Mind the Assumptions: Quantify Uncertainty and Assess Sensitivity

Emanuele Borgonovo
Department of Decision Sciences, Bocconi University
Milan, Italy
emanuele.borgonovo@unibocconi.it

August 25, 2022

Abstract

This chapter considers how to treat spoken and unspoken assumptions and how to use uncertainty quantification and global sensitivity analysis to make them transparent. We analyze the broad role of assumptions (or hypotheses) in scientific modeling. We investigate how they impact alternative elements of a model. We address models of data (machine learning) and models of phenomena (simulators). We then discuss the impact of varying assumptions on the output of a mathematical model highlighting the role of uncertainty quantification and sensitivity analysis. We single out four main sensitivity analysis goals emerging from the literature.

Keywords: Scientific Modeling; Scientific Hypotheses; Sensitivity Analysis; Models of Data; Models of Phenomena

Assumptions are a fundamental pillar of any scientific investigation. There is no theory or modeling exercise that can live without assumptions. A search of the term "assumption"in the Stanford Encyclopedia of Philosophy does not find a dedicated entry but finds the presence of this term in 839 entries, that comprise "scientific method", "scientific discovery", "scientific logic". Specifically, Assumptions play a fundamental role in scientific modeling and scientific simulations. Then, what are Assumptions? Is there a way in which we can converge on a definition of the term? While this is certainly out of the reach of this chapter, we will try to provide some initial reflections and insights.

According to the Cambridge dictionary, an assumption is *something that you accept as true without question or proof.* In a scientific modeling exercise, an assumption is a statement that establishes the value of a parameter or takes away elements that may not be of concern. Let us refer for a second to the Archimedean model of the lever.

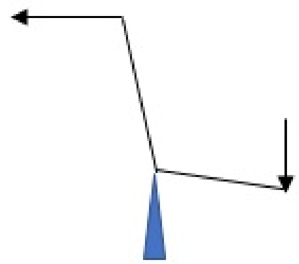

Figure 1: A possible representation of the Archimedean model of the lever. We draw the effort arm, the force arm, the pivot and the two forces at their extremes. The model does not have any reference to the color or the material of the lever, but assumptions strip away irrelevant elements.

Its typical graphical representation is composed by two lines, possibly one triangle and two arrows, representing, respectively, the effort, the load arm, the pivotal point and the forces. This model was first conceived in the Opus On the Equilibrium of Planes by Archimedes and has, since then, been fundamental in statics. The model is, for instance, at the basis of trolleys we carry in airports. However, Archimedes does not consider the type of material, the color, or the brand of the trolley. These aspects have a relevance in marketing the trolley, but not in its physics. The discussion of Frigg and Hartmann (2020) is illuminating in this respect: assumptions help stripping away unessential elements from scientific models.

On the other side, we can regard assumptions as at the set of hypotheses (or statements) that need to hold before we can substantiate the use of a given scientific model. In other words; suppose that in the problem at hand a set of conditions has to be realized simultaneously in order for a certain equation to become applicable; then an equation with a known form would be adequate to simulate the system or phenomenon at hand. Here, the definition of adequacy would need a separate chapter and we refer to Frigg and Hartmann (2020) and Winsberg (2019) for a thorough treatment. We content ourselves with an illustration: Einstein's prediction of the deviation of light is based on the theoretical principles of general relativity. These principles yield the mathematical equation from which to calculate the deviation. The light indeed undergoes the deflection forecasted by the model, as physical experiments indicate. We have all three corners of the scientific triangle: a theory (first vertex) that produces a mathematical model (second vertex) that actually predicts a real world phenomenon.

However, in several scientific investigations, we do not have all three vertices: Scientists build mathematical models that try and reproduce a given phenomenon, often before a scientific theory is developed. In some cases, it will be impossible to achieve a complete scientific theory with axioms from which propositions are derived and that, in turn, yield the equations of the phenomenon at hand. This is likely the case in artificial intelligence investigations where a complete theoretical background about the phenomenon of interest (think of the case in which data come from social networks) is likely to remain out of reach. Under these conditions, we then encounter all those assumptions that a modeler makes in creating a simulation: they range from the selection of the mathematical form of the equations to the value of inputs and parameters.

We are then dealing with the vast class of models represented by computer simulators. Simulators help the scientist in obtaining a numerical portrait of a system or phenomenon of interest.

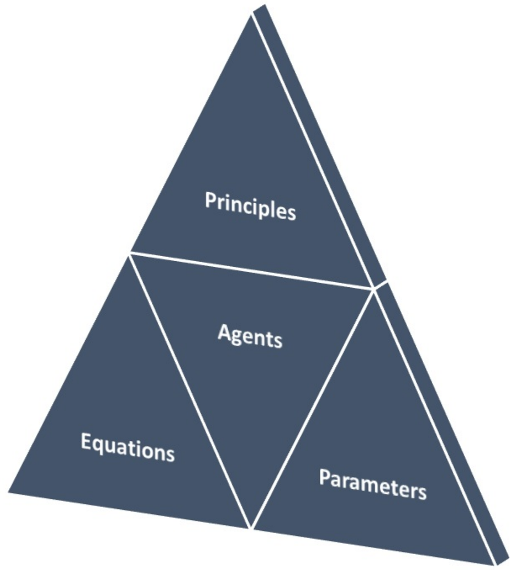

Figure 2: A non-exhaustive set of elements of a scientific modeling problem.

Equations in computer codes aid engineers, physicists, medical doctors with a variety of tasks. As defined in Winsberg (2019), we usually divide simulators into two main categories: equation-based simulators and agent-based simulators. An outstanding example of the former type of model is the simulator of the human hart developed by Alfio Quarteroni and his collaborators in a series of mathematical works (Quarteroni et al., 2017). An example of the second model is the agent based model of the financial market, developed by Le Baron (2006).

Assumptions impact all aspects of scientific modeling. For instance, assumptions are at the basis of the theoretical principles that govern the model development itself. At the same time, they concern the values we attribute to parameters entering the final equations of the model. In a sense, it may well be handy to make a classification of assumptions, depending on the model element to which they refer to. Assumptions can refer to the theoretical foundations of a model. We can call these assumptions *principles*, to give them a stronger connotation (Figure 2). Principles are the conceptual basis of a scientific model or investigation: They comprise axioms and hypothesis of an underlying theory that grounds the model. As noted in Borgonovo et al. (2020), *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.*

Regarding principles, the distinction between models of phenomena and models of data becomes relevant. In models of data, analysts "assume" a form for the input-output mapping, typically selecting a parametric family and then fitting the model through data. This assumption is not necessarily dictated by a theory supporting the model. This is referred to as absence of theory in Begoli et al. (2019), and is one of the main reasons of concern about the reliability of results of machine learning models. Models of phenomena, instead, are developed by researchers following theoretical principles and, if the assumptions concerning the equations being formulated are verified, then these equations will produce an adequate description of the system behavior.

To bring tangible examples, let us consider two representative works, one in the realm of risk assessment, (Apostolakis, 1990), and the other in the business realm . These works deal with models developed to forecast the reliability of complex technological systems. Apostolakis introduces the notion of model of the world (MOW), to denote the system of equations that model the phenomenon of interest — of course, we should read this as the *model of the world of interest* in the specific investigation. To illustrate the notion, Apostolakis uses the Darcy equation for ground-water flow in saturated media:

$$q = -K\frac{\partial h}{\partial x},$$

where q is the specific discharge in the x direction, h is the hydraulic head and K is the hydraulic conductivity. Ford W. Harris introduces the economic order quantity model (Harris, 1913), which is a cornerstone in operations research. Harris's problem is that of determining the optimal lot size in a production system. While his work is written for a business magazine, his approach is strikingly scientific. Harris first states a sharp list of assumptions that clearly pin the reader down to the heart of his problem. For instance, he assumes the *movement is regular* (Harris, 1913, p. 948)(i.e., stationarity). He then formulates an objective function (a total cost) based on the stated elements of the problem, and derives the now famous economic order quantity (EOQ) formula from a minimization:

$$Q = \sqrt{\frac{240MS}{C}}, \tag{1}$$

where Q is the optimal order quantity, S is the setup cost of an order, M is the monthly demand (*the movement of the stock* (Harris, 1913, p. 947)), and C is the unit price of items in the stock. Under Harris's setting, this is the quantity that a rational manager should order. Nonetheless, Harris himself recognized the limitations of the model and warns us about a careful use of the numbers (Cárdenas-Barrón et al., 2014). "*But in deciding on the best size of order, the man responsible should consider all the factors that are mentioned [...] Hence, using the formula as a check, is at least warranted* (Harris, 1913, p. 947)."

Such rule calls for a careful interpretation of any numerical indication obtained by models. This warning matches recent alarms heard in the scientific community. Examples are the ethical guidelines in scientific modeling of the manifesto of Saltelli et al. (2020), or the concerns related to the use (or abuse) of machine learning methods in Rudin (2019) and Begoli et al. (2019).

These concerns call us to the application of methods that can perform a proper uncertainty quantification of the results of the model. Let us take a step back and regard Equation 1 as a black box model in which three factors, namely, C, M and S are fed into in a simulator that produces Q. Clearly, if we fix C, M and S at nominal values C^0, M^0 and S^0 we have an optimal order quantity Q^0.

At level closer to the practice of modeling, besides assumptions about the characteristics of the system, i.e., stationarity or the fact that agents behave rationally, we have assumptions that regard parameters. As discussed in Borgonovo et al. (2020), parameters are *are cardinal quantities that influence the evolution of the model but are determined outside of the simulation run.*

In general, one regards the model as a black box, in which a set of inputs X is fed into a simulator g(X) to produce an output Y of interest (in our case it would be $X = [C, M, S]$) (Figure 3).

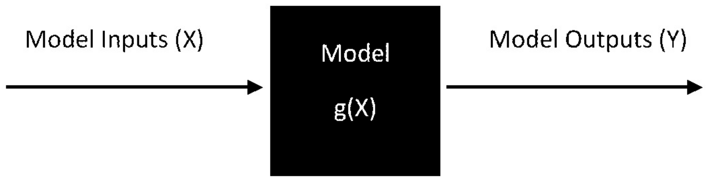

Figure 3: The black-box view of the input-output mapping of a computer model.

Consider now an analyst needing (willing) to determine the EOQ, and whose available information allows her to set the following values for the inputs: $M^0 = 4500$, $S^0 = 3.5$, $C^0 = 75$. The corresponding EOQ calculated by the model is $Q = 224.5$ units. Note that if we were to be commuting this quantity to management, we would be exposed to the risk of a-critically accept the results of the model. Reporting a point value in a context full of uncertainties is evidently a poor scientific practice (see the recent manifesto of Saltelli et al. (2020)). The Nobel prize for economics Paul Samuelson, in his work that established the foundations of sensitivity analysis in Economics, writes: "If no more than this could be said, the economist would be truly vulnerable to the gibe that he is only a parrot taught to say "supply and demand" (Samuelson, 1947, p. 97)."The risk is that of undermining the whole efforts of the modeling exercise.

In this context, we are interested in the variation of the model output that follow variations in the inputs. This task comprises two important subtasks: sensitivity analysis and uncertainty quantification.

Regarding uncertainty quantification, we refer to the monographs of Saltelli et al. (2008), Sullivan (2015), Borgonovo (2017), as a complete treatment is out of our reach in this chapter. However, let us briefly review the notion and procedure. The purpose of an uncertainty quantification is to assess the variability of the output of a model. If the model output is numerical, this variability can be expressed in the form of a variation range and, if possible, of a probability distribution of the model output over this range. These quantities are obtained by propagating the uncertainty in the model inputs through the model input output-mapping. To illustrate, the analyst assigns plausible ranges to the inputs and, if possible, a corresponding probability distribution. Uncertainty is then propagated through the model via a Monte Carlo simulation[1]. To illustrate, suppose that the scientist dealing with the Harris EOQ model sets the following ranges: $M \in [2000, 7000]$, $S \in [2, 5]$, $C \in [50, 100]$ and assigns uniform distributions to the inputs over these ranges. Monte Carlo propagation would lead to the distribution of Q whose density is reported in Graph a) of Figure 4.

[1]For a characterization of Monte Carlo uncertainty propagation, please refer to Metropolis (1987) for a historical account, to Glasserman (2003) for a technical treatment, and to Winsberg (2019) for the epistemological viewpoint.

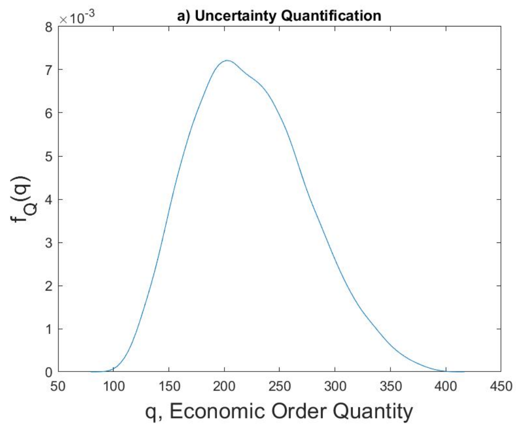

Figure 4: Model output density, as a possible visualization of the outcome of the uncertainty quantification for our EOQ model in Equation 1.

Graph a) Figure 4 shows that the EOQ varies between 100 and 400 units. It is then up to the analyst or the manager to understand whether such variation is reasonable enough to proceed with the order, or if new analysis needs to be carried out to reduce the variability.

In general, the derivation of insights from sensitivity analysis can be made formal by the notion of sensitivity analysis setting. A setting is *a way of framing the sensitivity quest in such a way that the answer can be confidently entrusted to a well-identified measure*(Saltelli et al., 2008, p. 24). In the case of Figure 4, we the setting a *stability setting.* This setting indeed comprises all those situations in which analysts wish to determine whether the robustness of the indications they are providing based on the numerical output of the model. For instance, consider an analyst dealing with a linear programming problem. Here, an important information to the modeler is the range of variation in the model inputs over which the optimal plan remains the same. A specific methods that helps the analyst answering this question, besides the uncertainty propagation we have seen, is Wendell's tolerance sensitivity (Wendell, 1985). [2] Similarly, in a decision analysis problem expressed in the form of a decision-tree or an influence diagram or even in a generic optimization problem, we find several techniques for assessing stability. One in particular is value of information Oakley (2009). If the value of information of a certain input is null, then the analyst can conclude that the optimal plan is insensitive to variations in that input.

A second task with which analysts are (or should) be frequently concerned is the answer to the question: how does the response of the model depend on each input marginally? Formally, we could call this setting *marginal behavior determination.* The machine learning as well as the simulation literature have made available several tools to answer this question, especially in association with the urge to increase interpretability of machine findings. Much attention and several advancements have been made in partial dependence indicators, and we can point the reader towards the works of Friedman et al. (2001), Goldstein et al. (2015) and Apley and Zhu (2020), for further details.

In this chapter, we content ourselves with a visual interpretation. Consider Figure 5.

[2] Tolerance sensitivity is an approach tailored to linear programs. By exploiting the geometric properties of linear optimization, tolerance sensitivity determines the range of simultaneous variation in the inputs such that the optimal solution of the linear program remains stable. The inputs are the objective function coefficients, the right hand sides or the elements of the coefficient matrix.

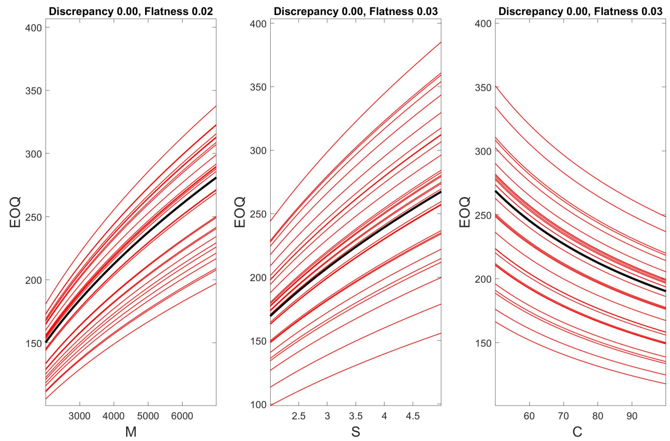

Figure 5: Comparative statics with individual conditional expectation plots (Goldstein et al., 2015). The black line displays the average marginal behavior of the EOQ as a function of each of the three inputs.

The graphs in this Figure 5 immediately suggest us that the EOQ is increasing in M and S and decreasing in C. The analyst can then infer answers to questions such as: Is this behavior consistent with intuition? What should we suggest to a decision maker based on this behavior? Answering the first question also provides a way to enter into a debugging mode. If an underlying theory suggests that a model response should be decreasing in a given input and the sensitivity analysis suggests the opposite, we need further investigation. We might either be at the verge of a scientific discovery in which a common (mis)-conception is falsified or, more simply, there might be a bug, with a minus placed instead of a plus at some point of the (usually complex) code.

The tool displayed in Figure 5 is one of the possible choices analysts have to determine marginal behavior. Analysts can count on local approaches based on differentiation (Griewank, 2000), on the sign of Newton ratios (Rakovec et al., 2014), or on one-way sensitivity functions van der Gaag et al. (2007). Global approaches comprise methods such as partial dependence functions Friedman et al. (2001) and individual conditional expectation plots Goldstein et al. (2015), as well as the recently introduced Accumulated Local Effect plots Apley and Zhu (2020). A review and critical analysis of these methods from a decision-making viewpoint is recently given in Borgonovo et al. (2021b).

An insight which is consider also crucial in machine learning for interpretability is input importance. For analysts it is relevant to know the factors that drive a model response. Knowing for instance that a given input is extremely relevant for the prediction of an algorithm would help analysts understand that the algorithm might be easily biased. If the input is discriminatory in nature, then the algorithm predictions themselves would become questionable from an ethical standpoint (see the discussion in Rudin (2019)). Understanding feature importance goes under the setting of *factor prioritization* (Saltelli, 2002; Saltelli et al., 2008). Analysts can find several methods that accomplish this in the literature. On the one hand we have local methods such as elasticity or the differential important measure Borgonovo and Apostolakis (2001), in which the indications of partial derivatives are synthesized in dimensionless indicators (indeed, partial derivatives themselves may not be directly comparable because they carry the units on which the input and output are denominated and inputs may have different units). Or, in the well known representation of Tornado Diagrams Howard (1988), one plots the effects of one-at-a-time variations in the inputs. These methods are local. Alternatively, one can use global methods, with the families of regression-based methods (Kleijnen and Helton, 1999), variance-based methods (Saltelli and Tarantola, 2002), moment-independent methods (Borgonovo and Apostolakis, 2001; Da Veiga, 2015) and Shapley Values (Owen, 2014). The literature here is vast and we refer the reader to works such as Razavi et al. (2021), Saltelli et al. (2008), Borgonovo (2017), and the handbook of Ghanem et al. (2017) for further details. To illustrate the insights, let us present a graphical result obtained for our EOQ model.

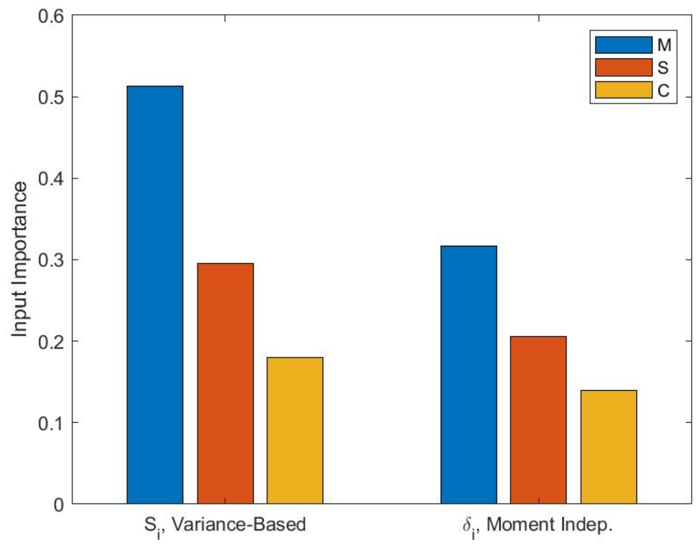

Figure 6: Input Importance measured with variance-based (left bars) and moment-independent sensitivity measures (right bars).

The first triplet of bars in Figure 6 represents the relative importance of the three inputs of our EOQ example computed using variance-based importance measures. The second triplet displays the relative importance using a moment-independent sensitivity measure (the δ-importance — see Borgonovo and Apostolakis (2001)). Both indices concur in evidencing M as the most important input, followed by S and C.[3]

The last setting we address is *model structure determination*. In this setting, the analyst wishes to understand whether the variations of the output are the sum of the variations induced by the individual inputs. If this the case, then the model is said to be additive; conversely, interaction effects are important. The study of interactions is a huge chapter in modelling. Often, interactions are studied in association with causality, as detailed in the book of Vanderweele (2015). Relevant

[3] For the technically inclined reader, we have computed these sensitivity measures from the same input-output dataset used for uncertainty quantification in Figure 4. The technique is called given-data estimation and is documented in works such as Strong et al. (2012) and Plischke et al. (2013).

to the determination of interactions is the scale at which the analyst is looking at. We may consider interactions at a global scale or at the local or at the infinitesimal scale. To illustrate, let us carry out an analysis of interactions at the global scale for our problem. It can be seen that the difference between 1 and the sum of variance-based sensitivity measures is an indicator of global interactions (Owen, 2003; Saltelli, 1999). In our example, this difference is equal to 0.01, signalling that the model behavior is additive over the range of interest.

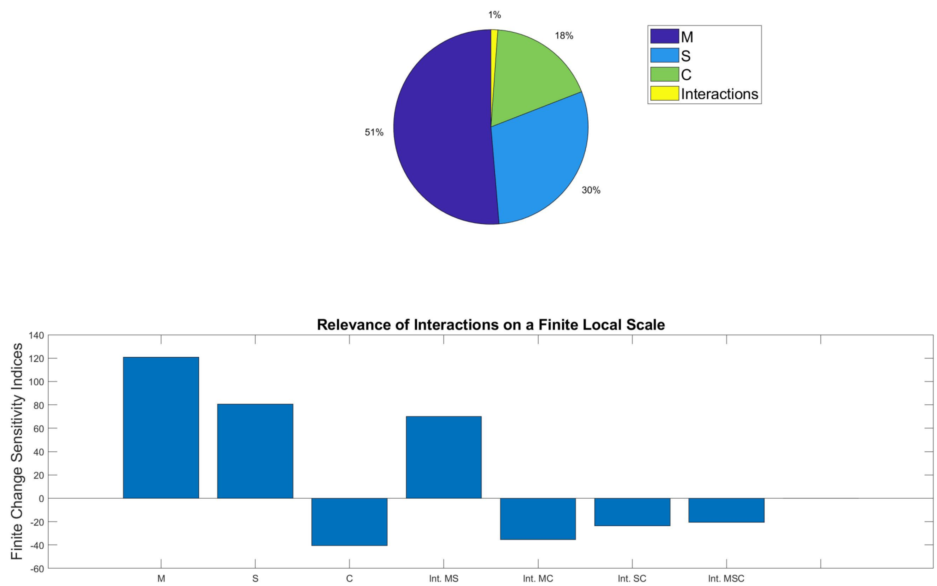

Figure 7: Upper graph: Interactions on a global scale for our EOQ model running example. Lower graph: Interaction on a local scale for the same model.

However, consider carrying a local analysis of the response of the EOQ model when the inputs are between their extremes. That is, we consider the difference between the EOQ computed when the inputs are fixed at [2000,2,50] and at [700,5,100]. Such difference is equal to 151.26. The literature has shown that this change can be exactly apportioned to the individual effects of M, S and C, to their pairwise interactions and to the remaining triple interaction. The result is displayed in the lower graph of Figure 1. This graph shows that interactions matter in explaining the EOQ variation, when the inputs jump from the lower to their upper extreme of their ranges. In particular, we record positive (the interaction between M and S) as well as negative interactions (the ones between M and C, S and C and M, S and C). The main message here is that the relevance of the scale at which we carry out a model structure analysis. insights on a local or global scale might be different. Space limitations do now allow us to enter into the technical aspects of these results, but would point the reader to the recent work of Borgonovo et al. (2021a).

We would like to close with some further reflections concerning the choice of the assumptions inspired by the critical review in Pfleiderer (2020). (Hornberger and Spear, 1981, p 8) state *... most* *simulation models will be complex with many parameters, state variables and non-linear relations. Under the best circumstances, such models have many degrees of freedom and, with judicious fiddling, can be made to produce virtually any desired behaviour, often with both plausible structure and parameter values.* In his essay, Pfleiderer (2020) observes that scientists might cherry pick assumptions to create models that produce a result they (consciously or unconsciously) aim at. Indeed, in principle, the right set of assumptions will lead to a model of the form that produces a pre-determined outcome Pfleiderer (2020) actually illustrates several examples and proposes the use of filters on assumptions to prevent this from occurring. Pfleiderer (2020) underlines also the ethical implications of such choices. Similarly, it is possible to manipulate inputs so that a given model produces a desirable numerical value. In this chapter, we have made an introductory effort to illustrate that carefully dealing with assumptions, and thoroughly performing sensitivity analysis should help analysts preventing this type of pitfalls. The analysis must be as transparent as possible, even at the peril of displaying contradictory behavior from the same model. However, this would allow stakeholders to make a fully informed decision about whether to retain or discard a model's results and predictions.

References

Apley, D W, J Zhu. 2020. Visualizing the effects of predictor variables in black box supervised learning models. *Journal of the Royal Statistical Society. Series B: Statistical Methodology* **82**(4) 1059–1086.

Apostolakis, G E. 1990. The concept of probability in safety assessments of technological systems. *Science* **250**(4986) 1359–1364.

Begoli, E, T Bhattacharya, D Kusnezov. 2019. The need for uncertainty quantification in machine-assisted medical decision making. *Nature Machine Intelligence* **1**(1) 20–23.

Borgonovo, E. 2017. *Sensitivity Analysis: An Introduction for the Management Scientist*. I ed. International Series in Management Science and Operations Research, Springer New York.

Borgonovo, E, G E Apostolakis. 2001. A new importance measure for risk-informed decision making. *Reliability Engineering & System Safety* **72**(2) 193–212.

Borgonovo, E, M Pangallo, J Rivkin, L Rizzo, N Siggelkow. 2020. Sensitivity Analysis in Agent-Based Modeling. *Work In Progress* 1–58.

Borgonovo, E., E. Plischke, G. Rabitti. 2021a. Interactions and Computer Experiments. *Scandinavian Journal of Statistics* 1–30doi:10.1111/sjos.12560.

Borgonovo, Emanuele, Manel Baucell, Elmar Plischke, John Barr, Herschel Rabitz. 2021b. Trend Analysis in the Age of Machine Learning. *SSRN* 1–24.

Cárdenas-Barrón, Leopoldo Eduardo, Kun-Jen Chung, Gerardo Treviño-Garza. 2014. Editorial: {C}elebrating a century of the economic order quantity model in honor of {F}ord {W}hitman {H}arris. *International Journal of Production Economics* **155** 1–7.

Da Veiga, S. 2015. Global Sensitivity Analysis with Dependence Measures. *Journal of Statistical Computation and Simulation* **85**(7) 1283–1305.

Friedman, Jerome, Trevor Hastie, Robert Tibshirani, Others. 2001. *The elements of statistical learning*, vol. 1. Springer series in statistics New York.

Frigg, Roman, Stephan Hartmann. 2020. Models in Science. Edward N Zalta, ed., *The Stanford Encyclopedia of Philosophy*, spring 202 ed. Metaphysics Research Lab, Stanford University, 1–10.

Ghanem, Roger, David Higdon, Houman Owhadi, eds. 2017. *Handbook of Uncertainty Quantification*. Springer Verlag.

Glasserman, Paul. 2003. *Monte Carlo Methods in Financial Engineering*. Springer Verlag, New York, NY, USA.

Goldstein, A, A Kapelner, J Bleich, E Pitkin. 2015. Peeking Inside the Black Box: Visualizing Statistical Learning With Plots of Individual Conditional Expectation. *Journal of Computational and Graphical Statistics* **24**(1) 44–65.

Griewank, A. 2000. *Evaluating Derivatives, Principles and Techniques of Algorithmic Differentiation*, *Frontiers in Appl. Math.*, vol. 19. SIAM, Philadelphia.

Harris, F W. 1913. How Many Parts to Make at Once. *Factory, The Magazine of Management* **10**(2) 135–136.

Hornberger, G. M., R. C. Spear. 1981. An approach to the preliminary analysis of environmental systems. *Journal of Environmental Management* **12** 7–18.

Howard, R A. 1988. Decision Analysis: Practice and Promise. *Management Science* **34**(6) 679–695.

Kleijnen, J.P.C., J.C. Helton. 1999. Statistical analyses of scatterplots to identify important factors in large-scale simulations: 1: Review and comparison of techniques. *Reliability Engineering & System Safety* **65** 147–185.

Le Baron, B. 2006. Agent-Based Computational Finance. *Handbook of Computational Economics* **2** 1187–1233.

Metropolis, N. 1987. The beginning of the Monte Carlo Method. *Los Alamos Science* 125–130.

Oakley, J E. 2009. Decision-theoretic Sensitivity Analysis for Complex Computer Models. *Technometrics* **51**(2) 121–129.

Owen, A B. 2003. The Dimension Distribution and Quadrature Test Functions. *Statistica Sinica* **13** 1–17.

Owen, Art B. 2014. {S}obol' Indices and {S}hapley values. *SIAM/ASA Journal on Uncertainty Quantification* **2**(1) 245–251.

Pfleiderer, Paul. 2020. Chameleons: The Misuse of Theoretical Models in Finance and Economics. *Economica* **87**(345) 81–107.

Plischke, E, E Borgonovo, C L Smith. 2013. Global Sensitivity Measures from Given Data. *European Journal of Operational Research* **226**(3) 536–550.

Quarteroni, A, A Manzoni, C Vergara. 2017. The Cardiovascular System: Mathematical Modelling, Numerical Algorithms and Clinical Applications. *Acta Numerica* **26** 365–590.

Rakovec, O, M C Hill, M P Clark, A H Weerts, A J Teuling, R Uijlenhoet. 2014. Distributed evaluation of local sensitivity analysis (DELSA), with application to hydrologic models. *Water Resources Research* **50**(1) 409–426.

Razavi, S, A Jakeman, A Saltelli, C Prieur, B Iooss, E Borgonovo, E Plischke, S Lo Piano, T Iwanaga, W Becker, S Tarantola, J H A Guillaume, J Jakeman, H Gupta, N Melillo, G Rabitti, V Chabridon, Q Duan, X Sun, S Smith, R Sheikholeslami, N Hosseini, M Asadzadeh, A Puy, S Kucherenko, H R Maier. 2021. The Future of Sensitivity Analysis: An essential discipline for systems modeling and policy support. *Environmental Modelling and Software* **137**.

Rudin, C. 2019. Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *Nature Machine Intelligence* **1**(5) 206–215.

Saltelli, A. 1999. Sensitivity analysis. Could better methods be used? *Journal of Geophysical Research* **104(D3)**(37) 89–93.

Saltelli, A. 2002. Sensitivity Analysis for Importance Assessment. *Risk Analysis* **22**(3) 579–590.

Saltelli, A, G Bammer, I Bruno, E Charters, M Di Fiore, E Didier, W N Espeland, J Kay, S Lo Piano, D May, R Jr. Pielke, T Portaluri, T M Porter, A Puy, I Rafols, J R Ravetz, E Reinert, D Sarewitz, P B Start, A Stirling, J P van der Sluijs, P Vineis. 2020. Five Ways to Ensure that Models Serve Society: a Manifesto. *Nature* **582** 482–484.

Saltelli, A, M Ratto, T Andres, F Campolongo, J Cariboni, D Gatelli, M Saisana, S Tarantola. 2008. *Global Sensitivity Analysis – The Primer*. Chichester.

Saltelli, A, S Tarantola. 2002. On the Relative Importance of Input Factors in Mathematical Models: Safety Assessment for Nuclear Waste Disposal. *Journal of the American Statistical Association* **97**(459) 702–709.

Samuelson, P. 1947. *Foundations of Economic Analysis*. Harvard University Press, Cambridge, MA.

Strong, M, J E Oakley, J Chilcott. 2012. Managing Structural Uncertainty in Health Economic Decision Models: a Discrepancy Approach. *Journal of the Royal Statistical Society, Series C* **61**(1) 25–45.

Sullivan, T J. 2015. *Introduction to Uncertainty Quantification*. Springer Verlag.

van der Gaag, L C, S Renooij, V M H Coupe. 2007. Sensitivity Analysis of Probabilistic Networks. *Advances in Probabilistic Graphical Models, Studies in Fuzziness and Soft Computing Volume 214, 2007, pp 103-124* **214** 103–124.

Vanderweele, T J. 2015. *Explanation in Causal Inference*. Oxford University Press, New York, NY.

Wendell, R E. 1985. The Tolerance Approach to Sensitivity Analysis in Linear Programming. *Management Science* **31**(5) 564–578.

Winsberg, Eric. 2019. Computer Simulations in Science. Edward N Zalta, ed., *The {Stanford} Encyclopedia of Philosophy*, {w}inter 2 ed. Metaphysics Research Lab, Stanford University.