim is to classify the “moving parts” of an ABM that can be subjected to sensitivity analysis. We first describe the conceptual structure at a theoretical level, and then illustrate our framework using the Garbage Can Model.

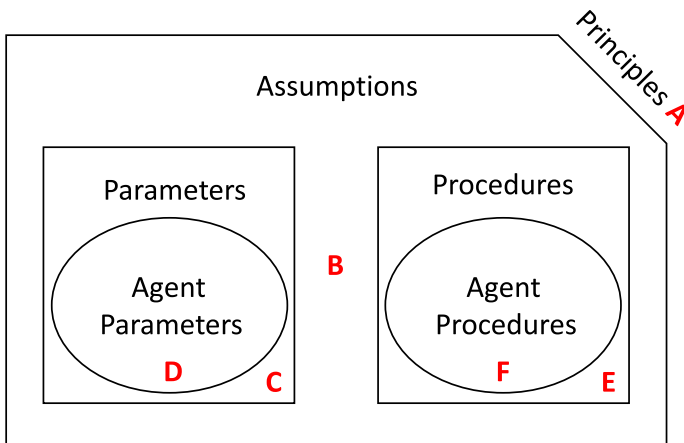


Fig. 1 We identify four types of elements: principles, assumptions, parameters, and procedures. Some parameters are associated with agents, as are some procedures. The capital letters in the diagram identify relevant sets discussed in the text

3.1 Theoretical structure

Our conceptual framework is represented in the diagram of Fig. 1.

All elements of the ABM are contained in the outer rectangle. They can be classified into six sets and subsets:

- *Principles* (A) are high-level elements that define the nature of the ABM. Principles do not include specific algorithmic implementations; rather, they are conceptual guidelines that influence the modeler in formulating specific procedures or in choosing certain parameters. This is illustrated in Fig. 1, in which principles are represented outside of the rectangle containing all practical assumptions of an ABM. We consider principles to be out of the scope of sensitivity analysis, in the sense that varying principles would lead to a different ABM rather than a sensitivity analysis of a given ABM.
- *Assumptions* are low-level elements that define a specific implementation of an ABM. As such, we would like to include as many assumptions as possible within the scope of a sensitivity analysis. While most assumptions are instances of procedures and parameters, some are not. Such assumptions are identified with the letter B in Fig. 1. Telling principles from assumptions is not always straightforward in practice: changing some assumptions may lead to analyzing a new model. We address this issue further in the discussion section.
- *Parameters* are a specific subset of assumptions. They are cardinal quantities that influence the evolution of the model but are determined outside of the simulation run, either at the initialization stage or through a direct intervention of the modeler. Some parameters (C) characterize the environment or define general properties of the simulation (e.g., the number of time steps), while other parameters are closely associated to agents and determine their properties (D).

- *Procedures* are also a subset of assumptions. They can be defined as algorithmic prescriptions that determine the time evolution of the ABM. Unlike parameters, procedures are not cardinal, nor are they drawn from an easily defined set of possibilities. Procedures and parameters are sometimes intertwined, in the sense that certain procedures regulate the distribution of parameters, and certain parameters co-determine the effects of procedures. Certain procedures define the main mechanisms through which the model works or through which agents' attributes are initialized (E), while other procedures define what an agent “does”, and are also commonly known as behavioral rules (F).

Changing the assumptions that characterize agents makes it possible to conduct a sensitivity analysis “at the agent level”, that is to evaluate the impact of agents on simulation outcomes.¹

3.2 Illustrating the structure: the garbage can model

To illustrate the conceptual structure outlined in Sect. 3.1, we consider the Garbage Can Model (GCM). We first provide a summary description of the GCM and then we map out the elements of the model following the conceptual structure of Fig. 1. We focus on the Fioretti and Lomi (2008, 2010) implementation, which takes the most careful steps to match the ABM to the original verbal theory of Cohen et al. (1972). In Appendix 2 we include more details on the software implementation, while we refer to Fioretti and Lomi (2008, 2010) for greater details.

3.3 Description of the GCM

Fioretti and Lomi model four classes of agents: *participants*, who can make decisions; *choice opportunities*, openings for participants to make decisions; *problems* concerning conditions or people inside or outside the organization; and *solutions* in search of problems.² The organization is depicted as cells on a grid, and over time, participants, choice opportunities, problems, and solutions move—randomly and independently—on the grid. When at least one choice opportunity, at least one participant, and at least one solution happen to be collocated in a cell, the participants can make a decision. This is easy if, by chance, no problem is in the cell. The participants just declare a decision, even though it solves no problem! Cohen et al. (1972) refer to this as a *decision by oversight*.

Matters are more challenging when one or more problems are in the cell along with at least one choice opportunity, participant, and solution. Then participants can make a decision only if the solution is good enough to solve the problem(s). In

¹ An intuitive definition of agent is provided by Jennings (2000): An agent is an encapsulated computer system that is situated in some environment and that is capable of flexible, autonomous action in that environment in order to meet its design objectives.

² It is debatable whether choice opportunities, problems and solutions should be considered “agents”. However, because these units match the technical definition of agent.

particular, participants are characterized by a level of *ability*, a cardinal variable that takes a value between two extremes (minimum and maximum ability). Likewise, solutions have a certain *efficiency*, and problems have a certain *difficulty*, both taking values between a minimum and a maximum. If the sum of the abilities of the participants who are present multiplied by the efficiency of the most effective solution in the cell is greater than the sum of the difficulties of problems in the cell, then the participants can make a decision that solves the problems. Cohen et al. (1972) refer to this as a *decision by resolution*. A decision by resolution is the most desirable outcome, as it associates a solution to a problem.

If the participants in a choice opportunity lack a solution good enough to solve their problems (because their *ability* is too low and/or the *efficiency* of the solutions is too low), they are blocked—unable to make a decision or to move on. They remain stuck until another choice opportunity, moving randomly, happens into their cell. The newly-arrived choice opportunity then grabs the most difficult problem and wanders off with it, freeing up the participants and solutions. Cohen et al. (1972) refer to this as a *flight*. In subsequent wandering, the now-freed participants might very well stumble onto the same vexing problem again.

A final twist in the model concerns organizational structure. The model described so far corresponds to an *anarchy* setting. In their ABM, Fioretti and Lomi (2008, 2010) also allow a *hierarchy* setting. In a hierarchy, choice opportunities are always assigned a ranking from most important to least important. Participants can also be ranked relative to one another, as can problems and solutions. The rankings are then used as follows:

- When hierarchy is applied to participants (that is, the participant structure is hierarchical rather than anarchic), each participant is allowed to be part of choice opportunities as important as, or less important than, herself. So the most important participant can take part in all choice opportunities, but less important participants can take part only in less important choice opportunities.
- When hierarchy applies to problems (i.e., the problem structure is hierarchical), each problem can be considered in choice opportunities as important as, or less important than, itself. The most important problems can be considered in all choice opportunities, but less important problems are on the table only in less important choice opportunities.
- When hierarchy applies to solutions (i.e., the solution structure is hierarchical), each solution can be considered in choice opportunities as important as, or less important than, itself. The most important solutions can enter all choice opportunities, but less important solutions are up for consideration only in less important choice opportunities.

The model allows researchers the flexibility to apply hierarchy rather than anarchy just to participants, just to problems, just to solutions, or to any combination of the three.³

³ Fioretti and Lomi (2008, 2010) used the terms decision structure, access structure, and availability structure to indicate whether participants, problems, and solutions—respectively—had anarchic or hier-

3.4 Elements of the GCM

To illustrate our conceptual structure of Fig. 1, in this section we map various elements of the GCM.

- Principles (A): A key principle of the GCM is organized anarchy: decision makers work on problems they have stumbled upon, using solutions that happen to be available. This can be contrasted with game-theoretic models of industrial organization, in which agents with clear preferences seek rational solutions to well-defined problems. Another distinguishing principle of the GCM is independence between the objects of decision-making. Indeed, choice opportunities, participants (decision-makers), solutions and problems exist independently of one another. This principle distinguishes the GCM from other models where, instead, solutions exist only attached to problems.
- Assumptions that are neither parameters nor procedures (B): The assumption that agents move on a regular square grid as opposed to a more sophisticated structure such as a network fits into this class.
- Parameters that are not agent parameters (C): A first example of an element in this class is the number of choice opportunities, as it determines how many agents of this type exist, but not their properties. Note that the number of choice opportunities is an assumption rather than a principle: in the original implementation of the GCM (Cohen et al. 1972) there was only one choice opportunity, but some scholars have since argued that multiple choice opportunities are more in line with the verbal theory of the GCM (Bendor et al. (2001; Fioretti and Lomi (2010)). Thus, the number of choice opportunities does not seem to lead to conceptually different models. Another example of an element in class C is the size of the grid on which agents wander.
- Agent parameters (D): The minimum and maximum levels of ability, efficiency and difficulty determine key attributes of the agents, so we consider them as “agent parameters”.
- Procedures that are not agent procedures (E): As an example of a procedure in this class, we consider the rule for assigning values to ability, efficiency and difficulty. This procedure assigns these values to agents by sampling the interval delimited by the minimum and maximum values above uniformly at random. While it determines an agent’s parameter, it does not define what an agent “does”, and so we do not think of it as an agent procedure.
- Agent procedures—or behavioral rules (F): The participant, problem and solution structures determine the access of participants, problems and solutions to choice opportunities, and so clearly influence how agents behave. Thus, we consider these procedures as behavioral rules.

Footnote 3 (continued)

archical access to choice opportunities. We prefer, and use, more intuitive terms: participant structure, problem structure, and solution structure.




Fig. 2 Six steps for sensitivity analysis of agent-based models

4 A systematic process for sensitivity analysis: six steps

In this section, we outline a process in six steps to make sensitivity analysis of ABMs systematic. We deliberately keep the discussion as general as possible, leaving all mathematical details and most references to Appendices 2 and 3. Also, for illustration purposes, we shall focus on traditional parametric elements in this section, leaving the extension to non-parametric elements to Sect. 5. We visualize the steps in Fig. 2.

4.1 Output of interest

The first step of the process is choosing the output(s) of interest. Indeed, as noted in Lee et al. (2015), usually ABMs produce a multiplicity of outputs. One or more of the outputs could be quantities of interest, as long as they are deemed relevant by the analyst or by the decision maker. Moreover, the ABM response is frequently stochastic. For instance, in the Garbage Can ABM, the number of decisions by resolution or by oversight is stochastic with respect to the simulation inputs. Formally, in stochastic simulators the output Y is the conditional distribution of the quantity of interest given inputs X . Not infrequently, analysts are interested in one or more summary statistics or functions of such distribution. For example, quantities of interest can be a moment (the expected value, the variance of Y), a quantile (the median, the 95 or 99 quantile) or the probability that Y is above a given threshold. All these quantities can be the target of a sensitivity analysis.

4.2 Goal

After the analyst has determined the quantities of interest, she decides on the *goal* of the sensitivity analysis (this is also known as *setting* in the sensitivity analysis literature). A variety of goals are possible. The analyst may be wishing to increase her own understanding of the model behavior, or might be asked to test theoretical aspects as part of a broader research investigation, or may be required to deliver robust managerial insights to a stakeholder. Broadly, one may be interested in the relative importance of alternative model elements (*factor prioritization*), or in whether they increase or decrease the quantity of interest (*direction of change*), or in how they interact (*interaction quantification*), or in whether conclusions drawn from the model are robust with respect to variations in the inputs (*robustness analysis*).

4.3 Elements

A third step is to decide which elements of the ABM to vary. This phase requires a critical review of the model, its main principles, assumptions, parameters, and procedures. For instance, varying a given element may not make managerial sense in a given application, and in this case it would be safe to exclude that element from the sensitivity analysis. The structure laid out in Sect. 3 is a useful reference to the researcher for classifying the element(s) at hand. It helps in balancing which elements are left out of the analysis (e.g., it is hazardous to focus entirely on parameters and to ignore procedures—see Sect. 7 for further discussion).

4.4 Sensitivity method/design

A fourth step is to choose the most appropriate method for each goal-element combination, and the related experimental design, that is, the choice of how to sample points in the input space. The choice of the method also implies the choice of the scale, local or global. Let us start from a local sensitivity analysis, namely a sensitivity analysis around a particular point in the input space (also known as *scenario*). A natural choice for Agent-Based Models is to use *finite difference* methods. Finite differences are given by the difference between the output at a base scenario and the output at alternative scenarios, obtained by varying one or more of the inputs. The relevant sensitivity measures are *main, interaction and total effects* (mathematically defined in Appendix 2). They capture, respectively, the individual influence of each input, the influence of an input when varied together with other inputs net of its individual effect, and the sum of the two. Total effects can answer a factor prioritization goal, while main and interaction effects can be used to address direction of change and interaction quantification. Also, the examination of model output across scenarios addresses robustness.

The evaluation of the model at several locations in the input space is the basis for a global sensitivity analysis. Differently from a local analysis, a global sensitivity analysis implies considering the response of the model with inputs sampled across the entire input space, or at least across the portion of the input space that the analyst deems most relevant Saltelli et al. (2008)—see Appendix 2 for mathematical aspects. This exploration is often performed through various forms of Monte Carlo experimental designs. The selection of the global sensitivity method depends on the goal. For factor prioritization, *variance methods* are a classic set of tools. According to these methods, the most important inputs are those that explain most output variance. *Moment-independent methods* are a more recent approach that exploits the full distribution of the output given the inputs. For direction of change, popular techniques are, among others, gradients computed at randomized locations and partial dependence functions. For interaction quantification, the analyst can use high-order variance-based methods: instead of looking at which individual inputs explain variance the most, these methods consider the contribution to variance of two or more inputs that are varied together. Finally, the examination of the output behavior across

the Monte Carlo-generated scenarios helps the analyst in conducting a robustness analysis.

4.5 Assignment of values

Once the goal and the method/design of the analysis have been identified, the researcher selects numerical values for the parameters. The type of assignment depends on the design. For instance, if the analyst is considering a local analysis, she needs to assign reference values that form a basic scenario (usually called base case) and variation ranges/levels to the inputs to reach other scenarios (sometimes called best case or worst case). Conversely, if the analysis is global and the design is probabilistic, the analyst needs to specify values in the form of ranges (supports) and distributions for the inputs. Value assignment is a delicate step in ABMs: One may have constraints on parameter changes dictated by procedures or other model elements. That is, a choice for a procedure may limit the range of “reasonable” values that a parameter can take—Sect. 7 provides further discussion.

4.6 Results communication/visualization

The last step is the choice of how to visualize the results of a sensitivity analysis. For any combination of goal/element/method, there may be one or more appropriate visualization tools. For example, for factor prioritization and interaction quantification, parametric elements and calculation of main, total and interaction effects (local analysis), bar-charts or *tornado diagrams* Eschenbach (1992) are suitable. If instead one is interested in direction of change, a variety of tools are available, ranging from spiderplots Eschenbach (1992) to ICE plots Goldstein et al. (2015). In this respect, we note that ICE plots have been defined for simulators whose output is deterministic. Most ABMs have a stochastic response. We then propose a modification of ICE plots that takes the stochastic nature of the response into account. Specifically, because the response of a (stochastic) ABM simulator is, in fact, a distribution, we introduce a test for the statistical significance of the difference in mean values. The test allows us to distinguish whether a small but non-null change in a mean (or conditional mean) is indeed significant or just the consequence of numerical noise.

5 An experimental design for including non-parametric elements

Local and global methods have been mainly conceived for inputs which are continuous or discrete real numbers. However, ABMs include also elements such as behavioral rules and interaction procedures, that do not belong to a well-defined numerical space.

While it is not always necessary to vary all the elements of an ABM simultaneously, it is important to stress that a researcher who is performing a sensitivity analysis involving only parameters is implicitly fixing a substantial portion of an agent-based simulation. Suppose one performs a global sensitivity analysis

on all of the model parameters. As Fig. 1 highlights, only a fraction of the ABM would be tested (subsets C and D). Using the example of the GCM, the effect of the participant, solution and problem structures on simulation outcomes would be ignored.

To address these issues, we treat the alternative specifications of a non-parametric element as the levels of a categorical variable (see Appendix 3 for all technical details). The levels may or may not have an ordinal meaning. We then propose to perform a full factorial design that comprises all possible combinations of (a) levels of the categorical variables associated with non-parametric elements and (b) two or more discrete parameter values for each parameter (the choice of a full factorial design is not restrictive; computationally cheaper designs can be also exploited—see Sect. 7).

This design makes it possible to address sensitivity analysis goals in a way that is in between a local and a global approach. In particular, it allows the researcher to compute finite difference measures (main, interaction and total effects) at several locations of the input space and to obtain global measures as appropriate functions of these local effects.

For example, as we show in Appendix 3, from the variance of the main effects it is possible to estimate the total order variance-based sensitivity indices, a recommended global importance measure (Saltelli et al. 2008). This allows the researcher to address the factor prioritization goal. The same sample allows also the calculation of moment-independent global sensitivity measures. We should emphasize that obtaining global measures as appropriate functions of local measures is by no means novel in sensitivity analysis (Morris 1991). Nonetheless, our approach extends this idea to non-parametric elements, making it possible for example to assess the contribution of a behavioral rule to output variance.

Our design also allows the researcher to address additional goals besides factor prioritization. Regarding direction of change, for a given input, the design allows one to consider all pairs of points that are only different because of that input. One can then compute the resulting Newton quotients at all such pairs of points. We propose to visualize this information in Stochastic Individual Conditional Expectation (S-ICE) plots that take into account the stochastic nature of the ABM response. We will describe these plots when presenting our example of the GCM in Sect. 6 (Appendix 3 reports additional details).

Our design also allows the researcher to quantify interactions of a parameter and a procedure. To do so, the researcher evaluates the simulator response at all pairs of points that differ in two or more inputs. It is then possible to calculate all corresponding interaction effects as well as second-order Newton quotients. We plot them as histograms to facilitate visualization. Such visualization shows heterogeneity of interactions across scenarios, enriching the insights coming from reporting a single number that represents a global interaction between inputs.

Finally, regarding robustness, we show how S-ICE plots can be used to test the robustness of conjectures on the model (we refer the reader to the example of the GCM in the next section).

6 The garbage can model as an illustration

This section offers an illustration of our approach through its application to the GCM.

6.1 Applying the six steps

6.1.1 Output of interest

Recall from Sect. 3.3 that decisions in the GCM can happen by oversight, flight or resolution. A notable outcome of many simulations is that decisions by resolution are far from ubiquitous. That is, decisions as envisioned in more traditional models of organizations, in which participants gather to apply a solution directly to a problem, are not at all the only thing that happens. Quite often, simulated participants make empty, symbolic decisions that solve no problems (oversight) or get bogged down in intractable issues (flight). The portion of decisions by resolution is therefore a natural quantity of interest in GCM simulations and a focus of previous literature (Takahashi 1997; Fioretti and Lomi 2008, 2010). In this section, our sensitivity analysis focuses solely on the fraction of decisions by resolution for reasons of space; the same analysis can be conducted for other outputs. To deal with stochasticity, for each set of inputs we perform 100 simulations with different random seeds, and report the average fraction of decisions by resolution across these simulations. Each simulation lasts 1000 time steps; this is enough for results to be sufficiently stable, in the sense that we observe that the output of interest shows little variability across simulation runs. Concerning this issue, the recent work of Vandin et al. (2021) proposes a way to automatically select the number and duration of simulation runs, provided some requirements.

6.1.2 Goal

We consider obtaining insights regarding all four main goals: factor prioritization, direction of change, interaction quantification and robustness analysis.

6.1.3 Elements

We select some elements of the GCM to illustrate the range of types of elements shown in Fig. 1. Of the elements that we consider (see Table 1), the number of opportunities and the minima and maxima of ability, efficiency and difficulty are parametric elements (D and F in Fig. 1, respectively), while participant, solution and problem structures (E) are procedures and thus non-parametric elements.

In contrast, we decide *not* to vary other elements of the GCM: these are the principles (A), which by definition are not candidates for sensitivity analysis; the size of the grid (D), which comprises 195 cells (15 in one direction and 13 in the other

Table 1 Inputs for the GCM sensitivity analysis

Element (procedure)	Name	Values	Element (parameter)	Name	Values
Participant structure	X_1	A, H	Number of choice opportunities	X_4	25, 40
Solution structure	X_2	A, H	Minimum ability	X_5	0, 1, 2
Problem structure	X_3	A, H	Maximum ability	X_6	8, 9, 10
			Minimum efficiency	X_7	0.0, 0.1, 0.2
			Maximum efficiency	X_8	0.8, 0.9, 1.0
			Minimum difficulty	X_9	0, 1, 2
			Maximum difficulty	X_{10}	8, 9, 10

A anarchic, H hierarchic

direction), and the number of participants, solutions and problems (D), which are 25 each; and the assignment rule of ability/difficulty/efficiency (C). We also do not vary the structure of the grid—a regular square lattice (B).

Looking back at Fig. 1, we notice that we vary elements in all main sets of our classification, suggesting that our exploration of the model touches all main classes of elements. Our conceptual structure also suggests which parts of the model are left untested. Since we do not subject the assignment rule of ability/difficulty/efficiency to our sensitivity analysis, we are not varying any non-agent procedure (C in Fig. 1).

6.1.4 Sensitivity method/design

We use two variance-based indicators (first order and total order indexes) and two moment-independent methods, namely the δ and β^{Ku} importance measures Baucells and Borgonovo (2013) (see Appendix 2 for the definitions). Using alternative measures has the advantage that they look at alternative properties of the output distribution. If alternative measures concur in indicating an input as important, this indication is more robust than the one coming from a single sensitivity measure.

6.1.5 Assignment of values

We assign participant structure, problem structure, and solution structure (which are non-parametric elements) to categorical variables X_1 , X_2 , and X_3 . These can take two values: $X_i = \{A, H\}$, with $X_i = A$ representing anarchy and $X_i = H$ representing hierarchy ($i = 1, 2, 3$). Only considering these two options is consistent with the organizational structures of the original Garbage Can Model and with the findings of Fioretti and Lomi (2008), who identified Anarchy and Hierarchy as the most relevant structures. The remaining elements that we test are parametric. According to our design, we select a few discrete values that the parameters can take. In particular, the number of choice opportunities can take a base value of 25 (as for all other classes of agents), and an incremented value of 40, which is sufficiently different to the baseline to generate a substantial change in the output. We assign three values to each of the remaining parametric variables: the two extremes and the central value




Fig. 3 Ranking of the GCM inputs according to first- and total-order variance-based, δ_i and β_i^{Ku} moment-independent measures. Sorting is based on the total-order indexes τ_i^2

of their respective ranges—see Table 1. Ranges are selected to avoid overlap of the minimum and the maximum values of attributes.

We consider all possible combinations of elements, giving rise to a full factorial design with 11, 664 input configurations (scenarios). Evaluating the ABM at these scenarios, allows one to compute 5,832 main effects for each element taking two values, $X_1 : X_4$, and 11,664 for each element taking three values, $X_5 : X_{10}$. Also, it is possible to calculate a total of 664, 848 pairwise interaction effects.

6.1.6 Results communication/visualization

For result visualization we use the tools discussed in Sect. 5.

6.2 Results of the numerical experiments

Across the 11,664 scenarios, we obtain an average of 10.3% decisions by resolution, 65.4% by oversight and 24.3% by flight, corroborating the original insights of the GCM. From now on, we focus on the share of decisions by resolution.

6.2.1 Factor prioritization

Figure 3 displays a Pareto-chart visualizing the ranking of the elements. Consistently across the four sensitivity measures, problem structure is the most important element. The number of choice opportunities is the second most important element in affecting the share of decisions by resolution, while solution and participant structure are ranked third and have a similar importance.

Overall, assumptions regarding non-parametric elements impact the output of the GCM more than assumptions regarding parameters. One may argue that this result is driven by the interval we chose for parameters — that is, between 0 and 2 (or 0.0 and 0.2) for minimum values, X_5, X_7, X_9 , and between 8 and 10 (or 0.8 and 1.0) for maximum values, X_6, X_7, X_{10} . The argument would be that choosing relatively




Fig. 4 S-ICE plots for the share of decisions by resolution as a function of the elements. Red: negative change. Blue: positive change. Grey: change is not statistically significant. (Color figure online)

narrow intervals in which parameters are varied could artificially result in procedures looking more important than parameters. To test the robustness of our result, we assigned alternative ranges to minimum and maximum values, $X_5 : X_{10}$. In particular, we repeated the calculations letting X_5 and X_9 take values between 0 and 4 and X_6 and X_{10} take values between 6 and 10 (we also let X_7 and X_8 take values between 0.0 and 0.4, and 0.6 and 1.0, respectively). These intervals are almost as wide as they can be relative to each other, as it would not make sense to have a minimum value of, say, ability, be larger than its maximum value. Re-running the analysis, the findings remained unchanged.

6.2.2 Direction of change

Figure 4 displays the S-ICE plots for the inputs. These graphs can be read as follows. Consider the top left panel as a reference. The horizontal axis reports participant structure ($X_1 = A$ is anarchy, while $X_1 = H$ is hierarchy). The vertical axis reports the expected share of decisions by resolution conditional on the value of X_1 in all scenarios (Because there are several scenarios, it is difficult to distinguish small black dots on the vertical axis, which instead appear as a continuous vertical black line). The large black dots on each vertical line represent the average percentage of decisions by resolution across all scenarios. For instance, given a participant structure that is anarchic (A), on average about 10% of the decisions are taken by resolution. This fraction increases on average to about 11% when participant structure equals H —technically, these two dots form the graph of the corresponding partial dependence function, which, in our design case, is discrete. This information is complemented by the lines joining the smaller black dots at $X_1 = A, H$, whose purpose is to show the change in the portion of decisions by resolution, scenario by scenario. These lines join pairs of small black dots such that an input in a given scenario is switched to an alternative value in the other scenario, with all other

elements unchanged. More specifically, in the first panel, the lines join pairs of conditional expectations obtained when participant structure is A in a given scenario and switched to H in the other scenario, with all other elements fixed. As we mentioned, for each scenario we run 100 simulations to estimate the conditional mean of the share of decisions by resolution in that scenario. Hence, the value we are dealing with is a sample mean; to assess whether the observed change is different from zero, we need to perform a statistical test. We use a two-sample t-test at 5% significance level. If the null hypothesis is not rejected the corresponding line in Fig. 4 is gray; otherwise it is blue or red depending on whether the change is positive or negative. In the first panel, we observe a majority of blue lines, indicating that the switch of participant structure from anarchy to hierarchy increases the portion of decisions by resolution in the majority of scenarios. At the same time, not all lines are blue, i.e., there are scenarios in which the opposite occurs and thus the output is not monotonic in this variable.

We now comment on the other panels. The solutions structure panel (first row, second panel in Fig. 4) delivers insights similar to the participants structure panel. The problem structure panel (first row, third panel) yields instead a different message. When problem structure is A (anarchy), we expect about 11.5% decisions by resolution; this number decreases to about 9% when problem structure is hierarchical. Thus, the switch of problem structure from A to H leads to a decrease (significant percentage-wise) in the average portion of decisions taken by resolution. Note that virtually all lines are red: for almost all combinations of all other inputs, the switch of problem structure from anarchy to hierarchy leads to a decrease in the portion of decisions by resolution. (See the paragraph on robustness below for additional comments on this property.) Moving to the following panel, it is possible to see that the number of choice opportunities, on average, has a positive effect on decisions by resolution; however, the plot shows two well clustered sets of positive and negative realizations. This is a sign that this element is involved in interactions with other inputs. In fact, it can be shown (see Appendix 3) that if an input is binary and the model is additive in that input, then there is only one possible slope for the one-way lines in an S-ICE plot. Thus, changes in slope denote non-additivity. Finally, panels for minimum and maximum values of ability, efficiency and difficulty ($X_5 : X_{10}$) exhibit a weaker but regular impact, with ability and efficiency parameters having on average a positive effect on decisions by resolution, and difficulty parameters having on average a negative effect. Overall, the panels in Fig. 4 signal a non-monotonic behavior of the quantity of interest as a function of the inputs. Traditional local sensitivity analysis methods that vary one input at a time would thus be inadequate to study this output.

6.2.3 Interaction quantification

We then come to our third goal, understanding interaction effects quantitatively. Overall, there are 45 possible pairwise interactions among the 10 variables we focused on. The first question we address is whether pairwise interactions are significant in determining the model outcome. To answer this question, we fit a linear regression model with all 45 interaction terms on the input-output data (we use the




Fig. 5 The five strongest interactions for decisions by resolution

subroutine `fitlm.m` available in Matlab). Results show that 43 pairwise interactions are statistically significant at the 1% threshold for decisions by resolution. Thus, interactions do matter in the GCM. Another indication in this sense comes from the sensitivity measures in Fig. 3. A visual inspection of the values of the first- and total-order variance-based sensitivity indexes shows several cases with a large discrepancy between these measures. For example, the values of the first- and total-order indexes for number of opportunities and participant structure in the Pareto chart of Fig. 3 reveals low first-order indexes and large total-order indexes, suggesting that these inputs owe a significant part of their importance to their interactions with the remaining inputs.

We gain information about the signs of interactions through second-order finite difference interaction effects. In Fig. 5, we report the histograms of the normalized second-order Newton quotients for the five most important pairwise interactions. For each interaction, in the legend we also display the mean value, the standard deviation, and the percentage of the scenarios that lead to a positive interaction. Note that interactions are considered positive when increasing both inputs together leads to an increase in output, by more than what could be expected by raising the two inputs individually. This definition only applies to elements that have a notion of ordinality; when considering categorical variables, this interpretation does not apply. To study interactions between numerical and categorical elements, by convention we consider $\{A, H\}$ as an ordered set in which A is lower. Therefore, positive interactions between a parameter and a categorical variable mean that output increases when the parameter is increased together with one out of problem/solution/participant structures switching from anarchy to hierarchy across two scenarios.

We find that the number of opportunities is involved in three of the five most important interactions. Specifically, the strongest interaction is between number of opportunities and problem structure. The interaction is positive in 98.8% of the cases and has an average effect of +2.8% on the share of decision by resolution, a large impact when considering that the average number of decision by resolutions is 10.4%. The second most important interaction is the one between participant structure and problem structure. This interaction is mostly positive (96% of the scenarios). Note that, had the analyst focused only on parameters, the sensitivity analysis would not have revealed the most important interactions.

6.2.4 Robustness

Let us suppose for a moment that the experiments on the GCM were carried out with the goal of validating the following two conjectures:

1. C1 *A hierarchical problem structure decreases efficiency in decision making.*
2. C2 *A higher number of opportunities increases efficiency in decision-making.*

Here, what we mean by an increase in efficiency is an increase in the fraction of decisions taken by resolution.

We can check if these conjectures are true by considering each possible line in the S-ICE plots. Conjecture C1 is confirmed by the experiments of the ABM. Indeed, in the third panel of the top row of Fig. 4, all lines are red and none are blue. There would a few cases in which the change in problem structure from anarchy to hierarchy slightly increases the expected number of decision by resolutions. However, it turns out that none of these increases is statistically significant, and indeed these lines are colored grey in our representation. These cases are a byproduct of the stochasticity of the model and we expect them to drop to zero for a large enough number of replications. In sum, we cannot find any statistically significant case in which a hierarchical problem structure decreases the share of decisions by resolution. Consequently, there is not enough statistical evidence to reject the first conjecture.

Conversely, conjecture C2 is not robust. Indeed, the cluster of red lines in the fourth panel in Fig. 4 shows that the opposite of the conjecture occurs in 26.5% of the scenarios, and many of these cases are statistically significant. In particular, we find that this behavior is critically dependent on the interaction with problem structure. When problem structure is hierarchic, the number of opportunities has a consistently positive effect (in 100% of the scenarios); yet, the effect is mixed when problem structure is anarchic, with a positive effect in 47% of the cases and a negative effect in the remaining.

6.3 Some managerial insights

Managerial insights are obtained when a question of interest for an operations research/management science investigator is answered by interrogating the ABM at hand. To illustrate, consider a researcher who is using the GCM to understand the combination of elements necessary to achieve a higher share of decisions by resolution (this number might be considered a proxy of the effectiveness of an organization). As is observable in the problem structure panel of Fig. 4, while the impact of a hierarchic problem structure on decisions by resolution is almost always negative, sometimes this effect is milder. This is because of an interaction effect: the negative impact can be mitigated by an increased number of opportunities and anarchic participant and solution structures (Fig. 5). Thus, the combination that leads to the highest share of decisions by resolution is the one with hierarchic participants and solutions structures, anarchic decisions structure, and a low number of choice

opportunities. Note that, had the researcher performed a simple series of one-at-a-time sensitivities (equivalent to considering only individual effects), the researcher would have naively picked the structure in which problems and solutions freely float in the organization as optimal.

More generally, interaction analysis reveals a further insight: the strongest interaction effect is the one between number of opportunities and problem structure, which can be seen as complementary variables given the sign of their interaction. Suppose that an organization has a hierarchic problem structure and a large number of choice opportunities. Because the interaction between these two variables is positive, the decrease in the decision-by-resolution share generated by the hierarchic problem structure is mitigated by the simultaneous variation of the number of opportunities.

A further insight for a research using the GCM is understanding what determines the share of decisions by resolutions. All in all, our analysis shows that since the presence of a problem is the defining difference between a decision by resolution and a decision by oversight, it is problem hierarchy that leads to the greatest decline in decisions by resolution. An organization that wants to avoid a proliferation of problem-free decisions should take care to spread problems out widely and avoid having lots of problem-poor, solution-rich decision makers looking for work.

To learn a broader lesson from these results, our sensitivity analysis approach can suggest where further modeling efforts could be fruitful. In our example, given the importance of problem structure, one may think of building sub-models that simulate access of problems to choice opportunities at a finer granularity than the anarchy-hierarchy procedures in the current GCM implementation.

7 Discussion

We now discuss in depth a number of possible issues that we mentioned in the previous sections.

7.1 The case of computationally heavy models

Thanks to the fast execution time of the GCM, we were able to explore all input combinations, but our approach may help the analyst also when this is not possible. On the one hand, consider an analyst wishing to apply our design to a time consuming model. Given a budget of available simulator runs, she can reduce the number of model evaluations in several ways, for instance by considering groups of inputs rather than individual assumptions (Saltelli et al. 2008); as an alternative, she might apply a so-called fractional factorial design, which allows the exploration of the model input space at a lower number of locations than in a full factorial design. In using fractional factorial designs, the recommended procedure would be that the investigator pre-identifies potential interactions of interest (for instance, all second order interactions), and then chooses the design that allows the desired level of resolution. Here, the analyst has available a variety of choices based on orthogonal

arrays (Morris et al. 2008) or supersaturated designs (Lin 2000). Moreover, (or alternatively) the analyst may fit the model response with an emulator, and focus on parametric sensitivity in a predetermined region of the model input space. If, according to statistical performance measures, the emulator fit is accurate, then the original (and time consuming) mathematical model can be replaced by the fast emulator, eliminating the computational burden restrictions (see Appendix 1 on works applying emulation in agent-based modeling).

On the other hand, the methodological part of our approach can provide guidance to analysts dealing with models for which computational burden does not allow one to apply the proposed design. Given high resource constraints, it becomes even more delicate to carefully examine the triplets element/goal/method to find the combination(s) that could maximize the insights on model behavior and aid managerial intuition. The analyst could apply our sensitivity analysis iteratively, in conjunction with model building. This would make it possible to avoid computationally heavy assumptions, if cheaper assumptions lead to similar results with respect to the goals of interest.

7.2 When is it sensitivity analysis and when is it a new model?

In Sect. 3.1, we defined principles as those elements of an ABM that characterize it: changing principles would lead to a new model. As it could be difficult to tell principles from assumptions, especially when considering alternative procedures, this might lead the analyst to the border between exploring a new model versus performing a sensitivity analysis of the same model. If the use of different procedures corresponds to the use of different theories, then, indeed, the analyst might be interpreting results that are in between a sensitivity analysis of a model and the comparison between two alternative models. To be concrete, consider the choice in the GCM between anarchy and hierarchy for, say, participant structure. Both anarchic and hierarchic participant structures are consistent with the principles of the GCM listed in Sect. 3.4. However, a procedure that computes the optimal assignment of participants in a way that the number of decisions by resolution is maximized would be in contrast with the first principle from Sect. 3.4, and thus lead to a new model. Our investigation would suggest that it is not possible to draw a sharp line for all ABM elements and all ABMs. Telling whether an assumption is in conflict with a principle depends on domain knowledge, and is outside the scope of sensitivity analysis.

7.3 Selection of elements

Choosing which elements are subjected to sensitivity analysis and which are ignored depends on many aspects. First, the analyst could disambiguate between elements that are central to the research question, which we call elements *of interest*, and elements that are needed for internal consistency of the model, which we call *incidental*. Under this distinction, the analyst could focus the sensitivity analysis on elements of interest. For instance, in the GCM, anarchy vs. hierarchy is certainly an

element of interest, while the assignment of ability/efficiency/difficulty (our example element in class E in Sect. 3.4) is an incidental element, which we do not vary.

Other classifications of elements are possible. In terms of parameters, Smith and Rand (2018) suggest four possible criteria to select the ones appropriate for a robustness test. Their analysis suggests that model-altering parameters and parameters about whose value the research is uncertain should be part of a sensitivity analysis. Additionally, it can be possible to add controlling parameters, which have policy-intervention potential, and environmental parameters. The researcher decides which subset of these parameter groups should be included in the analysis, depending on what the research question is.

Finally, some combinations of elements may yield identical results simply through the mechanics of the model. For example, in the GCM, for values of grid size and number of agents around the current assignment, it turns out that it is equivalent to increase grid size or decrease the number of agents of each type, as what determines the number of potential decisions is the “density” of agents.—We actually performed additional experiments increasing the number of agents from 25 to 40 and keeping the density fixed, and results confirmed this assumption.—However, testing this effect for a larger number of agents might result in different insights: for instance the model might evidence the appearance of self-organizing behavior. More generally, understanding these equivalences requires a detailed knowledge of the simulator, which may not be possible for large and complicated ABMs.

7.4 Assigning values to input variables

To carry out a sensitivity analysis, the analyst is actually asked to make decisions on what element(s) to vary, and on what values to assign to the element(s) under scrutiny. This is a critical step, because the indications of the model as well as the results of the sensitivity analysis depend on the quantitative assumptions. However, this is a common problem for any sensitivity analysis method. We can distinguish alternative situations. In a modeling phase in which information collection on the inputs is partial, the analyst may be willing to assign wide or conservative (albeit subjective) ranges to the inputs, to gain preliminary insights about, among others, the correctness of the model behavior. In a second phase, the researcher might hone in on parameter ranges in which the model changes its behavior, e.g., when a previously found result switches its sign. For complex ABMs that attempt to replicate real world phenomena, relevant parameter ranges might be suggested by field experts. Varying procedures might further lead to the observation that the domain of procedures is actually infinite and that while an infinite domain for a parameter is tractable, such a domain is intractable for procedures. In this respect, our approach can even help the analyst to recognize that she is not in a position to make decisions concerning numerical assignments. Then, the quantitative part of our approach would not be applicable. It would remain an open research question whether such an impasse would be pointing the analyst towards using an alternative sensitivity analysis approach or towards additional modeling efforts or information collection, which are needed before a sensitivity analysis can be fully informative.

Also, the assignment of values cannot be separated from the structure of the model: constraints may impose to bind together certain elements, limit the magnitude of relative changes or even their possibility of varying individually. In certain situations one might not be able to disentangle individual effects by varying each input separately. These structural constraints then impact the design that can be chosen for a meaningful sensitivity analysis.

Our approach also opens further research questions. First, while sensitivity analyses on parameters mostly lead to further data collection, sensitivity analyses on procedures and behavioral rules can guide further research on how agents interact or behave. In this respect, our method could increase the synergy between lab experiments and ABMs, highlighted in the recent work of Smith and Rand (2018). Second, the issue about the border of an ABM, i.e., whether changing an element leads to a new model, actually raises the question of what is the border of sensitivity analyses.

8 Conclusion

This work contributes to the use of ABMs by studying their sensitivity analysis. We have proposed a general conceptual structure that classifies the main elements of an ABM and a six-step approach that allows the researcher to introduce an element-method oriented analysis. We have studied a design that enables the variation of potentially all ABM elements simultaneously. Our approach borrows from the literature on the design of computer experiments and adds the calculation of a variety of sensitivity measures; in particular, the design allows the randomized evaluation of finite change sensitivity indices for determining individual and interaction effects. For direction of change, we have introduced a convenient graphical representation modifying the well known ICE plots to account for the stochastic nature of the ABM response.

To illustrate our method, we have carried out a thorough sensitivity analysis of the GCM, varying three procedures and seven parameters. Among other results, the analysis reveal that the most important element is non-parametric and interaction occurs between a parameter and a non-parametric assumption; this interaction would have been overlooked if the sensitivity had focused only on parameters. Limitations of the approach and direction of future research have been underlined in the discussion section.

Appendix

1. Sensitivity analysis in agent based simulation: a concise overview

Agent-based modeling is a versatile simulation technique vastly applied in the social, physical and natural sciences. Its popularity in the management sciences is increasing with time. A search in the database Scopus,⁴ reveals 319 papers

⁴ Search updated on January 2021, using [SUBJAREA (deci) OR SUBJAREA (busi)] AND TITLE-ABS-KEY (“Agent-based modeling” OR “Agent-based simulation”).

concerning agent-based modeling in the field of decision and management science during the period from 2000 to 2009. In the following decade (2010–2019), the same searching criteria reveal 1261 articles, with an increase of +295%. However, sensitivity analysis remains a challenging step that is often missing or performed in a sparse and unsystematic way. In this section, we review the literature on sensitivity analysis of ABMs according to three criteria: the goal (also called setting), the technique or method adopted, and the element under investigation.

Regarding goals or settings, the most frequently adopted setting in the literature is robustness. In a robustness analysis the modeler is interested in gaining an understanding of how stable a result is, given plausible variations of the model inputs/assumptions. In practice, this is done by checking whether parameter values different from those in a baseline scenario impair the main qualitative/theoretical conclusions of the paper. For instance, Smith and Rand (2018) study the dynamics of a network of workers that either narrow or widen their social network upon being fired, and test whether this choice has an effect on wealth inequality. They show that it does, and that this finding is robust to different unemployment rates and initial network conditions. Similarly, in the context of a network of innovation diffusion, Stummer et al. (2015) study the adoption dynamics of a second-generation biofuel in Austria. The authors perform a robustness analysis on the size of the population and on several network parameters, showing that the path of the biofuel market share is similar for all distinct parameter values (the only exception is the case of very sparsely connected networks). Other examples of this use of sensitivity analysis can be found in Lenox et al. (2006), Puranam and Swamy (2016), Jaspersen and Peter (2017), Keuschnigg and Ganser (2017), Clement and Puranam (2018), and Hernandez and Menon (2018).

While only a minority of authors consider other goals beyond robustness, another relatively popular goal of sensitivity analysis is input prioritization (Dancik et al. 2010; Jiang et al. 2016; Broeke et al. 2016). For example Delre et al. (2016) study the effects of shared consumption on the life cycle of new product launches and, using sensitivity analysis, they show that shared consumption has the largest impact on the output. A similar but opposite goal for sensitivity analysis is the identification of least important parameters, which can then be fixed to save computational cost and reduce output complexity (Ligmann-Zielinska 2018). Another setting that can be found in the literature is direction of change. For instance, Levinthal (1997) shows that a decrease in the smoothness of the fitness landscape negatively affects the ability of organizations to survive in changing environments. Other examples of studies that frame the sensitivity analysis around this setting are Lenox et al. (2006) and Sauvageau and Frayret (2015). An important but often overlooked sensitivity setting is interaction quantification. While this analysis is carried out in few works, an exception is represented by the study of Dancik et al. (2010), where the interaction effects between parameters are thoroughly explored and the most important interactions are reported. Other studies instead rely on the exploration of the behavior of the model by plotting a two-dimensional surface obtained varying two arbitrarily chosen inputs, while keeping all others fixed (e.g., Troitzsch 2012). Usually researchers are interested only in two-element interactions, since higher-level interactions are

computationally expensive to evaluate. We note that authors adopt sensitivity analysis to answer one of the mentioned settings, without a-priori stating the goal of the analysis.

Regarding techniques, the most widely adopted method is the variation of one-parameter-at-a-time. This is not surprising, since it is the most common technique through which the robustness of a model is tested. It is used for example in works such as Levinthal (1997), Zhao and Ma (2016), Hassani-Mahmooei and Parris (2013) and Lenox et al. (2006). A further widely adopted method is scenario analysis. An illustrative example is the work of Leitner et al. (2017): the authors combine variations in 5 parameters to obtain 1296 scenarios on which they evaluate their model. The scenarios are then analyzed graphically, without a formal quantification of sensitivity measures. A similar exercise, but with a smaller number of scenarios, is performed in He et al. (2017) and Zhao and Ma (2016). Global sensitivity methods are applied in fewer studies. Segovia-Juarez et al. (2004) employ uncertainty and sensitivity analysis with Latin hypercube sampling and the partial rank correlation coefficient. In their study, Marino et al. (2008) extend the sensitivity analysis of the Segovia-Juarez et al. (2004) model by performing variance-based sensitivity analysis. For computationally expensive models, we find works that use metamodeling techniques before performing global sensitivity analysis. An example is represented by Dancik et al. (2010), who approximate their model with Gaussian Processes to perform a variance-based sensitivity analysis. Gaussian Process/Kriging emulation is also used in Parry et al. (2013) and Dosi et al. (2018), while Happe et al. (2006) adopt a regression-based metamodel. More recently, Fadikary et al. (2018) propose the use of quantile Gaussian processes as emulators to deal with the stochasticity of the ABM response. Support vector machines are used as a metamodeling tool in Fonoberova et al. (2013).

A third way to analyze the literature is to consider the elements of the ABM that are subjected to investigation. (See Sect. 3 in the main text for the conceptual map). In the majority of cases, the sensitivity analysis concerns only parameters, while all other elements of the model are fixed (e.g., Segovia-Juarez et al. 2004; Fonoberova et al. 2013; Leitner et al. 2017). The parameters that are varied can fall into several categories: some parameters are attributes of agents (e.g., the number N of decisions that organizations face in the NK model), other parameters are properties of the environment (e.g., the number of locations in a spatial model), yet other parameters quantify how many agents of each class exist (Restocchi et al. 2018), etc. Some parameters characterize a behavioral rule, thus impacting a non-parametric element (Riggs et al. 2008). For example, Laurie and Jaggi (2003) study the effect of a vision parameter in the moving rule of a modified version of the Schelling model. Finally, some studies go beyond parameters and change non-parametric elements using scenario analysis, where two or more scenarios are compared while keeping all other elements fixed. For instance, to optimize the distribution of initial alternative fuel vehicles refueling stations, Zhao and Ma (2016) compare four scenarios derived from different targeting strategies. We observe that investigations that involve both parameters and procedures are rarely done, possibly due to the lack of a systematic framework.

We finally note two works that review the use of sensitivity analysis in agent-based modeling before us, namely Lee et al. (2015) and Broeke et al. (2016). Lee et al. (2015) review alternative local and global sensitivity methods and explain the challenges/limitations related to their use. They highlight several technical aspects and the challenges of carrying out sensitivity analyses, given the stochastic, spatial and dynamic nature of the output of ABMs. Broeke et al. (2016) compare one-factor-at-a-time methods with global methods (regression-based and variance-based). Both works focus on methods and on sensitivity analysis of ABMs to variations in input parameters only.

2. Sensitivity methods: quantitative details

In this appendix we provide background information on Sect. 4. We mathematically define sensitivity analysis methods and give more details about visualization tools. Here we only review tools already existing in the literature; our contributions and some specific details of our analysis are described in Appendix 3. As a guide throughout this section, the reader can refer to Table 2.

Preliminaries

For the purposes of sensitivity analysis, we regard any model as a black box input-output mapping that relates inputs X to output Y . Given the scope of our work, we call this black box a “simulator”. Here and in what follows, we use capital letters to indicate the generic name of variables, and small letters to indicate specific values that those variables can take. We write the simulator as $Y = g(\mathbf{X}) + \epsilon(\mathbf{X})$, where $g : \mathcal{X}_{par} \mapsto \mathbb{R}$, \mathcal{X}_{par} is the parametric input space (the subscript *par* stands for parameter), in general a subspace of \mathbb{R}^d , d is the number of simulator inputs, the symbol \mathbf{X} denotes the inputs vector $\mathbf{X} = \{X_1, X_2, \dots, X_d\}$. The term $\epsilon(\mathbf{X})$ is a stochastic error term which, if present, makes the response of the simulator stochastic. The simulator is called deterministic otherwise. In case the simulator is deterministic, the output or quantity of interest is $Y = G = g(\mathbf{X})$. If the simulator is stochastic, then we consider as quantity of interest $Y = \mathbb{E}[G|\mathbf{X}]$, or some other function of the simulator output distribution, such as a given quantile. If the simulator inputs are uncertain, we denote by $(\mathcal{X}_{par}, \mathcal{B}(\mathcal{X}_{par}), \mathbb{P}_{\mathbf{X}})$ the corresponding probability space and by $F_{\mathbf{X}}$ the simulator input (cumulative) distribution.

Local methods

Let \mathbf{x}^0 and \mathbf{x}^1 denote two possible values of the simulator inputs. The corresponding values of the quantity of interest are $y^0 = g(\mathbf{x}^0)$ and $y^1 = g(\mathbf{x}^1)$. We can regard \mathbf{x}^0 and \mathbf{x}^1 as two scenarios, one corresponding to all inputs at the base case and the other with all inputs at the sensitivity case. If the analyst shifts the inputs across these two scenarios, she obtains a variation in the model output response

Table 2 Mapping from goals to methods, depending on the scale (local, global, and our design)

Goal	Method		
	Local	Global	Our design
Factor prioritization	Main effect (2), total effect (6)	Variance-based (10), (11), Moment-independent (15), (14)	Global measures from local effects (17)
Direction of change	Newton quotient (3)	Partial dependence function (16)	Newton quotients at all locations of the design
Interaction quantification	High-order Newton	High-order variance-based (12)	High-order Newton quotients at all locations
Robustness	Scenario analysis	Monte Carlo scenario analysis	Scenario analysis with statistical tests

Numbers refer to the equation in which each method is defined

equal to $\Delta y = y^1 - y^0$. In the literature, methods have been developed to quantitatively explain the change Δy , apportioning it to the main (individual) effects of each input, to the interactions of pairs, triplets of inputs, etc.. Formally, several works have shown that Δy can be expanded as (Li et al. 2001)

$$\Delta y = \sum_{i=1}^d \phi_i^{0 \rightarrow 1} + \sum_{i < j} \phi_{i,j}^{0 \rightarrow 1} + \dots + \phi_{1,2,\dots,d}^{0 \rightarrow 1}, \tag{1}$$

where

$$\begin{cases} \phi_i^{0 \rightarrow 1} = g(x_i^1 : \mathbf{x}_{-i}^0) - g(\mathbf{x}^0), \\ \phi_{i,j}^{0 \rightarrow 1} = g(x_{i,j}^1 : \mathbf{x}_{-(i,j)}^0) - g(\mathbf{x}^0) - \phi_i^{0 \rightarrow 1} - \phi_j^{0 \rightarrow 1}, \\ \dots \end{cases} \tag{2}$$

In the above equation, the symbol $(x_i^1 : \mathbf{x}_{-i}^0)$ denotes the point in \mathcal{X}_{par} obtained by setting x_i at the sensitivity case, while keeping all remaining inputs at the base case. Similarly, the symbol $(x_{i,j}^1 : \mathbf{x}_{-(i,j)}^0)$ denotes the point in \mathcal{X}_{par} obtained by setting x_i and x_j at the sensitivity case, while keeping the remaining inputs at the base case.

The first order terms in Eq. (1), denoted by $\phi_i^{0 \rightarrow 1}$, are called *main effects* of the model inputs Li et al. (2001). These effects are obtained by what is known as a “one-at-a-time” experimental design, in the sense that they are obtained by varying one input only. Main effects can be visualized through Tornado Diagrams (Eschenbach 1992). The normalized version of main effects,

$$\omega_i = \frac{\phi_i^{0 \rightarrow 1}}{x_i^1 - x_i^0} = \frac{g(x_i^1 : \mathbf{x}_{-i}^0) - g(\mathbf{x}^0)}{x_i^1 - x_i^0}, \tag{3}$$

is the *Newton quotient* of g at \mathbf{x}^0 for the change $\Delta = x_i^1 - x_i^0$. When considered as a function of x_i , the numerator in Eq. (3) contains the univariate function $z_i(x_i) = g(x_i : \mathbf{x}_{-i}^0)$, which is the one-way sensitivity function of g with respect to x_i .

The second order terms in Eq. (1), denoted by $\phi_{i,j}^{0 \rightarrow 1}$, quantify the residual interaction between x_i and x_j . These terms are *interaction effects* Kleijnen (2015). Dividing the second order effects by $(x_i^1 - x_i^0)(x_j^1 - x_j^0)$, we obtain the *second order Newton quotients*:

$$\omega_{i,j} = \frac{\phi_{i,j}^{0 \rightarrow 1}}{(x_i^1 - x_i^0)(x_j^1 - x_j^0)}. \tag{4}$$

If inputs are denominated in different units, the second order Newton quotients are not directly comparable. In order to homogenize the scale on the horizontal axes of the graphs in Fig. 5, we reported the distribution of

$$\omega_{i,j}^* = \frac{\phi_{i,j}^{0 \rightarrow 1} |(x_i^1 - x_i^0)(x_j^1 - x_j^0)|}{(x_i^1 - x_i^0)(x_j^1 - x_j^0)}. \quad (5)$$

The sign of $\omega_{i,j}$ (or $\omega_{i,j}^*$) delivers information about the sign of the pairwise interaction between X_i and X_j at \mathbf{x}^0 . A positive $\omega_{i,j}$ implies a synergistic interaction, a negative $\omega_{i,j}$ implies an antagonistic interaction. A similar interpretation holds for higher interactions and for higher order terms in (1).

In the literature, one also defines the *total finite change effect* of x_i (Li et al. 2001). This effect is the sum of all the terms in Eq. (1) that contain index i :

$$\tau_i^{0 \rightarrow 1} = \phi_i^{0 \rightarrow 1} + \sum_{j \neq i} \phi_{i,j}^{0 \rightarrow 1} + \dots + \phi_{1,2,\dots,d}^{0 \rightarrow 1}. \quad (6)$$

The total effect $\tau_i^{0 \rightarrow 1}$ is the fraction of Δy that can be apportioned to x_i , due to its individual effect and to all its interactions with the remaining inputs. A visualization that makes it possible to see main, total and interaction effects is the Generalized Tornado Diagram (Borgonovo and Smith 2011). In order to compute all terms in the decomposition in Eq. (1), one needs to evaluate the model at all vertices of the hypercube joining \mathbf{x}^0 and \mathbf{x}^1 . Thus, one is effectively considering a full factorial design between the two model inputs (see below).

Towards global methods: experimental designs

In a global sensitivity analysis the model is explored at N locations, x^0, x^1, \dots, x^N . If parameters are discrete then a finite (albeit large) N may suffice in exploring the entire input space, while if parameters are continuous there are infinitely many points to inspect. For space limitations, we present intuition about some designs, referring the reader to Chapter 2 in Saltelli et al. (2008) and to Barton (2015) for a more extensive treatment. We start from discrete input spaces. A standard choice is to consider all possible combinations of the values that model inputs can take—this is known as a full factorial design (Montgomery 2000). Of course, running the simulator at all these input points can be extremely expensive, so computationally cheaper solutions have been proposed. For example, one may consider fractional factorial designs (Montgomery 2000). These make it necessary to run the simulator at only a subset of the points of the full factorial design, selected so as to be able to calculate interaction effects up to a certain order. In principle, the use of an experimental design does not require that inputs are assigned a probability distribution. In the case inputs are assigned a probability distribution, then points are sampled from this distribution. Several sampling techniques are available, among which crude Monte-Carlo, Quasi-Monte Carlo and Latin Hypercube Sampling (Le Guiban et al. 2018). As mentioned in the main text, some sampling techniques ensure that the input space is explored uniformly, without oversampling a portion of the input space and undersampling another part. These sampling techniques are called *low-discrepancy sequences*; a popular one is the Halton sequence (Saltelli et al. 2008).

Global methods: variance-based

Among global sensitivity methods, variance-based methods play an important role. Let us assume, for the moment, that inputs are independent. One then has the possibility to decompose the simulator output variance in an equation analogous to Eq. (1). As proven in Efron and Stein (1981), we can write

$$\sigma_Y^2 = \sum_{i=1}^d \sigma_{Y:i}^2 + \sum_{i < j} \sigma_{Y:ij}^2 + \dots + \sigma_{Y:1,2,\dots,d}^2, \quad (7)$$

where

$$\begin{aligned} \sigma_{Y:i}^2 &= \int g_i(x_i)^2 dF_i(x_i), \\ \sigma_{Y:ij}^2 &= \int g_{ij}(x_{ij})^2 dF_{ij}(x_i, x_j). \\ &\dots \end{aligned} \quad (8)$$

In the above equation, $\sigma_{Y:i}^2$ is the portion of the variance attributed to X_i alone, $\sigma_{Y:ij}^2$ the portion of the variance attributed to the residual interaction between X_i and X_j , etc. The function F denotes the cumulative distribution with respect to its arguments, and the functions $g_i(x_i)$, $g_{ij}(x_i, x_j)$ are obtained by taking partial expectations of Y through the following equations:

$$\begin{cases} g_i(x_i) = \mathbb{E}[Y|X_i = x_i] - \mathbb{E}[Y], \\ g_{ij}(x_i, x_j) = \mathbb{E}[Y|X_i = x_i, X_j = x_j] - g_j(x_j) - g_i(x_i) - \mathbb{E}[Y], \\ \dots \end{cases} \quad (9)$$

We observe that $\mathbb{E}[Y|X_i = x_i]$ is the conditional regression function of Y with respect to X_i . In practice, explicitly computing the integrals in Eq. (8) and the expected values in Eq. (9) is unfeasible, so practitioners using variance-based methods use techniques such as partitioning the input space Saltelli et al. (2008). We discuss in Appendix 3 how we compute variance-based measures for our design.

In the literature, one considers the normalized version

$$S_i = \frac{\sigma_{Y:i}^2}{\sigma_Y^2}, \quad (10)$$

called *first order Sobol's index*, giving the relative contribution of input i to total variance. One also defines the total sensitivity index as (Homma and Saltelli 1996):

$$\tau_i^2 = \frac{\sigma_{Y:i}^2 + \sum_{j=1, j \neq i}^d \sigma_{Y:ij}^2 + \dots + \sigma_{Y:1,2,\dots,d}^2}{\sigma_Y^2}, \quad (11)$$

giving to the relative contribution to variance of input i , taking into account its interactions with other inputs. The total effect τ_i^2 has been proposed as a sensitivity measure in the operations research field by Wagner (1995). Let $u \subseteq \{1, 2, \dots, n\}$ be a subset of indices, say $u = i_1, i_2, \dots, i_k$, with $k \leq n$. The variance-based index

$S_u = \frac{\sigma_{Y:u}^2}{\sigma_Y^2}$ is a global measure of the interaction among the parameters in group u .

Specifically, when $u = \{i, j\}$,

$$S_{ij} = \frac{\sigma_{Y:ij}^2}{\sigma_Y^2} \quad (12)$$

is a global measure of the strength of the pairwise interaction between x_i and x_j . We recall that when inputs are dependent, first order variance-based sensitivity measures S_i remain well posed. However, the interpretation of higher and total order indexes becomes less clear (Kucherenko et al. 2012).

Global methods: distribution-based

A further class of sensitivity measures whose definition remains well posed under input dependence is the class of “distribution-based” or “moment-independent” sensitivity measures. These sensitivity measures are obtained considering the entire distribution of model output instead of a particular moment (e.g., variance). In general, a distribution-based sensitivity measure can be written as

$$\xi_i = \mathbb{E}[d(F_Y, F_{Y|X_i})], \quad (13)$$

where $d(\cdot, \cdot)$ is a distance or divergence between distributions, F_Y is the marginal distribution of the simulator output and $F_{Y|X_i}$ is the conditional distribution of the simulator output given X_i . Depending on the choice of $d(\cdot, \cdot)$, one obtains alternative ways of measuring importance. For instance, if d is set equal to the Kuiper distance between cumulative distributions functions, one obtains the global importance measure (Baucells and Borgonovo 2013)

$$\beta_i^{Ku} = \mathbb{E}[\sup_y \{F_{Y|X_i}(y) - F_Y(y)\} + \sup_y \{F_Y(y) - F_{Y|X_i}(y)\}]. \quad (14)$$

Alternatively, if Y is absolutely continuous, one can choose the L^1 -norm between densities f_Y and $f_{Y|X_i}$, calculating the so called δ -importance measure Borgonovo et al. (2014):

$$\delta_i = \frac{1}{2} \mathbb{E} \left[\int |f_Y(y) - f_{Y|X_i}(y)| dy \right]. \quad (15)$$

We observe that, in the case the model output distribution is unimodal, then β_i^{Ku} and δ_i coincide [see Borgonovo et al. (2014)]. It is important to note that sensitivity measures τ_i^2 , β_i^{Ku} and δ_i possess the nullity-implies-independent property. That is, a null value of any of these sensitivity measures occurs if and only if Y is independent of X_i . First order variance-based sensitivity measures (S_i) do not possess this property.

Finally, other global methods (which we cannot describe due to space limitations) include regression-based methods Storlie et al. (2009), Shapley values Owen (2014) and others Ghanem et al. (2016).

Global methods: direction of change

All methods listed so far provide indications concerning factor prioritization. However, global methods exist for direction of change and interaction quantification as well. For example, a standard method for direction of change is based on integrating g over the marginal distributions of all inputs but X_i . This leads to *partial dependence functions* Friedman (2001), a family of well known indicators used in machine learning. The expression of a partial dependence function is

$$h_i(x_i) = \int_{\mathcal{X}_i} g(x_i; \mathbf{x}_{-i}) dF_{\mathbf{X}_{-i}}(\mathbf{x}_{-i}). \quad (16)$$

This function is visualized through a partial dependence plot. We note that $h_i(x_i)$ in Eq. (16) is the expectation of the functions $g(x_i; \mathbf{x}_{-i})$ with respect to $F_{\mathbf{X}_{-i}}(\mathbf{x}_{-i})$. Consider a given point \mathbf{x}_{-i}^0 ; the function $z_i(x_i) = g(x_i; \mathbf{x}_{-i}^0)$ is the one-way sensitivity function of X_i at \mathbf{x}_{-i}^0 [see, among others (Borgonovo (2017, Ch. 4)]. Note that $g(x_i; \mathbf{x}_{-i}^0)$ is termed an individual conditional expectation (ICE) function in Goldstein et al. (2015). Thus, a partial dependence function is the average of one-way sensitivity functions (or of ICE functions).

3. Our design, sensitivity measures and S-ICE plots

This section presents the proposal for a design that allows the simultaneous variation of several (possibly all) elements of an ABM and the estimation of global sensitivity measures.

Input space

All local and global sensitivity methods listed in Appendix 2 are defined on numeric input spaces. To include non-parametric elements of an ABM into the sensitivity analysis toolbox, consider the categorical variable X_p with support $\mathcal{X}_p = \{0, 1, \dots, x_p\}$. In the case $X_p = x_p$, the simulator is run under the $p^{th} + 1$ available procedure. For example, in the case of the GCM, X_p could be problem structure. We would have $\mathcal{X}_p = \{0, 1\}$, corresponding to anarchy and hierarchy.⁵ We can then enlarge the input space to account for all combinations of parameter values and possible assigned procedures. More generally one can assign categorical variables to any non-parametric element of an agent-based simulator, considering the enlarged input space $\mathcal{X}^{ABM} = \mathcal{X}_{par} \times \mathcal{X}_{non-par}$. The analyst can then assign any desired distribution on the

⁵ The assignment of the digit 0 or 1 to either state is a matter of convention. To avoid interpretation of the levels of the categorical variable with numeric inputs, in the main text we use letters A and H.




Fig. 6 The points of the experimental design used in this work. Left graph: two inputs, x_1 and x_2 with two values, 0 and 1, and one level for each value. Right graph: two inputs, x_1 and x_2 ; x_1 with two values, 0 and 1, and one level; x_2 with three values, 0, 1 and 2, and two levels. Blue arrows: individual effects; red arrows, interaction effects. (Color figure online)

measurable space $(\mathcal{X}^{ABM}, \mathcal{B}(\mathcal{X}^{ABM}))$. This opens the door to using any fully fledged global sensitivity analysis method and design.

Definition of the lattice

We aim at a scheme that leads to a full factorial design. To do so, we propose to take as support for the parameters the set $\mathcal{X}_{par} = \mathcal{X}_1 \times \mathcal{X}_2 \times \dots \times \mathcal{X}_d$, with $\mathcal{X}_i = \{x_i^{low}, x_i^{mid}, x_i^{high}\}$, where x_i^{low} and x_i^{high} are the lowest and upper boundary values assigned to X_i , and x_i^{mid} is an intermediate value. Then, \mathcal{X}^{ABM} becomes a lattice, whose points are all combinations of the ABM input elements. In case the simulator is expensive to run, one can use cheaper experimental designs, such as the ones listed in Appendix 2.

Local effects at multiple locations

Evaluating the simulator on this lattice, the analyst is effectively evaluating the model at all points necessary to obtain the full finite change decomposition in Eq. (1). Not only, but the lattice allows the analyst to obtain several replicates of first order effects and higher order finite change effects for each input. Evaluating many local effects is by no means novel in sensitivity analysis. For example, it is at the basis of the classical method of Morris (1991). By that method, the analyst computes main effects at a certain number of randomly sampled points in the lattice, and then obtains global measures as appropriate averages of the random main effects. However, the Morris method only makes it possible to address the factor prioritization

goal. Our design instead makes it easy to also address direction of change and interaction quantification.

As a reference, consider Fig. 6. The left graph displays a two-dimensional design, with two inputs and two values (one level) per input. We have four possible combinations, displayed as vertices of the rectangle. The arrows evidence shifts of the inputs from one point to another on the grid. Movements parallel to the vertical and horizontal axis are “one-at-a-time shifts”; diagonal movements are instead “interaction shifts”. Evaluating the model at all points in the lattice allows us to obtain: a) the complete decomposition of the finite change $\Delta^{A \rightarrow C} y = g(x_1^1, x_2^1) - g(x_1^0, x_2^0)$, with the main effects $\phi_1^{A \rightarrow C} = g(x_1^1, x_2^0) - g(x_1^0, x_2^0)$, $\phi_2^{A \rightarrow C} = g(x_1^0, x_2^1) - g(x_1^0, x_2^0)$ and the interaction effect $\phi_{1,2}^{A \rightarrow C} = g(x_1^1, x_2^1) - g(x_1^0, x_2^0) - \phi_1^{A \rightarrow C} - \phi_2^{A \rightarrow C}$. But we can also find the main effects $\phi_1^{B \rightarrow D} = g(x_1^0, x_2^1) - g(x_1^0, x_2^0)$, $\phi_2^{B \rightarrow D} = g(x_1^1, x_2^1) - g(x_1^1, x_2^0)$ and the interaction effect $\phi_{1,2}^{B \rightarrow D} = g(x_1^1, x_2^1) - g(x_1^1, x_2^0) - \phi_1^{B \rightarrow D} - \phi_2^{B \rightarrow D}$. These effects are associated with the arrows displayed in the graph (blue for individual and red for interaction effects). The evaluation of the model at all points of the lattice, by the above rearrangement, allows us to find two main effects and two finite change interaction indices. (These are all the distinct effects, however one could consider additional effects or alternative versions of these effects by flipping the sign—inverting the direction of the arrows in Fig. 6.) The right graph displays a two dimensional design in which x_1 has one levels while x_2 has two levels. In this case, the total number of points in the grid is six, and combination of these six model evaluations allows us to find six individual effects for each input and six second order interaction effects.

In general, letting N_i denote the number of main effects ($\phi_i^{0 \rightarrow 1}$) associated with x_i , we have a lattice of $N = \prod_{i=1}^d k_i$ points, where k_i is the number of levels of each variable. The number of first order effects associated with input x_i is $N_i = \frac{k_i}{2} \prod_{l=1}^d (k_l + 1)$. The number of second order effects for the pair x_i, x_j is $N_{i,j} = \frac{k_i k_j}{2} \prod_{l=1}^d (k_l + 1)$. One can verify these formulas for the examples in Fig. 6. For the case in the left graph, both variables x_1 and x_2 have only one level ($k_1 = k_2 = 1$). Thus, the number of first order effects for each variable is $N_1 = N_2 = 1/2 \cdot 2 \cdot 2 = 2$, giving a total of four first order effects. The total number of interactions is $N_{1,2} = 1/2 \cdot 2 \cdot 2 = 2$. In the right graph, variable x_1 has again one level, while x_2 has two levels ($k_1 = 1, k_2 = 2$). So, $N_1 = 1/2 \cdot 2 \cdot 3 = 3$ and $N_2 = 2/2 \cdot 2 \cdot 3 = 6$, giving a total of nine first order effects. The total number of interactions is $N_{1,2} = 2/2 \cdot 2 \cdot 3 = 6$.

For the GCM analysis in Sect. 6, in the notation of this section, we have $k_1 = \dots = k_4 = 1$ and $k_5 = \dots = k_{10} = 2$, so that $N = 5832$, from which we compute $N_1 = \dots = N_4 = 5832$ and $N_5 = \dots = N_{10} = 11664$ first order indices and a grand total of 664, 848 second order interaction effects.

Sensitivity measures from the design

Once these effects are available, we calculate sensitivity measures. So far, we have been defining global sensitivity measures in an abstract way. However, when we are dealing with a simulation model with no explicit functional form for g , it is impossible to evaluate the integrals in Eq. (8). Several methods have been proposed for this task, with corresponding experimental designs (Saltelli et al. 2008). Given the

discrete nature of the input space implied by our experimental design (Appendix 3), we use a method that obtains global measures (such as the total index τ_i^2), from local measures (such as finite change sensitivity indexes $\phi_i^{0 \rightarrow 1}$), evaluated at many points in the input space. Consider $N + 1$ locations x^0, x^1, \dots, x^N and compute a series of indexes $\phi_i^{0 \rightarrow 1}, \phi_i^{1 \rightarrow 2}, \dots, \phi_i^{N-1 \rightarrow N}$. It is possible to see that the following equation holds Gamboa et al. (2016):

$$\tau_i^2 = \frac{\mathbb{V}[\Phi_i]}{2N}, \quad (17)$$

where Φ_i is a random variable denoting the population of first order effects (that is, Φ_i is a random variable whose realizations are $\phi_i^{0 \rightarrow 1}, \phi_i^{1 \rightarrow 2}, \dots, \phi_i^{N-1 \rightarrow N}$). This means that calculating the variance of the available first order effects and dividing by two leads to a proxy (an estimator) of the total index of the corresponding simulator input.

Specifically, from the main effects it is then possible to compute: (a) N_i Newton quotients $\omega_i^{0 \rightarrow 1}$ for each input and (b) one total variance-based effect for each input from Eq. (17). From interaction effects, one can calculate second order Newton quotients, gaining insights on the sign and magnitude of interactions.

Moreover, the design yields a dataset of $N \times d$ input-output realizations. This input-output sample can be used to calculate the total order variance-based index using Eq. 17, and one can also estimate the global sensitivity measures S_i (first order variance-based), β^{Ku} and δ through a so-called “given data approach” Plischke et al. (2013). [Specifically, we use the betaKS3.m function available at <https://zenodo.org/record/885332#.XgoB-kdKiUk>.]

S-ICE plots

An ICE plot is made of two main constituents: first, the graph of a partial dependence function $h_i(x_i)$ in Eq. (16), second the graphs of a number of ICE functions. The simultaneous use of these two constituents was suggested by Goldstein et al. (2015) to enrich the information provided by partial dependence functions. In fact, the marginal expectation at the basis of $h_i(x_i)$ produces a smoothing that hides patterns realized for specific values of \mathbf{x}_{-i}^0 . Instead, such patterns are potentially revealed if one plots also the $g_i(x_i; \mathbf{x}_{-i}^0)$ functions, as Goldstein et al. (2015) observe.

In our case, regarding the first constituent, we observe that marginal expectation coincides with conditional expectation in our design. Denoting the realizations of X_i with $(x_i^1, x_i^2, \dots, x_i^K)$, the partial dependence plot represents the set of pairs $(x_i^k, \mathbb{E}[Y|X_i = x_i^k])$, $k = 1, 2, \dots, K$. We shall use big black circles on the vertical axis to represent $\mathbb{E}[Y|X_i = x_i^k]$. Regarding the graph of ICE functions, consider pair of points of the type $(x_i^0, g(x_i^0; \mathbf{x}_{-i}^k))$ and $(x_i^1, g(x_i^1; \mathbf{x}_{-i}^k))$. Note that the two values $g(x_i^0; \mathbf{x}_{-i}^k)$ and $g(x_i^1; \mathbf{x}_{-i}^k)$ belong to the univariate function $z_i(x_i)$, whose graph is represented by two points on the vertical axis (one at x_i^0 and the other at x_i^1) in the case x_i is binary. To evidence pairs of corresponding points $(x_i^0, g(x_i^0; \mathbf{x}_{-i}^k))$ and $(x_i^1, g(x_i^1; \mathbf{x}_{-i}^k))$, we join such pairs with a line. We use blue if the output change is positive, red if the output

change is negative and gray if it is non-statistically significant. In fact, we recall that $g(x_i^0; \mathbf{x}_{-i}^k) = \mathbb{E}[Y | \mathbf{X} = (x_i^0; \mathbf{x}_{-i}^k)]$ is the result of an expectation over simulation replicates. That is, the quantities reported in the graph are $\hat{g}(x_i^0; \mathbf{x}_{-i}^k) = \hat{\mathbb{E}}[Y | \mathbf{X} = (x_i^0; \mathbf{x}_{-i}^k)]$. That is, due to numerical noise, we may register $\hat{g}(x_i^0; \mathbf{x}_{-i}^k) - \hat{g}(x_i^1; \mathbf{x}_{-i}^k) \neq 0$, albeit small, even if in reality it is $g(x_i^0; \mathbf{x}_{-i}^k) - g(x_i^1; \mathbf{x}_{-i}^k) = 0$. Then, visualizing the change with a blue or red line would be misleading, because, in fact, the change would not be statistically significant. We then use the Student t-test for differences in the two conditional means $\hat{\mathbb{E}}[Y | X_i = x_i^0]$ and $\hat{\mathbb{E}}[Y | X_i = x_i^1]$ to assess the statistical significance of the change. In the case the test provides a “not significant” report, a gray line is reported.

The prototypical representation of an ICE plots involves continuous functions with continuous inputs on an interval. In this case, the graphs of $h_i(x_i)$ and of each $g(x_i^0; \mathbf{x}_{-i}^k)$ are continuous lines. Implementation-wise, the graph is obtained by discretizing the values of X_i into a dense enough grid and then reporting the corresponding values of $h_i(x_i)$ and of each $g(x_i^0; \mathbf{x}_{-i}^k)$ on the y-axis to render the visual impression of continuity. Nonetheless, the definitions of ICE functions as well as partial dependence plots hold unchanged when inputs are discrete. The difference in the visual representation is that on the vertical axis one registers a sequence of points (dots) rather than a continuous line. To illustrate, in the case of binary inputs with values x_i^0 and x_i^1 , a partial dependence plot reduces to two points $(x_i^0, h_i(x_i^0))$ and $(x_i^1, h_i(x_i^1))$; the plot of any individual ICE function is a pair of points of the type $(x_i^0, g(x_i^0; \mathbf{x}_{-i}^k))$ and $(x_i^1, g(x_i^1; \mathbf{x}_{-i}^k))$.

We also note that ICE plots deliver information about whether g responds additively (or not) to variations in x_i . In fact we can prove that if g is additive in X_i , and X_i is binary, then all slopes in ICE functions are parallel.

Proposition *If model input x_i is binary and the output is additive in x_i , then all lines in the direction plot of x_i are parallel.*

Proof If x_i is binary, let x_i^- and x_i^+ denote its two possible values, with $x_i^- < x_i^+$. Let $a_i(x_i^-)$, $a_i(x_i^+)$ denote the values of the univariate function $a_i(x_i)$ at x_i^- and x_i^+ , respectively. If g is additive in x_i then it can be written as $g(\mathbf{x}) = a_i(x_i) + a_{-i}(\mathbf{x}_{-i})$. Consider next the one-way sensitivity function $z_i(x_i) = g(x_i; \mathbf{x}_{-i}^k)$. At any point $(x_i^+, \mathbf{x}_{-i}^k)$ we have: $g(x_i^+; \mathbf{x}_{-i}^k) = g_i(x_i^+) + g_{-i}(\mathbf{x}_{-i}^k)$ and at any point $(x_i^-, \mathbf{x}_{-i}^k)$, we have: $g(x_i^-; \mathbf{x}_{-i}^k) = g_i(x_i^-) + g_{-i}(\mathbf{x}_{-i}^k)$. Therefore, we find

$$\phi_i^{- \rightarrow +} = g_i(x_i^+) + g_{-i}(\mathbf{x}_{-i}^k) - (g_i(x_i^-) + g_{-i}(\mathbf{x}_{-i}^k)) = g_i(x_i^+) - g_i(x_i^-).$$

Note that the above equation implies the constancy of the Newton quotients, q.e.d. \square

Consequently, one or more changes in slope in the S-ICE plot of a binary variables implies that the input of interest is involved in interactions with the other inputs.

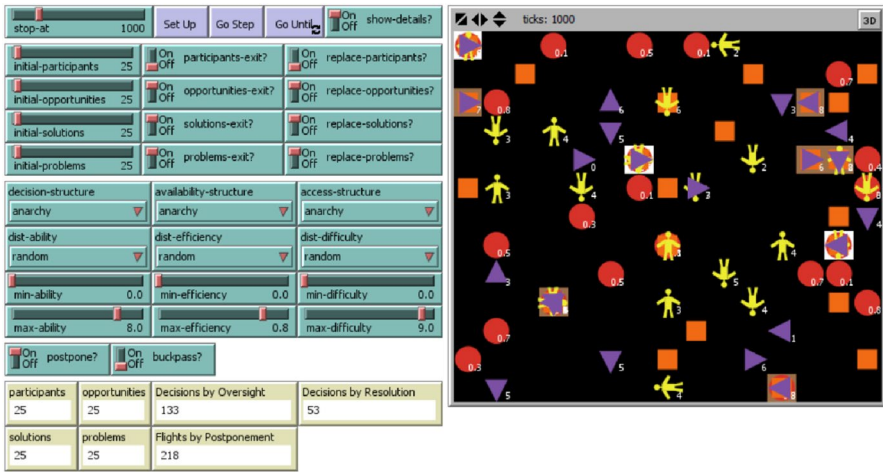


Fig. 7 NetLogo interface of the Garbage Can Model implementation

4. NetLogo implementation

In this section, we take a look at the NetLogo interface of the Garbage Can Model by Fioretti and Lomi⁶ (Fig. 7). This particular software implementation includes some extensions of the traditional GCM, such as the possibility to backpass difficult problems and to have a competence/incompetence system to assign parameters for ability, efficiency and difficulty. To remain as close as possible to the original theory of the GCM, we do not consider these extensions in our study. However, it would be straightforward to introduce these elements in the sensitivity analysis to gain more insights into their role within the model.

In our simulation, we adopt a 15×13 grid where 25 agents of each class move in a random fashion. The number of agents is proportional to the size of the grid so that there is a non-trivial amount of interactions between agents; note that, once the density of agents is fixed, the share of decisions of each kind is not affected by simultaneous changes in the number of agents and the size of the grid. Participants are represented on the grid with the man icon, solutions with a triangle, problems with a circle and choice opportunities with a square. Ability, efficiency and difficulty are shown at the bottom of the correspondent icon. A run of the simulation is composed of 1000 time steps, as this is enough to have stable values for the share of the different decision styles. Participants stay the same during the course of the simulation, while solved problems and adopted solutions leave the grid and are replaced with new ones with different difficulty and efficiency. Ability, efficiency and difficulty are assigned independently of the ranking of each agent within the organization.

⁶ http://www.cs.unibo.it/~fioretti/CODE/GC/GarbageCan_buck.nlogo.tar.gz.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10588-021-09358-5>.

Acknowledgements The authors wish to thank the editor, Professor Terril L. Frantz, for the editorial attention and three anonymous reviewers for their very perceptive comments that have greatly improved the manuscript.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Amini M, Wakolbinger T, Racer M, Nejad MG (2012) Alternative supply chain production-sales policies for new product diffusion: an agent-based modeling and simulation approach. *Eur J Oper Res* 216(2):301–311
- Anderson P (1999) Complexity theory and organization science. *Organ Sci* 10(3):216–232
- Ayer T, Zhang C, Bonifonte A, Spaulding AC, Chhatwal J (2019) Prioritizing hepatitis c treatment in us prisons. *Oper Res* 67(3):853–873
- Barnes S, Golden B, Wasil E (2010) MRSAT transmission reduction using agent-based modeling and simulation. *INFORMS J Comput* 22(4):635–646
- Barnes SL, Myers M, Rock C, Morgan DJ, Pineles L, Thom KA, Harris AD (2020) Evaluating a prediction-driven targeting strategy for reducing the transmission of multidrug-resistant organisms. *INFORMS J Comput* 32(4):912–929
- Barton RR (2015) Tutorial: simulation metamodeling. In: *Proceedings of the 2015 winter simulation conference. WSC '15*. IEEE Press, Piscataway, pp 1765–1779
- Baucells M, Borgonovo E (2013) Invariant probabilistic sensitivity analysis. *Manage Sci* 59(11):2536–2549
- Baumann O, Schmidt J, Stieglitz N (2018) Effective search on rugged performance landscapes: a review and outlook. *J Manage* 45(1):285–318
- Bendor J, Moe TM, Shotts KW (2001) Recycling the garbage can: an assessment of the research program. *Am Polit Sci Rev* 95(1):169–190
- Borgonovo E (2017) *Sensitivity analysis: an introduction for the management scientist*. Springer, New York
- Borgonovo E, Smith CL (2011) A study of interactions in the risk assessment of complex engineering systems: an application to space PSA. *Oper Res* 59(6):1461–1476
- Borgonovo E, Tarantola S, Plischke E, Morris MD (2014) Transformation and invariance in the sensitivity analysis of computer experiments. *J R Stat Soc Ser B* 76(5):925–947
- Brailsford SC, Eldabi T, Kunc M, Mustafee N, Osorio AF (2019) Hybrid simulation modelling in operational research: a state-of-the-art review. *Eur J Oper Res* 278(3):721–737
- Carley KM (2002) Computational organization science: a new frontier. *Proc Natl Acad Sci USA* 99(SUPPL. 3):7257–7262
- Clement J, Puranam P (2018) Searching for structure: formal organization design as a guide to network evolution. *Manage Sci* 64(8):3879–3895
- Cohen MD, March JG, Olsen JP (1972) A garbage can model of organizational choice. *Adm Sci Q* 17(1):1–25
- Currie CSM, Fowler JW, Kotiadis K, Monks T, Onggo BSA, Robertson DA, Tako AA (2020) How simulation modelling can help reduce the impact of COVID-19. *J Simul* 14(2):83–97

- Dancik GM, Jones DE, Dorman KS (2010) Parameter estimation and sensitivity analysis in an agent-based model of *Leishmania major* infection. *J Theor Biol* 262(3):398–412
- Delre SA, Broekhuizen TLJ, Bijmolt THA (2016) The effects of shared consumption on product life cycles and advertising effectiveness: the case of the motion picture market. *J Market Res* 53(4):608–627
- Dosi G, Levinthal DA, Marengo L (2003) Bridging contested terrain: linking incentive-based and learning perspectives on organizational evolution. *Ind Corp Change* 12(2):413–436
- Dosi G, Pereira MC, Virgillito ME (2018) On the robustness of the fat-tailed distribution of firm growth rates: a global sensitivity analysis. *J Econ Interact Coord* 13(1):173–193
- Efron B, Stein C (1981) The Jackknife estimate of variance. *Ann Stat* 9(3):586–596
- Eschenbach TG (1992) Spiderplots versus Tornado diagrams for sensitivity analysis. *Interfaces* 22(6):40–46
- Fadikary A, Higdon D, Chenx J, Lewis B, Venkatramanan S, Marathe M (2018) Calibrating a stochastic, agent-based model using quantile-based emulation. *SIAM/ASA J Uncertain Quantif* 6(4):1685–1706
- Fibich G, Gibori R (2010) Aggregate diffusion dynamics in agent-based models with a spatial structure. *Oper Res* 58(5):1450–1468
- Fioretti G, Lomi A (2008) An agent-based representation of the garbage can model of organizational choice. *J Artif Soc Soc Simul* 11(1):1
- Fioretti G, Lomi A (2010) Passing the buck in the garbage can model of organizational choice. *Comput Math Organ Theory* 16(2):113–143
- Fonoberova M, Fonoberov VA, Mezić I (2013) Global sensitivity/uncertainty analysis for agent-based models. *Reliab Eng Syst Saf* 118:8–17
- Friedman JH (2001) Greedy function approximation: a gradient boosting machine. *Ann Stat* 29(5):1189–1232
- Gallegati M, Richiardi MG (2009) Agent based models in economics and complexity. In: Meyers R (ed) *Encyclopedia of complexity and systems science*. Springer, New York, NY. https://doi.org/10.1007/978-0-387-30440-3_14
- Gamboa F, Janon A, Klein T, Lagnoux A, Prieur C (2016) Statistical inference for Sobol pick-freeze Monte Carlo method. *Statistics* 50(4):881–902
- Garcia R, Jager W (2011) From the special issue editors: agent-based modeling of innovation diffusion. *J Prod Innov Manage* 28(2):148
- Ghanem R, Higdon D, Owhadi D (eds) (2016) *Handbook of uncertainty quantification*. Springer, Cham
- Glynn PW, Greve HR, Rao H (2020) Relining the garbage can of organizational decision-making: modeling the arrival of problems and solutions as queues. *Ind Corp Change* 29(1):125–142
- Goldstein A, Kapelner A, Bleich J, Pitkin E (2015) Relining the garbage can of organizational decision-making: modeling the arrival of problems and solutions as queues. *J Comput Graph Stat* 24(1):44–65
- Happe K, Kellermann K, Balmann A (2006) Agent-based analysis of agricultural policies: an illustration of the agricultural policy simulator *agripolis*, its adaptation and behavior. *Ecol Soc* 11(1):49
- Harrison JR, Lin Z, Carroll GR, Carley KM (2007) Simulation modeling in organizational and management research. *Acad Manage Rev* 32(4):1229–1245
- Hassani-Mahmooui B, Parris BW (2013) Resource scarcity, effort allocation and environmental security: an agent-based theoretical approach. *Econ Model* 30:183–192
- He Z, Xiong J, Ng TS, Fan B, Shoemaker CA (2017) Managing competitive municipal solid waste treatment systems: an agent-based approach. *Eur J Oper Res* 263(3):1063–1077
- Hernandez E, Menon A (2018) Acquisitions, node collapse, and network revolution. *Manage Sci* 64(4):1652–1671
- Homma T, Saltelli A (1996) Importance measures in global sensitivity analysis of nonlinear models. *Reliab Eng Syst Saf* 52(1):1–17
- Jaspersen JG, Peter R (2017) Experiential learning, competitive selection, and downside risk: a new perspective on managerial risk taking. *Organ Sci* 28(5):915–930
- Jennings NR (2000) On agent-based software engineering. *Artif Intell* 117(2):277–296
- Jiang G, Tadikamalla PR, Shang J, Zhao L (2016) Impacts of knowledge on online brand success: an agent-based model for online market share enhancement. *Eur J Oper Res* 248(3):1093–1103
- Keuschnigg M, Ganser C (2017) Crowd wisdom relies on agents' ability In small groups with a voting aggregation rule. *Manage Sci* 63(3):818–828
- Kleijnen JPC (2015) *Design and analysis of simulation experiments*, 2nd edn. Springer, Cham

- Kucherenko S, Tarantola S, Annoni P (2012) Estimation of global sensitivity indices for models with dependent variables. *Comput Phys Commun* 183:937–946
- Laurie AJ, Jaggi NK (2003) Role of 'vision' in neighbourhood racial segregation: a variant of the schelling segregation model. *Urban Stud* 40(13):2687–2704
- Le Guiban K, Rimmel A, Weisser MA, Tomasik J (2018) The first approximation algorithm for the maximin Latin hypercube design problem. *Oper Res* 66(1):256–266
- Lee JS, Filatova T, Ligmann-Zielinska A, Hassani-Mahmooei B, Stonedahl, F, Lorscheid I, Voinov A, Polhill G, Sun Z, Parker DC, (2015) The complexities of agent-based modeling output analysis. *J Artif Soc Soc Simul* 18(4):1–27
- Leitner S, Rausch A, Behrens DA (2017) Distributed investment decisions and forecasting errors: an analysis based on a multi-agent simulation model. *Eur J Oper Res* 258(1):279–294
- Lenox MJ, Rockart SF, Lewin AY (2006) Interdependency, competition, and the distribution of firm and industry profits. *Manage Sci* 52(5):757–772
- Levine SS, Prietula MJ (2012) How knowledge transfer impacts performance: a multilevel model of benefits and liabilities. *Organ Sci* 23(6):1748–1766
- Levinthal DA (1997) Adaptation on rugged landscapes. *Manage Sci* 43(7):934–950
- Li G, Wang SW, Rosenthal C, Rabitz H (2001) High dimensional model representations generated from low dimensional data samples. I. mp-Cut-HDMR. *J Math Chem* 30(1):1–30
- Ligmann-Zielinska A (2018) "Can you fix it?" Using variance-based sensitivity analysis to reduce the input space of an agent-based model of land use change. *GeoComputational analysis and modeling of regional systems*. Springer, Cham, pp 77–99
- Lin Y (2000) Tensor product space ANOVA models. *Ann Stat* 28(3):734–755
- Lomi A, Harrison JR (2012) The garbage can model of organizational choice: looking forward at forty. *Res Soc Organ* 36:3–17
- Lorscheid I, Heine BO, Meyer M (2012) Opening the black box of simulations: increased transparency and effective communication through the systematic design of experiments. *Comput Math Organ Theory* 18:22–62
- Luo Y, Shah NB, Huang J, Walrand J (2018) Parametric prediction from parametric agents. *Oper Res* 66(2):313–326
- Marino S, Hogue IB, Ray CJ, Kirschner DE (2008) A methodology for performing global uncertainty and sensitivity analysis in systems biology. *J Theor Biol* 254(1):178–196
- Masuch M, LaPotin P (1989) Beyond garbage cans: an AI model of organizational choice. *Adm Sci Q* 34:38–67
- Montgomery DC (2000) *Design and analysis of experiments*, 5th edn. Wiley, Hoboken
- Morris MD (1991) Factorial sampling plans for preliminary computational experiments. *Technometrics* 33(2):161
- Morris MD, Moore LM, McKay MD (2008) Using orthogonal arrays in the sensitivity analysis of computer models. *Technometrics* 50(2):205–215
- Owen AB (2014) Sobol' indices and shapley value. *SIAM/ASA J Uncertain Quantif* 2(1):245–251
- Parry HR, Topping CJ, Kennedy MC, Boatman ND, Murray AWA (2013) A Bayesian sensitivity analysis applied to an agent-based model of bird population response to landscape change. *Environ Model Softw* 45:104–115
- Plischke E, Borgonovo E, Smith CL (2013) Global sensitivity measures from given data. *Eur J Oper Res* 226(3):536–550
- Prietula MJ, Carley KM (1994) Computational organization theory: autonomous agents and emergent behavior. *J Organ Comput* 4(1):41–83. <https://doi.org/10.1080/10919399409540216>
- Prietula MJ, Carley KM, Gasser L (1998) *Simulating organizations: computational models of institutions and groups*. AAAI Press, Menlo Park
- Puranam P, Swamy M (2016) How initial representations shape coupled learning processes. *Organ Sci* 27(2):323–335
- Rahmandad H, Sterman J (2008) Heterogeneity and network structure in the dynamics of diffusion: comparing agent-based and differential equation models. *Manage Sci* 54(5):998–1014
- Restocchi V, McGroarty F, Gerding E, Johnson JEV (2018) It takes all sorts: a heterogeneous agent explanation for prediction market mispricing. *Eur J Oper Res* 270(2):556–569
- Riggs T, Walts A, Perry N, Bickle L, Lynch JN, Myers A, Flynn J, Linderman JJ, Miller MJ, Kirschner DE (2008) A comparison of random vs. chemotaxis-driven contacts of T cells with dendritic cells during repertoire scanning. *J Theor Biol* 250(4):732–751

- Rivkin JW, Siggelkow N (2003) Balancing search and stability: interdependencies among elements of organizational design. *Manage Sci* 49(3):290–311
- Robertson DA (2019) Spatial transmission models: a taxonomy and framework. *Risk Anal* 39(1):225–243
- Saltelli A, Ratto M, Andres T, Campolongo F, Cariboni J, Gatelli D, Saisana M, Tarantola S (2008) *Global sensitivity analysis—the primer*. Wiley, Chichester
- Saltelli A, Aleksankina K, Becker W, Fennell P, Ferretti F, Holst N, Li S, Wu Q (2019) Why so many published sensitivity analyses are false: a systematic review of sensitivity analysis practices. *Environ Model Softw* 114:29–39
- Saltelli A, Bammer G, Bruno I, Charters E, Di Fiore M, Didier E, Espeland WN, Kay J, Lo Piano S, May D, Pielke RJ, Portaluri T, Porter TM, Puy A, Rafols I, Ravetz JR, Reinert E, Sarewitz D, Start PB, Stirling A, van der Sluijs JP, Vineis P (2020) Five ways to ensure that models serve society: a manifesto. *Nature* 582:482–484
- Sauvageau G, Frayret JM (2015) Waste paper procurement optimization: an agent-based simulation approach. *Eur J Oper Res* 242(3):987–998
- Segovia-Juarez JL, Ganguli S, Kirschner D (2004) Identifying control mechanisms of granuloma formation during *M. tuberculosis* infection using an agent-based model. *J Theor Biol* 231(3):357–376
- Simshauser P (2018) Garbage can theory and Australia’s national electricity market: decarbonisation in a hostile policy environment. *Energy Policy* 120:697–713
- Smith EB, Rand W (2018) Simulating macro-level effects from micro-level observations. *Manage Sci* 64(11):5405–5421
- Storlie CB, Swiler LP, Helton JC, Sallaberry CJ (2009) Implementation and evaluation of nonparametric regression procedures for sensitivity analysis of computationally demanding models. *Reliab Eng Syst Saf* 94(11):1735–1763
- Stummer C, Kiesling E, Günther M, Vetschera R (2015) Innovation diffusion of repeat purchase products in a competitive market: An agent-based simulation approach. *Eur J Oper Res* 245(1):157–167
- Takahashi N (1997) A single garbage can model and the degree of anarchy in Japanese firms. *Hum Relat* 50(1):91–108
- Ten Broeke G, Van Voorn G, Ligtenberg A (2016) Which sensitivity analysis method should i use for my agent-based model? *J Artif Soc Soc Simul* 19(1):5
- Thiele JC, Kurth W, Grimm V (2014) Facilitating parameter estimation and sensitivity analysis of agent-based models: a cookbook using NetLogo and ‘R’. *J Artif Soc Soc Simul* 17(3):11
- Troitzsch KG (2012) Team formation in the garbage can. *The garbage can model of organizational choice: looking forward at forty*, vol 36. Emerald Group Publishing Limited, Bingley, pp 229–252
- Utomo DS, Onggo BS, Eldridge S (2018) Applications of agent-based modelling and simulation in the agri-food supply chains. *Eur J Oper Res* 269(3):794–805
- Vandin A, Giachini D, Lamperti F, Chiaromonte F (2021) Automated and distributed statistical analysis of economic agent-based models. [arXiv:2102.05405](https://arxiv.org/abs/2102.05405)
- Wagner HM (1995) Global sensitivity analysis. *Opera Res* 43(6):948–969
- Wall F (2016) Agent-based modeling in managerial science: an illustrative survey and study. *Rev Manage Sci* 10(1):135–193
- Zhao J, Ma T (2016) Optimizing layouts of initial AFV refueling stations targeting different drivers, and experiments with agent-based simulations. *Eur J Oper Res* 249(2):706–716

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Emanuele Borgonovo is Full Professor at the Department of Decision Sciences of Bocconi University, co-editor-in-chief of the *European Journal of Operational Research*, and the current President of the Decision Analysis Society of INFORMS.

Marco Pangallo is a James S. Mc Donnell Foundation Postdoctoral Fellow at the Institute of Economics and EMbeDS at the Sant’Anna School of Advanced Studies.

Jan Rivkin is the C. Roland Christensen Professor of Business Administration at Harvard Business School.

Leonardo Rizzo is PhD candidate at the Department of Network and Data Science of the Central European University with a joint doctoral fellowship in Economics and Network Science and a researcher at the Complexity Science Hub Vienna.

Nicolaj Siggelkow is David M. Knott Professor of Management and Co-Director of the Mack Institute for Innovation Management at the Wharton School of Business of the University of Pennsylvania.

Authors and Affiliations

Emanuele Borgonovo¹  · **Marco Pangallo**² · **Jan Rivkin**³ · **Leonardo Rizzo**⁴ · **Nicolaj Siggelkow**⁵

Marco Pangallo
marco.pangallo@santannapisa.it

Jan Rivkin
jrivkin@hbs.edu

Leonardo Rizzo
leonardo@phd.ceu.edu

Nicolaj Siggelkow
siggelkow@wharton.upenn.edu

- ¹ Department of Decision Sciences and BIDS, Bocconi University, Via Roentgen 1, 20136 Pisa, Italy
- ² Institute of Economics and Department EMbeDS, Sant'Anna School of Advanced Studies, 56127 Milan, Italy
- ³ Harvard Business School, 237 Morgan Hall, Boston, MA 02163, USA
- ⁴ Department of Network and Data Science, Central European University, Quellenstrasse 51, 1100 Wien, Austria
- ⁵ Management Department, The Wharton School, University of Pennsylvania, Philadelphia, PA, USA

Terms and Conditions

Springer Nature journal content, brought to you courtesy of Springer Nature Customer Service Center GmbH (“Springer Nature”).

Springer Nature supports a reasonable amount of sharing of research papers by authors, subscribers and authorised users (“Users”), for small-scale personal, non-commercial use provided that all copyright, trade and service marks and other proprietary notices are maintained. By accessing, sharing, receiving or otherwise using the Springer Nature journal content you agree to these terms of use (“Terms”). For these purposes, Springer Nature considers academic use (by researchers and students) to be non-commercial.

These Terms are supplementary and will apply in addition to any applicable website terms and conditions, a relevant site licence or a personal subscription. These Terms will prevail over any conflict or ambiguity with regards to the relevant terms, a site licence or a personal subscription (to the extent of the conflict or ambiguity only). For Creative Commons-licensed articles, the terms of the Creative Commons license used will apply.

We collect and use personal data to provide access to the Springer Nature journal content. We may also use these personal data internally within ResearchGate and Springer Nature and as agreed share it, in an anonymised way, for purposes of tracking, analysis and reporting. We will not otherwise disclose your personal data outside the ResearchGate or the Springer Nature group of companies unless we have your permission as detailed in the Privacy Policy.

While Users may use the Springer Nature journal content for small scale, personal non-commercial use, it is important to note that Users may not:

1. use such content for the purpose of providing other users with access on a regular or large scale basis or as a means to circumvent access control;
2. use such content where to do so would be considered a criminal or statutory offence in any jurisdiction, or gives rise to civil liability, or is otherwise unlawful;
3. falsely or misleadingly imply or suggest endorsement, approval, sponsorship, or association unless explicitly agreed to by Springer Nature in writing;
4. use bots or other automated methods to access the content or redirect messages
5. override any security feature or exclusionary protocol; or
6. share the content in order to create substitute for Springer Nature products or services or a systematic database of Springer Nature journal content.

In line with the restriction against commercial use, Springer Nature does not permit the creation of a product or service that creates revenue, royalties, rent or income from our content or its inclusion as part of a paid for service or for other commercial gain. Springer Nature journal content cannot be used for inter-library loans and librarians may not upload Springer Nature journal content on a large scale into their, or any other, institutional repository.

These terms of use are reviewed regularly and may be amended at any time. Springer Nature is not obligated to publish any information or content on this website and may remove it or features or functionality at our sole discretion, at any time with or without notice. Springer Nature may revoke this licence to you at any time and remove access to any copies of the Springer Nature journal content which have been saved.

To the fullest extent permitted by law, Springer Nature makes no warranties, representations or guarantees to Users, either express or implied with respect to the Springer nature journal content and all parties disclaim and waive any implied warranties or warranties imposed by law, including merchantability or fitness for any particular purpose.

Please note that these rights do not automatically extend to content, data or other material published by Springer Nature that may be licensed from third parties.

If you would like to use or distribute our Springer Nature journal content to a wider audience or on a regular basis or in any other manner not expressly permitted by these Terms, please contact Springer Nature at

onlineservice@springernature.com