

# Models, Values, and Disasters

Michael H. Veatch (Gordon College)



Mike Veatch holds a B.A. in mathematics from Whitman College, an M.S. in Operations Research and Statistics from Rensselaer Polytechnic Institute, and a Ph.D. in Operations Research from Massachusetts Institute of Technology. Prior to his position at Gordon College he spent six years at The Analytic Sciences Corporation studying Air Force supply systems. His research interests include stochastic network control problems and humanitarian logistics.

## Abstract

Decision-support models have values embedded in them and are subjective to varying degrees. Philosophical and ethical perspectives on operations research models are used to describe this subjectivity. Approaches to model building are then suggested that take into account subjectivity and values. For the decisions to reflect the right values, the model must align with the decision-maker's values. I argue that it is appropriate and important for Christians applying mathematical models to be keenly aware of decision-maker's values and seek to reflect them in the model. Disaster response planning is presented as an example where incorporating values is challenging. The responding organizations have multifaceted goals. How is equity balanced with efficiency? How is cost and donor interest considered? I report on a study of how Christian relief organizations differ from non-faith based organizations in ways that can be reflected in their logistics procedures and in these models.

## 1 Introduction

This paper came about because, after 37 years of using mathematical models and reflection on modeling, I am still struck by how challenging it is to explain and justify the assumptions in a model, especially *prescriptive* models. My interest resurfaced during a recent study of how Christian relief organizations make operational decisions. Modeling in this context is particularly challenging and subjective. The study has raised more questions than it has provided answers. This paper is an attempt to lay out these questions as they relate to modeling more generally and suggest some pragmatic approaches to dealing with subjectivity in modeling. It is my hope that acknowledging and thinking through the subjective side of modeling will be useful for those trained in the mathematical sciences, where models are generally not viewed as subjective.

The unsuspecting undergraduate who starts modeling things beyond “hard” scientific phenomena steps into a long debate about the nature and quality of models. At Gordon College, a core science course includes thinking about the limitations and possible biases in science. When using models to support business or organizational decisions, many other issues about the use of models present themselves. The questions center around how well a model of a specific situation can describe it and the implications for the decision maker or the stakeholders in a decision.

Good models require data and an understanding of the dynamics or process being modeled. However, decision-support models also have values embedded in them. For the decisions to reflect the right values, the model must align with the decision-maker's values. This rather obvious statement raises several questions for the model builder, or analyst. How can the analyst know the values of

the decision maker? What if the analyst's values differ notably from the decision-maker's? When is it appropriate for the Christian (or other) faith of the decision-maker, which frames her overall values, to influence the model used? The author's position is that Christians applying mathematical models to support decisions should be keenly aware of decision-maker's values and seek to reflect them in the model. This principle is usually germane even when the decision-maker does not share the analyst's Christianpresuppositions.

This question about the decision-maker's and the analyst's values is explored in the specific context of decision-support models for humanitarian response after a disaster. The decisions that will be considered are about logistics: the transportation and distribution of goods. The responding organizations have multifaceted goals rather than an overarching profit objective. Because each disaster is unique and requires a rapid response, there are limitations on the information available to drive the model. Christian disaster response organizations (DROs) play a large and increasing role in global disaster response. Thus, it is natural to ask how their faith values influence their decisions and whether these values can be incorporated in mathematical models.

Rather than dip into the broad philosophical literature about subjectivity in science and social science, I will focus on the small, but noteworthy, literature on decision-support models in the operations research community. First, some ethical guidelines in modeling are reviewed and the role of values is considered. Section 3 describes the philosophical problem of justifying operations research models and Section 4 suggests some factors that affect the ability to construct useful models. Section 5 describes ways to incorporate values in the models used in disaster response. The author's study of Christian relief organizations and their values is described in Section 6. The final section suggests how students could be introduced to values in models.

## 2 Trusting modelers: ethics and values

The ethical pitfalls of applying operations research has received some attention in the field's journals. A lively earlier discussion of ethics in modeling is the edited volume [20]. Another early voice in this area was Saul Gass [7]. A survey of the literature on ethics is given in [19] and a report on the operations research community's stand on ethics is in [4]. Much of the literature and dialogue on this topic has been centered in Europe, including [6] and [14]. However, ethics is not emphasized as much in operations research as in professions such as engineering or business. Graduate students in operations research, business analytics, or related majors generally do not take courses in ethics. Business ethics would be the closest course offering, and is more likely to be taken by MBA students. There is no broadly endorsed code of ethics in the United States, suggesting a lack of agreement about what such a code would contain. There was a voluntary code proposed in Europe starting in 2000 [3], focusing on respect for the people affected by business decisions and environmental concerns.

When analysts develop a model, they certainly share the ethical responsibilities that come with the application area and decisions that are being made, and scientific standards of honesty and integrity. They also have a unique role because they are proposing a model, which they have chosen or developed, to a decision-maker or to stakeholders whose values may differ from the analyst's. Others may not understand the model, either because of the mathematics or not having access to the information that the analyst does. As discussed in the next section, there are limitations to the objectivity of models, leaving more room for the analyst to shape the results. The issue of trusting the model builder looms larger when the model is more subjective or opaque. Is the analyst simply

being trusted to be honest in creating an objective, scientific model or is the process permeated with the knowledge, experience, and biases of the analyst?

One common pitfall discussed in this literature is that the analyst may place too much trust in a model. What we can expect from an operations research model depends greatly on the system being modeled and available data. The gold standard of validating a model against exogenous data is not possible in many settings. Another dilemma is pressure from a client to produce a model or to reach a certain conclusion, biasing the results. To what degree should the analyst represent the interests of the client? Put another way, is it appropriate to use a model in a certain situation for advocacy? Concerns about subjectivity and bias presuppose that there is an objective reality that the analyst should seek to capture in the model and the data that drives it.

### 3 Trusting models: epistemology in operations research

John Mingers explains the various philosophical views of operations research models in [16] and [15]. The founders of operations research, with backgrounds in science and engineering, embraced realism and empiricism. They saw operations research models as subject to testing and refinement, broadly applicable, and capturing objective shared experience. These views appear in early textbooks, such as [1]. However, the systems being modeled include or are designed by humans, so the field has aspects of a social science. The positivist view of social science asserts that knowledge in these fields is empirical and can be objective, or theory-free. Realism leads to an emphasis on established models or problems, and a tendency to teach models, and how to justify them, rather than teaching a discovery process of modeling [17]. A course might present the Verhulst logistics growth model for a population, or a linear programming approach to choosing an investment portfolio, then ask students to apply this type of model to a situation. Empiricism emphasizes that models are based on data and observation. Statistical and simulation models exemplify this approach.

In contrast to empiricism, a pragmatic or instrumentalist philosophy views models as useful for prediction, but not as containing truth or explanatory power. The “system” is in the model, not the real world. Proponents of this view see operation research as technology, not science—which is still viewed as realist.

A third view, prevalent in the philosophy of science, is that theories are socially constructed. In the social sciences, theories—or paradigms—tend to coexist. Rather than one theory making predictions that could falsify the other, they may be incommensurable, essentially talking past each other. Mingers sees social construction as explaining the alternative methods that have flourished in operations research. He lists

- traditional optimization models
- cognitive mapping of qualitative models in decision making
- soft systems methodology [5], which is strongly relativist and constructs models based on the decision-maker’s views
- critical operations research, advocated by Mingers based on a philosophy of critical realism
- decision analysis [12], which recognizes multiple criteria and assesses a decision-makers values or utility function.

Observing the models used by individuals with different training within operations research, it is difficult to avoid the view that the models employed in a specific situation are partly chosen because of the person's background. To a degree, models are chosen by modelers, not derived from the reality being modeled.

The categories of models described above, and the focus of this paper, excludes the many “models” that are essentially statistical. While all of these models use data, a statistical model makes few structural assumptions and has no dynamics. Rather, its goal is to generalize, or draw inferences, from data. The most common example is a regression model, which might assume linear relationships between certain variables, but even these assumptions are usually tested. The model fits the data and seeks to predict outputs based on inputs, but offers little or no additional structure. The use of statistical models is growing rapidly. This is largely due to the availability of data and better algorithms from machine learning to draw inferences from large amounts of data. Machine learning and data science blur the lines somewhat between statistical and more structural models. Techniques like clustering and regularized regression (ridge or LASSO regression) can be considered “just” statistical, but they also identify patterns out of so many possibilities that one could view them as finding structure. Either way, data-driven projects that involve machine learning often include a more structural or dynamic model, to which the comments in this paper would apply.

Another growing use of models is when organizations classify and make decisions about individuals. The models tend to be of the machine learning type: large and data-intensive. For example, the credit scores used by companies to decide what loans and mortgages to offer to individuals are based on a complex model [18]. These data-driven systems are receiving increasing scrutiny because of their opaqueness, potential for bias or discrimination, and unintended consequences. While these issues are very important, this paper focuses on models used for organizational-level decisions and control systems.

## 4 Accuracy of models

Operations research models have continually been used in new domains and on more complex systems. However, for an application to be successful, some conditions must be met. Here is a partial list.

- The system must obey mathematical, or at least statistical, laws (law-like).
- The laws must be discernible. Encrypted messages, for example, contain information and conform to the syntactical laws of language. But this information cannot be extracted or the laws observed because the encryption precludes—as a matter of computing limitations but not in principle—detecting the regularities in the data.
- Data for the particular instance must be available.
- The phenomenon should be repeatable, either in a deterministic or statistical sense. This is a standard criterion in science. For knowledge to be scientific it must be public. In experimental science, this standard is maintained by requiring that other scientists be able to reproduce a result. In other sciences, the methodology (including models) and the existing data are subjected to review.

For models of systems that include people in certain roles, intentionality—actions that cannot seem to be described and predicted by scientific laws—creates limitations to modeling. These limits are widely recognized in the social sciences, where different theories of human agency place different limitations on the ability, in principle, to predict human behavior and its impact on models of human institutions and systems.

Some systems cannot be modeled without aggregating because of their sheer size. However, it is often possible to model aggregate behavior in these systems. One very successful example of modeling aggregate behavior is economic theory, where individual people's actions are not predicted but the relationships between market variables can be accurately modeled. Similarly, demographic and epidemiological models aggregate individuals into large classes, or cohorts, and apply a single rate parameter to the cohort to predict future population or the spread of a disease.

The ability to construct a predictive model of a system depends critically on its stability and related properties. Consider three examples.

1. An aircraft carrier is preparing to launch its full complement of aircraft. How long will it take until all (or almost all, since aircraft are not as reliable as the carrier) are launched?
2. A factory assembles many different, but closely related complex pieces of electronic equipment. They have plenty of business, so their goal is to produce as much as possible. How many items of each type will they produce in the next month?
3. A large group of people is gathered on the grounds of a rural home when a fire breaks out in the house. No help is nearby, but there is a well on the property and they attempt to put out the fire. How long does it take them to put out the fire?

These three examples have roughly comparable size, or complexity, as measured by the number of interacting people and objects that come to mind as we think about them. The fire brigade lacks complex equipment and so might be considered somewhat smaller. Yet even if we possessed all the data and knowledge about these systems that we can imagine, our ability to construct models of them would differ dramatically.

Carrier operations are a remarkable example of a very tightly designed and centrally controlled human/machine system. Although its individual elements are fraught with randomness and even some creativity, the overall behavior—the launching of aircraft—is very predictable, even though the system operates very nearly at capacity. Even the uncertainty in the model's prediction could be well captured by probability models of equipment reliability and human performance.

The fire brigade is at the opposite extreme. It has no predefined plan or control; it lacks design. Even given detailed knowledge of the house, the surroundings, the people present, how the fire started, and the weather, we would be at a loss as to how to construct a meaningful model. Given a large database of past fire brigades (the existence of which is problematic because the point is that the situation is unusual), we might find some correlations. Or we might establish bounds on water volume delivery and its efficacy on certain sizes of fires. But we would still be left with a huge uncertainty: Will the fire be brought under control while it is small or will it consume the house? This example differs not so much in the randomness exhibited by its components—both systems have that characteristic—but by its lack of an operational plan. We have no idea how the people will organize and how the system will operate.

The factory is in some sense in the middle. It is designed, with a fair degree of control over the actions of people, but it is controlled in a less centralized fashion than the carrier. Various people in the factory make decisions that affect scheduling, routing, and tempo. External circumstances and the response to them can easily alter the schedule. For example, a customer with a “hot” order may have their job expedited, resulting in other jobs (and total production) being delayed. A sole-source supplier may fail to deliver a part on time, making the production of certain items impossible. Models of the factory will be very helpful in understanding it and exploring sensitivities, but they are likely to miss some of the innovative behavior it exhibits. Predictions will contain randomness, but have the potential to be rather accurate in a statistical sense.

In summary, human agency without a clear purpose to constrain individual actions limits the ability to model. Such constraints occur in a system with a high degree of design and control. Repeatable situations are most easily modeled. Since systems with humans may evolve and learn, they can exhibit more change, so that a model no longer applies.

If accurate models are possible in a certain application, then as work continues on the models and more data is collected to drive them, they should perform better, converging on an accurate description of the system. Is this a realistic expectation? There are cases where accuracy has greatly improved over time and others where it has not. One problem that was the subject of intense study, particularly from about 1990 to 2010, is how to staff an inbound call center. Models in this area try to predict the proportion of calls made to customer service that will have to wait more than a specified time. The firm operating the call center may have contracted to provide this level of service to the company whose customers they are, or may be concerned about losing their own customers. Once the service level can be predicted from the staffing level, an optimal staffing level can be found. Early models were very inaccurate at predicting service level.

The first major improvement was to replace naive assumptions about the shape of the service time distribution with a shape fit to large amounts of data, which has a “heavier” tail. Then the same was done for customer impatience (how long they will wait on hold). A third major achievement was to develop good approximate models of the system dynamics, called many-server queues. The models study the limit as the number of servers (and calls) approaches infinity. The approximation is quite accurate and the resulting model can be studied and optimized more easily than simulation models. Forecasting call volume also improved. Today, these models are being used for staffing, resulting in steadier service levels and huge cost savings. This example has many of the characteristics described above: known system dynamics (laws), observability, voluminous data, and enough size to aggregate unpredictable human behavior.

## 5 Values in models

Next we describe how values appear in decision support models. These models are prescriptive or normative; they compare options, leading to a recommended decision or a ranking of alternatives. When there are many possible decisions, the general framework of constrained optimization is used. In this framework, there is a set of decision variables, constraints limiting the possible choices, and an objective function. Values may shape the choice of decision variables through the assumption of what things are subject to change and what is already fixed. Some constraints reflect logical requirements, operating procedures, or physical capacities; these are not likely to be influenced by values, once the choice of decision variables is made. Other constraints are goals set by the decision-maker, based on their values or need to accommodate other stakeholders. For example, an

airline scheduling flight crews may want the constraint that no one is away from home more than two weekends in a month. See [2] for a history of optimization models in the airline industry.

Objective functions are most clearly value-laden. Some objectives are singular, such as maximizing profit. In other situations, there are many competing objectives. Although multi-criteria decision making and goal programming can be used to find non-dominated decisions, conceptually we can assign weights to each criteria, creating a single utility function. Utility theory, expert systems, and more recently supervised learning all have methods to elicit the preferences from decision-makers. Still, often the desired objective cannot be measured or modeled, so surrogate objective functions are used.

The analyst needs to consider how to learn the decision-maker's values and the extent to which they will include them in the model. The client can also shape the analysis in many ways, some to be expected and some perhaps surprising. I offer an anecdote. In the 1990s I developed an inspection model for a large manufacturer of printers. Some batches of parts from suppliers were inspected by sampling a few parts. If defective parts were not detected until they were in subassemblies, or until the finished printer was tested, the cost of the defect was much larger. They wanted to better choose which parts to inspect before assembly to reduce their quality-related costs. They provided detailed data about their inspection procedures, costs involved, and the defects found over a two year period. This was precisely the data needed to measure their cost of quality, but contained very little information about how to improve it. The issue was that, if they do not use screening (100% inspection) of any items, the reduction in downstream defects due to sample inspection depends on common-cause defects. If defects occur independently in each part, then rejecting batches has very little effect (the uninspected items all have the same probability of a defect) and taking corrective action with the supplier has no effect. If common-cause defects affect whole batches or all future batches from the supplier, then sampling from batches and corrective action can be very effective. Their sampling procedures presumed common-cause defects, but they had very little relevant data: the distribution of defect rates across batches of the same item is needed.

The client didn't see the need for this data. When I persisted, they referred me to another department that designed their sampling plans. A statistician in this group fully understood the issue. No more data were available, but he guided me toward empirical Bayesian estimation procedures in the quality literature and what could be done with limited data. The client wanted results, so I developed a model using the dubious approach of pooling the data for all items, giving enough data to estimate the distribution. As is often the case in a large organization, the client had other needs and agendas. They wanted to use the model to justify doing less inspection, which had high labor costs. Reducing inspection would eliminate some jobs in that department. The model was used once to identify a list of items where inspecting batches was most cost effective, but ultimately had little effect on their decisions. This project illustrates that the client's view of the problem can drive which aspects are modeled, what data they make available, what activities they fund, and ultimately what models or results they use.

Analysts, trained in mathematical models, will tend to bring different motivations and values to their work than a typical client. The analyst may also differ from the client on broader values, such as social responsibility, environmental impact, or the importance of different stakeholders in the decisions being made. I would suggest that Christians have additional motivation to understand a decision-maker's values and priorities, as well as the larger values that underpin them. Within ethical boundaries and personal conscience, they should seek to reflect these values in the model, even if they do not hold the same values. The additional role, or calling, of the Christian in this

situation may be to think more deeply about the values implied by the modeling choices and how they impact other stakeholders, societies, the environment, and the broader notion of shalom.

## 6 Disaster response planning

Academics in operations research and logistics have been intensely studying how to better respond to humanitarian disasters, such as hurricanes and earthquakes [13]. This field of humanitarian logistics considers strategies for propositioning, transportation, warehousing, last mile distribution, debris removal, etc. Some of the methods apply to rapid onset man-made disasters as well.

Although many models have been developed, using them is challenging because of the unique features of disaster response. A large disaster requires an intense, rapid response, often needing supplies and personnel from far away. Infrastructure is usually damaged, vulnerable populations tend to be in remote areas, and control of the response is decentralized, with many organizations responding. Each disaster is different, and the needs in a disaster are difficult to assess and predict when it first occurs. On the spectrum in Section 4, disaster response is a hard situation to model accurately. Its unique features are discussed in [10].

Values are very important in these models. In the commercial sector, the objective is to minimize logistics cost, while in the military sector operating capability is maximized. In humanitarian logistics, the overall objective is to reduce loss of life and suffering due to lack of basic needs: water, sanitation, food, shelter, and emergency health care. One approach is to model deprivation costs, which are convex in the time until aid is delivered, and add them to logistics costs [11]. The secondary, more measurable objectives include the quantity and speed of aid, sending priority items to priority locations, cost, the impact of aid on local suppliers and the recovery phase, and media coverage to stimulate donations. Disaster relief organizations (DROs) need to consider all of the secondary objectives; most models combine several of them as in [8], or use constraints such as equity of distribution to different communities. More issues involved in these models are discussed in [9].

## 7 Christian disaster response organizations

Along with Paul Isihara and Danilo Diedrichs at Wheaton College (IL), I have been studying Christian DROs. Although several faiths are active in disaster response, we chose to only study medium to large, U.S. based, Christian organizations involved in disaster response. The primary research question is how their Christian mission and values influence the planning and implementation of disaster response. We have focused on their decision of whether to respond to a disaster and the operational decisions they make in distributing aid.

A small comparison group of non-faith-based DROs was included in the study; however, this group was not a random sample and is not very representative of all non-faith based DROs. Results from the surveys are not yet published, but here are some distinctives voiced by representatives of Christian organizations that we interviewed.

- A willingness to make more personal sacrifices when helping those in need.

	Speed	Cost	% of request	Priority items	Priority locations
Christian ( $n = 21$ )	3.0	3.8	4.7	1.9	1.7
Non-faith based ( $n = 6$ )	3.0	4.5	3.2	2.5	1.8

Table 1: Average priority given to attributes of a disaster response (1 is highest, 5 is lowest).

- A preference for meeting more needs, including emotional, psychological, for fewer people rather than meeting immediate physical needs for more people.
- Combining disaster relief with other activities, such as supporting local churches or witnessing to individuals.
- Staying longer in a disaster area, as relief transitions to recovery.
- A preference for a relational approach, where staff get to know beneficiaries.
- A desire to respond to disasters according to the needs, even when donors are not interested in a particular disaster.
- Some Christian DROs are very selective about who they partner with and do not work with governments or partner with organizations that do not share their Christian values.

One Christian organization we worked with has a system for classifying goods that are donated to the organization according to four levels of mission priority, depending on how well they align with their ongoing development programs. For example, because of their emphasis on children, new children’s clothing is usually classified in the highest mission priority. A simple model supporting the decision to accept the donation takes into account the mission priority.

In our survey, one question asked logistics managers to rank five factors of a typical first phase response to a large disaster, where requests for aid exceed the combined capacity of all respondents to deliver. The factors are taken from a study that assessed how much weight managers at large non-faith based DROs give to each factor [8]. Table 1 shows the average priority in our survey, on a five point scale where 1 is the highest priority. The Christian DROs give little importance to the total volume of aid, measured by the percent of the request met, and more importance to delivering high priority items to the locations with the greatest need. Cost is rated fairly low in importance; the difference between the two groups may be due to sampling or to the context of the question, rather than the overall concern for cost efficiency in the two groups. These results can be used as inputs to a model that optimizes the distribution of aid according to these values, e.g., truck or helicopter scheduling and routing, as done in [8].

## 8 Teaching modeling with values

In my course on optimization models for undergraduate mathematics majors, I don’t say much about values, ethics, or the ability to model some systems better than others. But maybe I should. Another course, Mathematical Models for Industry, brings up questions about the client’s values and what sort of model is possible. Students wrestle with a semester-long project from industry. Many students have said this course was very helpful. Other pedagogical changes could address values in models more directly.

One specific suggestion, which has been made repeatedly in forums about teaching operations research, is to teach fewer models and more modeling. When a model is taught, we are in a sense giving students the answer instead of the question. The quantity to be measured or optimized is usually identified for the student, as well as the general model structure. In contrast, the process of modeling includes choosing the type of model to use and the quantities to predict or optimize. This skill is often needed for real-world projects.

A second suggestion is to use messy, unstructured modeling projects that directly raise questions about values. A good starting place might be modeling in the public and nonprofit sectors, where there are often stakeholders with different values. Given the many ways that values influence models, and the desirability of getting the right values in the model to support the decision-maker, anything we can do to prepare students in this regard seems valuable.

## References

- [1] Russell L. Ackoff. *Scientific method: Optimizing applied research decisions*. Wiley, 1962.
- [2] C. Barnhart, P. Belobaba, and A.R. Odoni. Applications of operations research in the air transport industry. *Transportation science*, 37(4):368–391, 2003.
- [3] J Pierre Brans. OR, ethics and decisions: The oath of Prometheus. *European Journal of Operational Research*, 140(2):191–196, 2002.
- [4] Jean-Pierre Brans and Giorgio Gallo. Ethics in OR/MS: Past, present and future. *Annals of Operations Research*, 153(1):165–178, 2007.
- [5] C West Churchman. *The Design of Inquiring Systems: Basic Concepts of Systems and Organization*. Basic Books, 1971.
- [6] Giorgio Gallo. Operations research and ethics: Responsibility, sharing and cooperation. *European Journal of Operational Research*, 153(2):468–476, 2004.
- [7] Saul I Gass. Model world: Models at the OK corral. *Interfaces*, 21(6):80–86, 1991.
- [8] Erica Gralla, Jarrod Goentzel, and Charles Fine. Assessing trade-offs among multiple objectives for humanitarian aid delivery using expert preferences. *Production and Operations Management*, 23(6):978–989, 2014.
- [9] José Holguín-Veras, Miguel Jaller, Felipe Aros-Vera, Johanna Amaya, Trilce Encarnación, and Tricia Wachtendorf. Disaster response logistics: Chief findings of fieldwork research. In *Advances in managing humanitarian operations*, pages 33–57. Springer, 2016.
- [10] José Holguín-Veras, Miguel Jaller, Luk N Van Wassenhove, Noel Pérez, and Tricia Wachtendorf. On the unique features of post-disaster humanitarian logistics. *Journal of Operations Management*, 30(7-8):494–506, 2012.
- [11] José Holguín-Veras, Noel Pérez, Miguel Jaller, Luk N Van Wassenhove, and Felipe Aros-Vera. On the appropriate objective function for post-disaster humanitarian logistics models. *Journal of Operations Management*, 31(5):262–280, 2013.

- [12] Ralph L Keeney and Howard Raiffa. *Decisions with Multiple Objectives: Preferences and Value Trade-offs*. Cambridge University Press, 1993.
- [13] Nathan Kunz and Gerald Reiner. A meta-analysis of humanitarian logistics research. *Journal of Humanitarian Logistics and Supply Chain Management*, 2(2):116–147, 2012.
- [14] Marc Le Menestrel and Luk N Van Wassenhove. Ethics in operations research and management sciences: A never-ending effort to combine rigor and passion. *Omega*, 37(6):1039–1043, 2009.
- [15] John Mingers. *Systems Thinking, Critical Realism and Philosophy: A Confluence of Ideas*. Routledge, 2014.
- [16] John Mingers et al. The contribution of critical realism as an underpinning philosophy for OR/MS and systems. *Journal of the Operational Research Society*, 51(11):1256–1270, 2000.
- [17] William T Morris. On the art of modeling. *Management Science*, 13(12):B-707, 1967.
- [18] Cathy O’Neil. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Broadway Books, 2016.
- [19] Richard J Ormerod and Werner Ulrich. Operational research and ethics: A literature review. *European Journal of Operational Research*, 228(2):291–307, 2013.
- [20] William A Wallace. *Ethics in Modeling*. Pergamon, 1994.