



BIROn - Birkbeck Institutional Research Online

Vidgen, Richard and Hindle, G. and Randolph, I. (2019) Exploring the ethical implications of business analytics with a business ethics canvas. *European Journal of Operational Research* 281 (3), pp. 491-501. ISSN 0377-2217.

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/27315/>

Usage Guidelines:

Please refer to usage guidelines at <https://eprints.bbk.ac.uk/policies.html>
contact lib-eprints@bbk.ac.uk.

or alternatively

Exploring the ethical implications of business analytics with a business ethics canvas

Richard Vidgen*

UNSW Business School, University of New South Wales, Sydney NSW 2052,
Australia

School of Business, Economics and Informatics, Birkbeck University, London WC1E
7HX, UK

Email: r.vidgen@unsw.edu.au

Giles Hindle

Hull University Business School, University of Hull, Cottingham Road, Hull HU6 7RX,
UK

Email: giles.hindle@hull.ac.uk

Ian Randolph

Data Scientist

Email: ian.david.randolph@gmail.com

*corresponding author

Exploring the ethical implications of business analytics with a business ethics canvas

Abstract

The ethical aspects of data science and artificial intelligence have become a major issue. Organisations that deploy data scientists and operational researchers (OR) must address the ethical implications of their use of data and algorithms. We review the OR and data science literature on ethics and find that this work is pitched at the level of guiding principles and frameworks and fails to provide a practical and grounded approach that can be used by practitioners as part of the analytics development process. Further, given the advent of the General Data Protection Regulation (GDPR) an ethical dimension is likely to become an increasingly important aspect of analytics development. Drawing on the business analytics methodology (BAM) developed by Hindle and Vidgen (2018) we tackle this challenge through action research with a pseudonymous online travel company, EuroTravel. The method that emerges uses an opportunity canvas and a business ethics canvas to explore value creation and ethical aspects jointly. The business ethics canvas draws on the Markkula Center's five ethical principles (utility, rights, justice, common good, and virtue) to which explicit consideration of stakeholders is added. A contribution of the paper is to show how an ethical dimension can be embedded in the everyday exploration of analytics development opportunities, as distinct from a stand-alone ethical decision-making tool or as an overlay of a general set of guiding principles. We also propose that value and ethics should not be viewed as separate entities, rather they should be seen as inseparable and intertwined.

Keywords: business analytics; data science; business ethics canvas; Markkula; GDPR

Exploring the ethical implications of business analytics with a business ethics canvas

1 INTRODUCTION

Business analytics is playing a greater and greater role in our daily lives, impacting on job applications, medical treatment, parole eligibility, and loans and financial services. There are undoubted benefits to algorithmic decision-making in general, and artificial intelligence (AI) in particular. For example, AI is being used to detect the early stages of colorectal cancer, achieving 86% accuracy (Mukherjee, 2017). Such is the interest in AI for healthcare that the UK Government is pledging millions to AI applications for the early diagnosis of cancer and other chronic diseases, using patient data and lifestyle information to highlight patients at risk (Perkins, 2018). However, algorithmic decision-making is not without its dark side. Mann and O'Neill (2017) question the use of algorithms in hiring decisions, d'Alessandro et al. (2017) raise concerns about predictive policing. The Cambridge Analytica case has thrust data analytics squarely into the public domain. It is alleged that Cambridge Analytica collected data from more than 50 million Facebook users (without permission) and used that data to build a system to target US voters with personalized political advertisements with the aim of influencing the US election outcome (Greenfield, 2018). Unsurprisingly, algorithmic decision-making is attracting the interest of researchers as well as practitioner and regulators (e.g., Newell and Mirabelli, 2015; Kitchin, 2017).

The potential for harm – intended or unintended - arising from algorithmic decision-making indicates that an ethical dimension is needed. For example, Google, on discovering that its AI software was being used by the US military in its drone development programme, has pledged not to use AI for weaponry (Statt and Vincent, 2018). Google's CEO, Sundar Pichai, published a list of ethical principles for AI development, which include: be socially beneficial, avoid creating or reinforcing unfair bias, and be accountable to people (Pichai, 2018). These principles are prefigured by the Toronto Declaration, which is calling for governments and companies to ensure that algorithms respect basic principles of equality and non-discrimination (Brandom, 2018). Rights are being further encoded in legislation such as the General Data

Protection Regulation (GDPR), which requires organizations to protect the rights and privacy of individuals with associated constraints on data usage and a responsibility to provide a right to explanation and to address any presence of discrimination and bias.

We consequently argue that an ethical dimension to algorithm development and algorithmic decision-making is an essential aspect of the OR practitioner's professional profile. Indeed, concern for the ethical aspects of OR goes right back to the early pioneers of the discipline (for example, Churchman 1968, 1970; 1971; Ackoff 1974a, 1974b) who included a concern for stakeholders and wider society within their conceptualization of OR intervention. And a degree of concern for the ethical aspects of OR as a profession is reflected in early guidelines and codes created by OR societies such as the OR Society of America (ORSA) in the USA (Caywood et al., 1971) and the Fellowship for OR in the UK (Fellowship for Operational Research, 1974).

However, despite this evidence, we argue an explicit concern for ethics and the 'goodness' of OR practice has tended to fall outside of traditional or mainstream discussions on OR and, more latterly, data science. Such discussions have tended to focus on the efficacy of a range of quantitative modelling techniques and on how to achieve operational and process improvement in practice (Koch, 2000). The practice of OR has generally been regarded as a 'good thing' within OR communities due to the largely uncontroversial nature of process improvement within organisations and its scientific credentials and associations. Subsequently, a concern for ethics and ethical practice is not generally covered in a substantive way within OR textbooks or on courses in OR.

Thus, while ethics is a long-established and on-going concern for the OR community we argue that much of this work has been at too abstract a level for it to impact meaningfully on the lived day-to-day experience of OR and data science practitioners. For example, in the context of software development, McNamara et al. (2018) found no evidence that the Association for Computing Machinery (ACM) code of ethics influences software-related ethical decision making. Given the development of data science practice and the emerging regulatory environment this abstract approach to ethics is not sufficient – practitioners need practical tools, not just frameworks. Our aim, therefore, is to address the question: how can an ethical dimension be built into the process of business analytics development? We tackle this question using action

research. In the next section we provide the background to ethics in OR and in section three we describe the development of a framework for ethical guidance in data science. The research approach is outlined in section 4. In sections five, six, seven, and eight we present the action research project under the headings of diagnosis and planning, action taking, evaluation, and reflection. The paper concludes with a summary.

2 OR AND ETHICS

There has been an ongoing concern for the ethical aspects of OR within the community, which can be traced back to the early pioneers of the discipline (for example, Churchman 1968, 1970; 1971; Ackoff 1974a, 1974b). These researchers envisioned a broader role for OR than simply process engineering and proposed that a concern for stakeholders and wider society should be part of OR's methodology. However, despite this genuine concern, we argue ethical support for practitioners has tended to fall outside of mainstream discussions on OR practice (Koch, 2000). This is because the practice of OR has generally been regarded as a 'good thing' within the OR community due to the largely uncontroversial nature of process improvement within organizations and OR's scientific credentials and associations. Ethical issues in OR have been addressed by both professional societies and by academics.

2.1 The professional society perspective

The website of the Institute for Operations Research and the Management Sciences (INFORMS) in the USA contains ethical guidelines (INFORMS, 2018) and the Operational Research Society (ORS) in the UK lists ethical principles (OR Society, 2018).

The INFORMS guidelines are split into three sections: society, organizations and the OR profession (INFORMS, 2018). The 'society' section is concerned with aspiring to openness of assumptions, objectives and sponsors, to objectivity of analysis whilst being respectful of other views and values, and to undertaking work that provides positive benefits, such as progressing scientific understanding, organizational improvement and supporting the social good. The 'organizations' section is concerned with aspiring to be accurate, rigorous and realistic in conducting analysis, whilst being alert to the possible unintended consequences of recommendations and being aware of developments in other fields. The 'profession' addresses behaviour towards colleagues and to be vigilant and speak out against actions that may damage the profession.

The ethical principles of the OR Society are split into four sections: accuracy and rigour, honesty and integrity, respect for life, law and the public good, and responsible leadership (OR Society, 2018). Accuracy and rigour relates to practicing only within the sphere of our competence and delivering work openly, honestly and without bias. Honesty and integrity is concerned with being aware of how OR practice might affect other people and rejecting all forms of corruption and deception. Respect for life, law and the public good relates to acting lawfully, minimizing adverse impacts on society and the natural environment, and protecting the reputation of the profession. Finally, responsible leadership includes listening to the aspirations and concerns of others and promoting public awareness of the impacts and benefits of OR. The ORS claim these principles are fully compatible with the principles in the UK Government Chief Scientific Adviser's Universal Ethical Code for Scientists (GOV.UK, 2007).

The guidelines and ethical principles found on the INFORMS and ORS websites can only be viewed as a generic combination of professional best practice and traditional scientific values. There is limited practical guidance on how to deal with ethical problems and judgments within OR interventions, or even discussion of what the ethical issues might be in an age of analytics and data science.

2.2 The academic perspective

From an academic point of view, Ormerod and Ulrich argue there has been “sustained advocacy by a number of scholars” regarding the relationship between ethics and OR, although they admit the area is relatively underdeveloped (Ormerod and Ulrich 2013, p.292). They map scholarly contributions onto the core competences of OR to arrive at four aspects of OR and ethics: ethics and society (appreciating the purpose and context of OR practice), intervention ethics (managing OR projects), ethics in analysis (the application of OR techniques), and personal and professional ethics. They define an OR consultant as “an individual, engaged in analysis, within an investigation located in the context of an organization embedded in a particular society” (Ormerod and Ulrich 2013, p.292).

The first of the four aspects, *ethics and society*, covers topics such as the role of OR within society, the impact of OR interventions upon third parties and society at large (for example, Ackoff, (1974a); Churchman, 1979; Rosenhead, 1976; Midgley and Munro,

1998; Ulrich, 2007), and the questions of which organizations to work for and what problems to work on (for example, Ackoff, 1974b; Rosenhead and Thunhurst, 1982; Midgley and Ochoa-Arias, 2004). However, the review by Ormerod and Ulrich identified limited and sporadic discussion within the OR community under this heading.

The second aspect, *intervention ethics*, relates to how OR projects are conceived of and managed. "The process adopted for an OR intervention will influence the way that key ethical choices are made: what is the project scope, who is to be involved in what capacity, what are the organizational aims to be met and the societal norms to be complied with, what constraints apply?" (Ulrich 2006, p. 296). The topic has attracted an ongoing debate within OR, but has been situated mostly within the Soft OR and systems communities. Various approaches to OR interventions have developed over the years, which are located by Ormerod and Ulrich (2013) along a spectrum from "a single investigator working on a well-defined local problem for a single client" at one end to "a team of investigators working on a complex problem with extended social implications" (p. 296) at the other. They argue much of the debate about the process of OR has concerned interventions lying somewhere between these two extremes - termed "mid-range interventions" (p. 297). This is because interventions at the well-defined end of the spectrum (for example, queuing, inventory, scheduling, and logistics problems) have historically been considered relatively uncontroversial from an ethical point of view and interventions at the complex societal end are normally embedded within a wider political process involving a range of stakeholders, "which lays bare the ethical issues" (P.297). The advent of business analytics and data science, however, has problematised this narrative. Data science practice at the well-defined end of the spectrum can no longer be viewed as uncontroversial from an ethical point of view.

In the third aspect, *ethics in analysis*, Ormerod and Ulrich split their commentary into two sections. First, they consider the depiction of ethical issues within OR models. For example, the selection of an objective function to be optimized and the choice of constraints – for example, do we include CO₂ emissions as a constraint (or even an objective) when optimizing an electricity generating system? This type of modelling can raise a number of normative considerations and in some cases the elicitation of values can become problematic (Keeney 1994). Second, they consider the ethical issues that might arise during the activity of model building. Mason (1994) argues OR practitioners

have three ethical obligations: to represent reality adequately to clients, to incorporate the client's values into the model, and to ensure the client's actions are ultimately effective. Business analytics and data science are impacting OR in terms of the need for practical guidance in the area of ethics in analysis. The proliferation of application areas outside of the traditional process engineering focus of OR, the use of new types of software and statistical analyses and unprecedented access to new sources of data means that OR analysts now face a range of ethical issues and choices.

With regard to the fourth aspect, *personal and professional ethics*, Ormerod and Ulrich (2013) argue the OR community has always implicitly assumed practitioners would ensure ethical issues are appropriately dealt with in practice - and that any ethical dilemmas would be resolved through consultation with clients and colleagues. This explains why ethics is not routinely found on OR training courses or within OR textbooks, and why the ethical codes and guidelines of the professional societies have not been regarded as central to professional development activities.

In summary, despite attempts by the OR community to develop approaches to the treatment of ethics in OR and evidence of ethical codes of practice from OR societies around the world, in practice the OR community has largely tended to bypass the topic of ethics. This has been the result of various historical factors including viewing OR practice as non-contentious (e.g., in the case of process engineering applications such as logistics and inventory management), by viewing OR as an applied scientific activity and therefore implicitly a 'good thing' within society, or by treating ethics in the abstract through reliance on generic ethical checklists and codes of practice such as those published by INFORMS and the OR Society. However, we argue the development of business analytics and data science are changing the situation for OR. New application areas, new techniques and new sources of data are driving the need for practical guidance for OR practitioners and academics.

3 AN ETHICAL FRAMEWORK FOR DATA SCIENCE

Barocas and Boyd (2017) argue that ethical practice in data science is not simply an overlay that is dealt with independently or on top of technical practice (p. 24). Data scientists are involved in making trade-offs, judgments (e.g., is this error rate acceptable), and some will look explicitly at issues of fairness, e.g., by testing for gender

bias (Barocas and Boyd, 2017). However, in many situations there is no absolutely right or wrong answer and in favouring some fairness properties over others may reflect a difference in values, “rather than a failure to recognize the values at stake.” (Barocas and Boyd, 2017, p25). To make and defend moral judgments requires a normative ethics framework to provide arguments and rules for moral reasoning (Schwarz, 2005). Schwarz argues that the development and adherence of professional rules of conduct are essential but not sufficient to provide guidance in situations of ethical uncertainty. Schwarz also warns against ethical practice driven by a “religious-philosophical” agenda (p. 67). Following Schwarz, we argue that the practice of OR and data science needs an ethical toolset that: (1) goes beyond general guidelines for professional and ethical practice; (2) embodies a tool-set that does not require a deep background in philosophy; (3) reflects the normative status of ethical reasoning; and (4) can be applied practically to real-world ethical decisions.

In addressing these four concerns, Schwartz (2005) adopts the Markkula Center’s ethical framework (Markkula, 2018a), which identifies five sources of ethical standards:

Utilitarian. Ethical corporate action “is the one that produces the greatest good and does the least harm for all who are affected - customers, employees, shareholders, the community, and the environment.” (Markkula, 2018a). The utilitarian approach focuses on outcomes and attempts to maximise good done while reducing harm.

Rights. Humans should have the ability to choose freely what they do with their lives, i.e., we should respect human dignity. Human rights include being told the truth, to not be injured, and a right to privacy.

Fairness or Justice. Based in Greek philosophy, the justice approach proposes that all humans should be treated equally and where people are treated unequally then this is done on the base of defensible criteria (e.g., justification of why some people are paid more than others).

Common Good. Society is more than the sum of individuals: “the interlocking relationships of society are the basis of ethical reasoning and that respect and compassion for all others - especially the vulnerable” (Markkula, 2018a). There are common conditions (e.g., policing, health care, education) that are needed to protect the welfare of all members of society.

Virtue. Our actions should be consistent with ideal virtues that promote the full development of our humanity. Examples of virtues include truth, beauty, honesty, courage, compassion, generosity, tolerance, love, and integrity. Virtue ethics is concerned with us becoming the best person (and organisation) we can be.

The Markkula ethical framework proposes asking questions in these areas, relating to the five sources of ethical standards (Table 1). Schwarz (2005) has applied the Markkula ethical framework in the area of information assurance, and provided short summary questions for each of the five dimensions (Table 1). While Schwarz frames the application of the Markkula framework to information assurance, we frame the implications explicitly within the domain of business analytics and OR (Table 1).

Ethics source	Ethical questioning (Markkula, 2018a)	Implications for business analytics (adapted from Schwarz, 2005)
Utilitarian	Does this action produce the most good and do the least harm for all who are affected? What good and what harm will or may result? How will I measure a good outcome? Happiness? Financial impact? While the potential harm from this action may affect only a few people, is the harm so great that it would outweigh the good this action might bring to many others?	What are the benefits of the intended analytics use case? What are the harms created? Who benefits and who is harmed?
Rights	Does my action best respect the rights of all who have a stake? Does this action respect the dignity of others? If I take this action, am I treating others simply as a means to an end? Does the action hurt or help others in securing a minimum level of well-being?	Whose rights are respected or infringed by this analytics use case? What are those rights?
Justice	Does this action treat people equally or proportionally? Does it give each person affected his or her due? Might I have some prejudice or interest that might make me favor one person over another? Am I treating each individual the same way, or is there a valid reason to treat someone differently?	How fair are the outcomes of the analytics use case? Do they treat everyone in the same way or do they show favoritism and discrimination?
Common Good	Does this action best serve the community as a whole, not just some members? Will this option be equally to everyone's advantage? Does this action contribute to the conditions of social life that give everyone an opportunity to thrive? How will my action affect the resources everyone must share, such as the environment?	What is the community (or what are the communities) in which the analytics use case is to be developed? What constitutes the common good?
Virtue	Does this option lead me to act as the sort of person I want to be?	How does this analytics use case define me as a human person?

	<p>What character traits would I be exhibiting if I chose this action? Honesty or deceit? Compassion or selfishness? Prudence or irresponsibility?</p> <p>What habits of character would I be developing if I took this action? What would a person I respect say about this choice?</p>	<p>How does it define us as a company, an organization, a society, etc.? What do I or what do we want to be and become as a result of the outcomes of this analytics use case?</p>
--	--	--

Table 1: the Markkula ethical framework and its application in information assurance

While Table 1 addresses the content of an ethical investigation, the Markkula Institute also provides a five-stage process model (Markkula, 2018a). Firstly, there is recognition of an ethical problem. Secondly, information relevant to the problem (e.g., identifying stakeholders) is gathered. Thirdly, the problem is looked at questioningly from different perspectives. The fourth stage in the process is to come to a decision and the final stage is to revisit the outcomes of the decision.

The five principles and the five-stage process model offered by the Markkula Institute provide a useful starting point for investigating the ethics of algorithmic decision-making by data scientists and OR practitioners. The challenge for the current research is to apply the ethical principles as part of the business analytics development process in a practicable way that makes sense to working data scientists and managers.

4 RESEARCH APPROACH

In order to develop our ethical approach to business analytics, an action research framework has been employed (Susman and Evered, 1978; Eden and Huxham, 1996; Baskerville and Wood-Harper, 1996; Checkland and Holwell, 1998) involving a real world intervention. The primary purpose of the intervention is to explore the ethical governance of business analytics, but it’s worth noting that further development of the business analytics methodology (Hindle and Vidgen, 2018) was also viewed as a desirable outcome. According to Checkland and Poulter (2006), the key criterion of action research is to achieve recoverability, “that is to say, make the whole activity of the researcher absolutely explicit (including the thinking as well as the activity)” (p.177). In order to achieve this, they argue, the researcher must state in advance “the framework of language (the epistemology) in terms of which what counts as knowledge from the work will be expressed” (p.177). The definition of an epistemology also helps differentiate action research from consultancy (Baskerville and Wood-Harper, 1996). For the purposes of this research the epistemology is based on the concept of a soft

systems methodology (SSM) human activity system (Checkland, 1981; Checkland and Scholes, 1990). The framework of ideas is provided by the business analytics methodology (BAM) proposed by Hindle and Vidgen (2018) and the Markkula ethical framework.

BAM provides organisations with a process framework and a set of tools for establishing valuable business analytics practice. The BAM process entails four components for analytics development: (1) problem situation structuring using SSM rich picturing, (2) business model exploration using the business model canvas (Osterwalder and Peigneur, 2010) in conjunction with SSM purposeful activity system modelling, and (3) business analytics leverage analysis to identify and categorise development projects. The fourth stage is analytics development and deployment. The BAM provides an overarching framework of ideas for business analytics development in which we will explore ethical governance.

The action research cycle consists of a number of activities. For example, Susman and Evered (1978) proposed a model with five stages: diagnosis, planning, intervention, evaluation and reflection. As part of their canonical action research cycle Davison et al., (2004) redefined the five stages as: diagnosis, action planning, intervention (action taking), evaluation (assessment), and reflection (learning). Both the Susman and Evered and the Davison et al. models link reflection to the diagnosis stage of a subsequent round of action research. Davison et al. show an entrance point leading to diagnosis and an exit point from reflection. Thus, an action research project consists of one or more cycles of intervention. We structure our account of the analytics ethical governance action research using the stages defined by Davison et al.

5 DIAGNOSIS AND PLANNING

The client organization for the action research is EuroTravel (a pseudonym for the purposes of confidentiality), a European eCommerce travel business that sells travel tickets within the UK and Europe. The business supports the purchase of tickets using a full range of devices and also provides technical services to travel operators and administrative services to larger organizations for business travel. The business model is based upon a highly developed computer-based booking platform that provides unified booking across many travel operators, advance travel planning, travel

budgeting, and a range of other services. The business is supported by a data science team that develops products for the booking platform and provides insight work for marketing.

Central to EuroTravel operations is the booking platform, which is the gateway to customers. The data science team is active in gaining insights from customer data and building data products within the platform. Their objective is to improve the customer experience in order to drive customer retention, conversion ratios and frequency of purchase. Following the BAM (Hindle and Vidgen, 2018) we developed a rich picture and a range of SSM models to help us structure the problem situation and to think about the purpose of EuroTravel in terms of meaningful human activity. In conjunction with the SSM modeling we created a business model canvas (BMC) (Osterwalder and Peigneur, 2010). In the interests of brevity we focus here on the BMC as this provides a starting point for diagnosis and planning of the intervention. The BMC focuses on the sale of travel tickets to UK and European customers (Figure 1). The customers are segmented into leisure, business and commuter and the value propositions are listed as the provision of a unified booking facility across travel operators, advance travel planning, travel budgeting, travel updates and travel experiences. The key resource is the booking platform and key activities are based around the development of this platform as well as marketing analytics (e.g., SEO - search engine optimization, PPC - pay per click, ATL - above the line advertising, and social media).

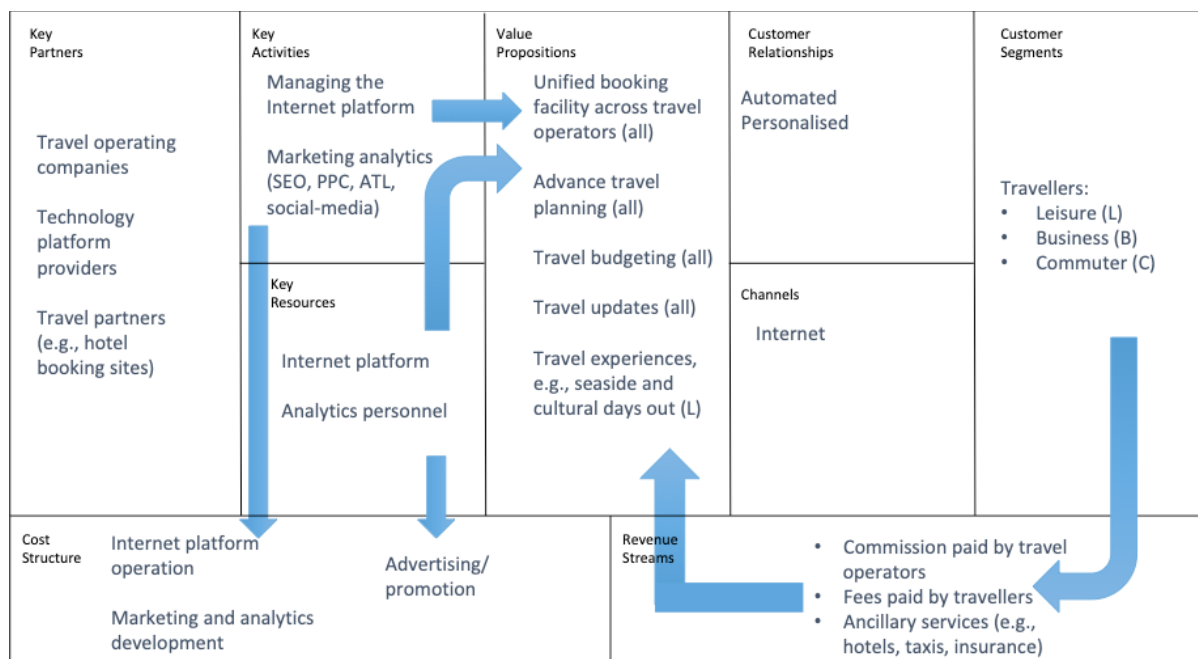


Figure 1: EuroTravel business model canvas

The next stage in the BAM is to use the BMC to identify analytics use cases. We identified a number of use cases (more than 10) and then selected one of these for further investigation of the ethical dimension of analytics development. The use case selected is “DaysOut”. The DaysOut project is concerned with making recommendations to customers and was defined to be: “As a user, I want to know of attractive, feasible daytrips for me so that I can have an adventure if I so choose”. This use case was selected as it is in live development, involves using analytics to make customer-specific recommendations, is a potentially valuable source of added revenue for EuroTravel, and the ethical aspects have not been systematically explored.

6 INTERVENTION (ACTION TAKING)

Having mapped the business model and selected the DaysOut project for further exploration we then extended the BAM to include an opportunity canvas (Patton, 2016). We create the opportunity canvas before looking at the ethical aspects of the use case in order to gain a more detailed understanding of the problem to be addressed using analytics and to understand business benefits afforded by the use case

6.1 Opportunity canvas

The opportunity canvas (Figure 2) is used to articulate the following (adapted from Patton, 2016):

Problems: What problems do prospective users and customers have today that the solution addresses? What needs, goals, and activities will the solution meet?

Solution ideas: What are the product, feature, or enhancement ideas that solve problems for the target audience?

Users and customers: Who are the users and customers that have challenges that will be addressed by the solution?

Users & Customers	Problems	Solution ideas	Solution use	User Metrics
Leisure travellers	Lack of awareness of adventure opportunities. Need for excitement and diversion	As a user, I want to know of attractive, feasible daytrips for me so that I can have an adventure if I so choose	Users will book adventure trips for discretionary travel	Number of travel adventures proposed to customers that are taken up
	Solutions Today Random and infrequent adventure travel		Adoption Strategy Above the line advertising Targeted, personalised communications via the Internet	
Business Challenges Without this user story customers will continue to have a transactional relationship with Trainline		Budget Additional revenue will be created to offset cost of development	Business Benefits and Metrics Increased revenues from ticket sales and ancillary services (e.g., hotel bookings) Transformation of customer relationship to experiential rather than purely transactional	

Figure 2: Opportunity canvas for use case ‘DaysOut’ (adapted from Patton, 2016)

Solutions today: How do users address their problems today? What competitive products or work-around approaches are available?

Business challenges: If we don’t solve these problems for our customers and users, how will it affect our business?

Solution use: What will customers and users do differently as a result of adopting the solution and how will that benefit them?

User metrics: What specific user behaviors can be measured that will indicate they try, adopt, use, and place value in the solution?

Adoption strategy: How will customers and users discover and adopt the solution?

Business benefits and metrics: What business performance metrics will be affected by the success of the solution?

Budget: What will our organization earn or save if the solution is created? What might it cost if the solution is not created?

In Figure 2 we have completed the opportunity canvas for the DaysOut use case. As can be seen, the focus is very much on a business problem (opportunity), a solution, and the user of that solution. The opportunity canvas places the use case squarely within a business context of benefits, challenges, and budget. Having developed an opportunity canvas for the target analytics use case DaysOut, we now turn to an exploration of ethical implications.

6.2 Developing the business ethics canvas template

To operationalize the five ethical concerns of the Markkula framework – and to show their interconnectedness - we overlay them on to a business ethics canvas (BEC) (Figure 3). The canvas was developed iteratively and through practical application. Numerous variations of format were tried until a stable version that was usable and ‘felt right’ was arrived at. The process also gave insight into the order in which the elements might be addressed.

Virtue <p style="text-align: center; font-size: 2em;">7</p> <p>How does this solution define me as a human person? How does it define us as a company, an organization, a society, etc.? What do I or what do we want to be and become?</p>	Users & Customers <p style="text-align: center; font-size: 2em;">1</p> <p>What types of users and customers have the challenges our solution addresses?</p>	Solution ideas <p style="text-align: center; font-size: 2em;">1</p> <p>What are the product, feature, or enhancement ideas that solve problems for our users and customers</p>	Stakeholders <p style="text-align: center; font-size: 2em;">2</p> <p>Who is affected by, or can affect, the proposed solution? What is their stake in the proposed solution?</p>	Utility <p style="text-align: center; font-size: 2em;">3</p> <p>What are the benefits of the intended solution? What are the harms created? Who benefits and who is harmed?</p>
Common good <p style="text-align: center; font-size: 2em;">6</p> <p>What is the community (or what are the communities) in which the decision has to be made? What is the common good?</p>		Justice <p style="text-align: center; font-size: 2em;">5</p> <p>How fair is the solution? Does it treat everyone in the same way or does it show favoritism and discrimination?</p>		Rights <p style="text-align: center; font-size: 2em;">4</p> <p>Whose rights are respected or infringed by this action? What are those rights?</p>

Figure 3: Business ethics canvas (BEC) template

In the developed form of the BEC, the first elements to be completed are the “Solution ideas” and “Users & Customers”, both of which are taken directly from the opportunity canvas (Figure 2), thus providing a direct link to the business rationale for the analytics development. In defining the content of the elements of the canvas we changed the language of the questions from “action” to “solution” to mirror the business language of the opportunity canvas and phrased each box as a question.

Secondly, stakeholder identification is conducted to understand who can affect, or is affected by, the proposed development (Freeman, 1984). This gives a pleasing symmetry of a solution surrounded by customer and stakeholders. Thirdly, having established the opportunity and stakeholder context, the five ethical dimensions are considered (we have included the definitions from Schwartz (2005) in Figure 3 as these are brief and to the point). We found that once solution and user had been copied from the opportunity canvas then the analysis flowed comfortably from utility to rights, right to justice, and then broadened out to the common good. Lastly, virtue is considered. Virtue is concerned with the organization (and the individuals in the organization) being the best it can be. It might also be couched as “What would happen if this appeared in the press tomorrow?” – would it enhance or damage our brand? Interestingly, the order we arrived at is the same order of ethical concerns that is adopted in the Markkula app for ethical decision-making (Markkula, 2018b.) The BEC arrangement depicts the ethical concerns “surrounding” the solution and the stakeholders, keeps the ethical subjects adjacent to every concern, and allows for an elegant clockwise workflow of exploration.

6.3 Applying the business ethics canvas template

While we have a logical sequence to the BEC development process, in practice it goes through many iterations in a non-linear way. After workshopping and discussion the completed canvas for the DaysOut analytics opportunity was arrived at (Figure 4). The contents of the boxes are developed using sticky notes. Writing directly on the canvas is a basic mistake, as it doesn’t allow for repositioning of ideas as thinking evolves. The stickies allow for colour coding, where:

- yellow = neutral
- green = opportunity

- red/pink = risk

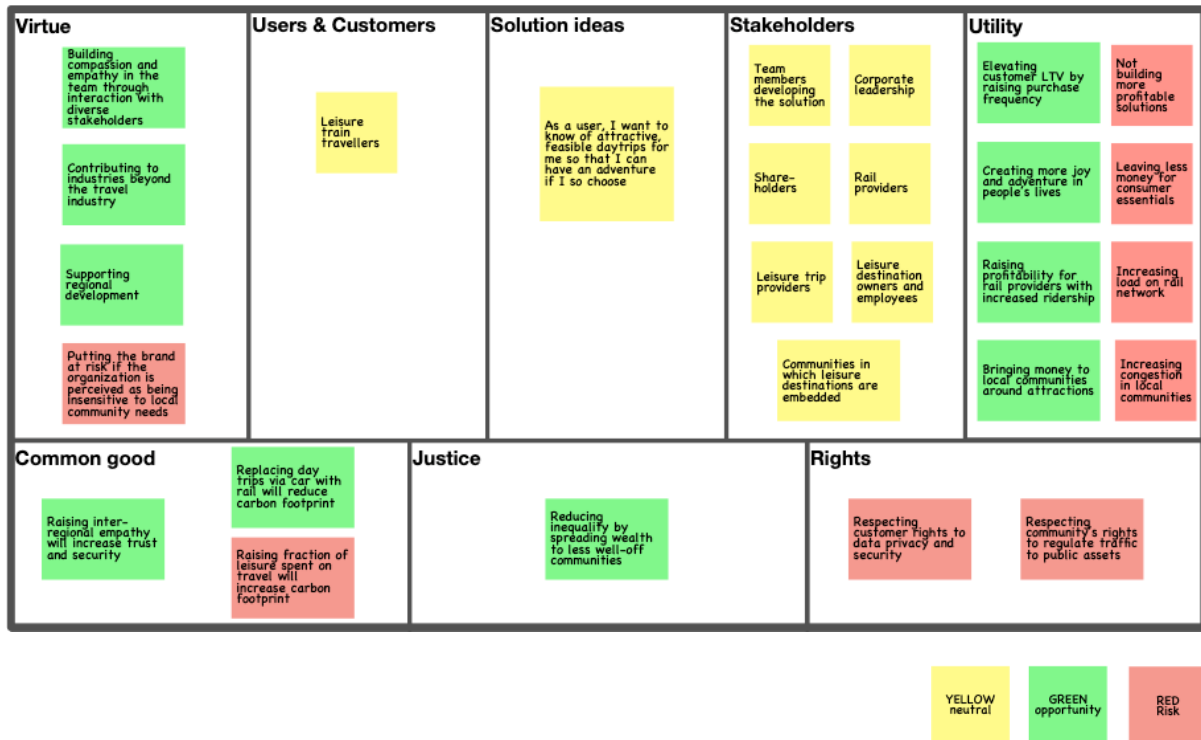


Figure 4: Business ethics canvas (BEC) for the use case DaysOut

The process of BEC development (leading to Figure 4) unfolded as follows.

(1) Preparation

- Identify user stories or solution concepts which look sufficiently attractive to action based on the opportunity canvass (Figure 2)
- Identify and gather stakeholders who understand the solution, the context or who have a stake in the proposed solution (i.e., are affected by or can affect the solution)
- Print out several ethics canvases on large paper (size A1 or A2 are ideal) and attached them to the wall
- Prepare yellow, green, and red/pink sticky notes and sharpies to write stickies and add votes

(2) Solution identification

For each analytics solution, with a facilitator scribing the notes, go through the canvas:

- Define solution idea and customer whose problem it will solve – in yellow stickies.
- Define all stakeholders who can affect or be affected by the proposed solution – in yellow stickies. Where stakeholders are not available then workshop members should role-play those stakeholders, expressing and testing assumptions.

(3) Ethical exploration. For each ethical area:

- Define positive opportunities that may arise from this solution to enhance the situation of one or more stakeholders along this ethical dimension. Write stickies in positive action language beginning with an '-ing' verb to focus on the outcome in terms of what would happen or change.
- Define risks that may result from the proposed solution that diminish the situation of one or more stakeholders along each ethical dimension. Write stickies in positive action language beginning with an '-ing' verb to focus on the outcome in terms of what would happen or change.
- Give each participant a limited number of votes and ask them to use their votes to put a dot against the most important risks to be addressed and opportunities to be developed. In this way the collective intelligence of the group is leveraged to focus attention on the most consequential ethical concerns.
- Rank concerns in order of votes and then explore and problem solve mitigation and exploitation strategies in priority order.
- If deal-breaker risks cannot be responsibly mitigated, put the solution to one side and re-visit it later. Otherwise focus on enhancing the ethical value of the initiatives.

(4) Revise the opportunity canvas in light of the BEC.

- For example, identification of leisure organisers and business owners in the BEC may lead to business value creation opportunities, such as providing dashboards to local attractions which help them predict footfall by monitoring incoming traveller traffic.

- Repeat the iteration between Opportunity Canvas and Business Ethics Canvas until a value-creating and ethically acceptable solution design is reached. Be open to the possibility that an OC may be systemically desirable but not culturally feasible.

(5) Be prepared to revisit these canvases post-implementation with appropriate regularity.

7 EVALUATION (ASSESSMENT)

In the action taking stage we have developed the BEC through exploration of the DaysOut use case. The next step in the action research cycle is to make an evaluation of the intervention. Would the BEC be adopted by EuroTravel and if not, why not? In order to consider the barriers and enablers to the adoption of the BEC we consider the BEC as an organizational innovation. Organizational innovation is commonly defined as the “[i]mplementation of a new organizational method in the firm's business practices, in the organization of its workplace or in its external relations, to improve the use of knowledge, workflows efficiency or quality of goods or services” (Zucoloto and Nogueira, 2016, p. 374). In the organizational innovation literature, two well-established models are particularly relevant: Rogers’ (2003) theory of innovation diffusion in organizations and the technology–organization–environment (TOE) framework introduced by Tornatzky and Fleischer (1990). These two models are consistent with each other and together provide a comprehensive and insightful framework for thinking about organizational innovation (Zhu et al., 2003). Drawing on the Rogers and the TOE model we frame our innovation model as COE: characteristics, organization, and environment and use this structure to guide our assessment of the BEC.

Rogers (2003) identifies characteristics (C) of an innovation as relative advantage, compatibility, complexity, observability, and trialability. The organization (O) perspective is concerned with the benefits, costs, and barriers to innovation (for example, structural and political issues). Under environment (E) regulatory and legal requirements, and relationships with external partners (e.g., customers and suppliers) are considered. Together the three dimensions – COE – provide a framework for investigating the innovation itself (the BEC), the organizational setting for the

innovation (EuroTravel), and the wider environment in which the organization operates. The academic members of the research team interviewed the lead practitioner using the COE framework as a guide to in order to provide a structured assessment of the action-taking.

The characteristics of the BEC indicate that it has a relative advantage since, to the best of our knowledge, there is not a comparable ethics methodology in use in business analytics. The BEC is compatible with analytics development at EuroTravel insofar as the use of workshops, canvasses, and aspects of design thinking is standard practice. The BEC is also relatively low in complexity, taking difficult ideas about ethics and simplifying them into five dimensions using a visual, canvas-based approach. Further, the BEC is observable and, as we have evidenced through action research, is eminently triable. In summary, the BEC would seem to score well on the *characteristics* of a successful innovation.

However, having positive innovation characteristics is no guarantee of organizational success. In the organizational reflections we first focused on the business benefits of using the BEC:

“With regards to concrete outcomes, what we can say is that using the BEC helped us to consider a stakeholder we hadn’t really considered before – the owner/operators of the attractions to which we would be driving our customers. We’ve now crafted a strategy for engaging and partnering with these stakeholders so that our customers get the best experience at these attractions and the owner-operators can give us feedback on issues such as congestion. This strategy includes an outreach campaign to attraction owner/operators to notify them of what we’re doing and offering to partner on co-promotion and ensuring a great experience for customers. This campaign alone created meaningful bottom line results for the company.” (Product Owner, DaysOut)

The BEC further enabled EuroTravel to design the value proposition for the use case giving explicit consideration to the ethical risks identified on the canvas, motivating the campaign to work with the owner-operators of attractions rather than to simply funnel travellers to leisure destinations.

However, adoption and implementation of the BEC faces organizational challenges. One barrier to adoption is a lack of clarity about when to use the business ethics canvas. Using the canvas entails an organisational cost in terms of time, energy, and potentially the coordination of a broad range of stakeholders, some of whom will be senior managers and some will be hard to access (e.g., local residents in a DaysOut travel

destination). Using the BEC for every project would likely represent an unacceptable overhead in terms of time and cost; therefore, a challenge to be addressed is whether to embody the BEC in every project or to have a selective mechanism for discerning when a product or initiative requires deeper ethical scrutiny.

Once it has been agreed that it is appropriate to create a BEC for a project, engaging in the exercise of building an ethics canvas requires a high degree of trust and openness between stakeholders. Without this trust, participants in the exercise won't feel comfortable speaking their minds about their concerns (for example, this might be due to a fear of isolation or reprisal). For individuals, speaking up on sensitive issues such as ethics can be a dangerous choice for their career. Using the BEC canvas also risks recasting past behaviour as unethical, creating cognitive dissonance for those responsible for past initiatives and likely to lead to defensive reactions.

In considering the environment (E) we focused on legislation and regulation, chief of which is the General Data Protection Regulation (GDPR). The GDPR defines and strengthens data protection for consumers and harmonizes data security rules within the European Union (EU). The GDPR came into effect on May 25, 2018 and regulates how organizations collect, store and process personally identifiable information (PII) about EU citizens. Any company that stores or processes data on EU citizens must comply with the GDPR legislation, including US, Asian, Australasian companies that sell goods and services in the EU. Dinsmore (2017) argues that GDPR affects data science practice in three areas: (1) it imposes limits on data processing and consumer profiling; (2) it creates a "right to an explanation" for consumers subjected to automated decision-making; (3) it holds firms accountable for bias and discrimination in automated decisions. Many aspects of the GDPR and its implications for organizations and the practice of business analytics are emerging as the legislation is interpreted and clarified by organizations. Given that the GDPR is concerned with the rights and privacy of individuals and the use of their personal data the BEC can help with GDPR compliance through explicit consideration of rights (e.g., to privacy, to an explanation) and justice (e.g., the avoidance of bias).

In summary, EuroTravel will continue to experiment with the BEC, applying it to new products and renewing it for products on which it has been applied (such as

DaysOut) while recognising that the BEC has yet to be incorporated into business as usual.

8 REFLECTION (LEARNING)

Our action research shows that ethical analysis should not be seen as a constraint or overhead to analytics development – exploring the ethical dimension and including multiple stakeholders provides a richer insight into business value creation, as well as providing greater confidence about emerging ethical implications. Our research suggests that we can position the business ethics canvas (BEC) and the opportunity canvas (OC) as counterparts and that placing them both within the context of the business model canvas (BMC) helps align analytics development with the mission and vision of the organization (Figure 5). The BMC is used to structure analytics development opportunities, which are then subjected to leverage analysis to identify projects that warrant further investigation. For each potential project an OC and a BEC are developed and explored jointly. Rather than see these as separate activities addressing separate questions (what is the business value creation opportunity? What are the ethical implications?) we position them as counterweights in a relationship of creative tension. The OC provides the BEC with a solution, users, and customers. The BEC provides the OC with stakeholders and an ethical perspective. Together, the two canvasses enable a joint exploration of value creation and ethical aspects, allowing a solution to be framed and reframed as the exploration of the problem situation unfolds.

Our action research approach aims to make ethics a tangible and lived part of the daily experience of analytics development. By integrating the BEC in the development process the ethical dimension becomes part of the exploration process rather than a ‘tick-box’ or ‘rubber-stamping’ activity. We found that the process of ethical exploration follows the five-stage process proposed by the Markkula Center (Markkula, 2018a). First, there is recognition of an ethical problem. This recognition should first happen when use cases are identified and the analytics portfolio created (Figure 5), which may lead to some use cases being rejected due to ethical concerns. As opportunity canvasses are developed then further ethical problems may be unearthed. Secondly, there is a gathering of information relevant to the problem (e.g., identifying stakeholders), which will inform the development of the business ethics canvas and stakeholder engagement. Thirdly, we look at the problem from different perspectives taking different

stakeholders into account and applying the five ethical principles. Fourthly, we come to a decision (action) about the shape of the use case to be implemented. Fifthly, we need to revisit the outcomes of that implementation decision and consider unintended consequences. This is especially needed if the Google principle of “Avoid creating or reinforcing unfair bias” (Pichai, 2018) is to be enacted, for example, in situation where algorithms learn from new data (avoiding bias is also a requirement of the GDPR).

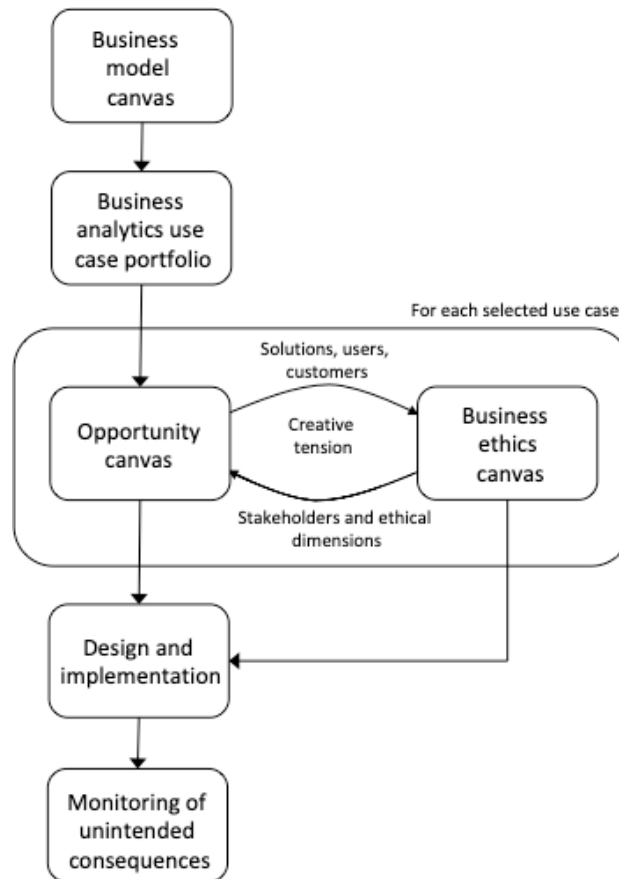


Figure 5: Enhancing the business analytics methodology (BAM) with an ethical dimension

The inclusion of stakeholders in the BEC helps move attention away from shareholder value and a narrow view of what constitutes business value. The origins of stakeholding can be traced back to the Stanford Research Institute who in 1963 defined stakeholders as “those groups without whose support the organization would cease to exist” (reported by Freeman (1984), p. 31). In applying this definition to analytics we recognize the need for analytics projects to move away from a unitary view in which a single stakeholder group is privileged (typically the shareholder) with a focus on a one-dimensional measure of success (such as profit maximization). Systems thinking

embraces stakeholding as a response to complexity and the need to consider the environment in which the organization operates, involving stakeholders as participants rather than treating them as constraints (Ackoff 1974a). Mitroff & Linstone (1993) define stakeholders as: “any individual, group, organization, or institution that can affect as well as be affected by an individual’s, group’s, organization’s, or institution’s policy or policies” (p. 141). The role of stakeholders has gained considerable attention in the OR literature (see de Gooyert et al., (2017) for a comprehensive review of the OR stakeholder literature) and in particular from a systems thinking perspective (e.g., see Vidgen (1997) and Wang et al., (2015), who both propose a method for stakeholder identification using soft systems). The inclusion of stakeholder identification and analysis techniques within the analytic is a fundamental plank in building a platform for exploration of the ethical dimension.

The research is not without limitations. First, it is based on a single action research project in a single organization. Further work is needed, therefore, to apply and develop the BEC in a broad range of organizations and contexts, together with a systematic assessment of the method’s efficacy, efficiency, and effectiveness. Second, while a BEC has been created for one use case (with demonstrable value add), embedding the BEC as a standard part of the analytics development process is a broader undertaking that will require senior management engagement and support. Third, it might be argued that our use case example, DaysOut, looks largely uncontentious at a first glance. However, the structured analysis afforded by the BEC demonstrates that ethical issues are indeed present in this use case and we would caution organizations making decisions about the ethical acceptability of analytics projects without conducting a systematic analysis first. Fourth, we recognize that it is difficult to publish research and accounts of practice in the area of analytics ethics due to the sensitivity of the subject matter - even though that research itself may lead to ethical outcomes. While EuroTravel rejects many potential analytics applications out of hand as being, *prima facie*, ethically unacceptable, it is still difficult to discuss the ethical issues arising from such cases due to concerns about how these discussions might be interpreted.

9 SUMMARY

With the rise of business analytics and AI the need to include an ethical dimension in the work of OR practitioners and data scientists is both urgent and important. Organizations operating in complex stakeholder environments, such as social media platforms, healthcare, and social care will likely be candidates for early adoption of the BEC technology as they are most likely to feel the pain of failing to take account of ethical aspects of analytics. While many of these complex environments will be public sector and not-for-profit organizations we expect ethical issues to come to the forefront for commercial organizations, too, particularly following the introduction of regulation such as the GDPR. Indeed, the justice and rights elements in the BEC address squarely the fundamental tenets of the GDPR: a right to privacy, a right to an explanation, and the removal of bias.

However, to be effective the exploration of ethical issues needs to be embedded in the organization's analytics development process and become a standard and everyday part of the OR and data science practitioner's toolkit. To address this challenge we extended the business analytics methodology developed by Hindle and Vidgen (2018) with the introduction of a business ethics canvas. The ethics canvas is intended to enhance the exploration of ethical issues and value creating opportunities rather than be used as a decision tool. Consequently, we are not proposing quantification leading to tables, scales, and scores. The aim of the ethics canvas is to uplift and de-risk ethical decision-making; as such it is using the Markkula framework as an ethical solution *exploration* tool rather than an ethical *decision* engine.

We have argued that value and ethics are inseparable. Rather than see ethics as an overlay to the analytics development process, in which value and ethics are separated, we see value and ethics as facets of the same problem: that is, value and ethics are intertwined and inseparable. However, when things go wrong, as with the Cambridge Analytica case (Greenfield, 2018), then they do indeed take on the appearance of being separable as the ethical issues come into sharp contrast. While ethics and value may be inseparable this does not mean that they coexist without tension. A particular challenge is the alignment of ethical priorities with commercial goals, which may be grounded in aggressive growth targets and incentive structures closely tied to commercial outcomes. Further research into managing the creative tension between analytics ethics and analytics pay-off (whether financial or non-financial) is needed.

REFERENCES

- Ackoff, R. (1974a). *Redesigning the Future: A Systems Approach to Societal Problems*. John Wiley & Sons: New York.
- Ackoff, R. (1974b). The social responsibility of operational research, *Operational Research Quarterly*, 25(3):361-371.
- d'Alessandro, B., O'Neil, C., and LaGatta, T. (2107). Conscientious Classification: A Data Scientist's Guide to Discrimination-Aware Classification. *Big Data*, 5(2): 120-134.
- Barocas, S., and Boyd, D. (2017). Computing Ethics: Engaging the Ethics of Data Science in Practice. *Communications of the ACM*, 60(11): 23-25.
- Baskerville, R., and Wood-Harper, A. T. (1996). A Critical Perspective on Action Research as a method for Information Systems Research. *Journal of Information Technology*, 11(3): 235-246.
- Brandom, R., (2018). New Toronto Declaration calls on algorithms to respect human rights. *The Verge*, 16 May. Retrieved May, 16 2018 from <https://www.theverge.com/2018/5/16/17361356/toronto-declaration-machine-learning-algorithmic-discrimination-rightscon>
- Caywood, T. E., Berger, H. M., Engel, J. H., Magee, J. F., Miser, H. J. and Thrall, R. M. (1971). Guidelines for the practice of operations research, *Operations Research*, 19, 1127-37.
- Checkland, P. (1981). *Systems Thinking, Systems Practice*. Wiley, Chichester.
- Checkland, P. and Holwell, S. (1998). Action Research: its nature and validity. *Systemic Practice and Action Research*, 11, 9-21.
- Checkland, P. and Poulter, J. (2006). *Learning for Action: A Short Definitive Account of Soft Systems Methodology and its Use for Practitioners, Teachers and Students*. Wiley: Chichester.
- Checkland, P. and Scholes, J. (1990). *Soft Systems Methodology in Action*. Wiley, Chichester.
- Churchman, C. W., (1968). *The Systems Approach*, Delacorte Press, New York.
- Churchman, C. W., (1979). *The Systems Approach and Its Enemies*. Basic Books, New York.
- Churchman, C. W., (1970). OR as a profession. *Management Science*, 17(2): B32.
- Churchman, C. W., (1971). *The Design of Inquiring Systems*, Basic Books, New York.
- Davison, R., Martinsons, M., and Kock, N. (2004). Principles of Canonical Action Research. *Information Systems Journal*, 14: 65-86.
- Dinsmore, T., (2017). How GDPR affects data science. *KDNuggets*, July 2017, <https://www.kdnuggets.com/2017/07/gdpr-affects-data-science.html>
- Eden, C. and Huxham, C. (1996). Action Research for Management Research. *British Journal of Management*, 7: 75-86.

- Fellowship for operational research (1974). *Professional conduct: guidance for Fellows of operational research*, London.
- Freeman, R. E., (1984). *Strategic Management, a stakeholder approach*. Pitman, Massachusetts.
- Gass, S. I. (1987). Managing the model process: a personal perspective. *European Journal of Operational Research* 31: 1–8.
- Gass, S. I. (2009). Ethical guidelines and codes in operations research. *Omega*, 37, 1044–1050.
- de Gooyert, V., Rouwette, E., van Kranenburg, H., and Freeman, E., (2017). Reviewing the role of stakeholders in Operational Research: A stakeholder theory perspective. *European Journal of Operational Research*, 262(2): 402-410.
- GOV.UK, (2007). *Guidance: Universal ethical code for scientists*. 12 September 2007. <https://www.gov.uk/government/publications/universal-ethical-code-for-scientists> (accessed 10 June 2018).
- Greenfield, P., (2018). The Cambridge Analytica files: the story so far. *The Guardian*. 26 March 2018. <https://www.theguardian.com/news/2018/mar/26/the-cambridge-analytica-files-the-story-so-far>
- Hindle, G. A. (2011). Teaching Soft Systems Methodology and a Blueprint for a Module. *INFORMS Trans Ed*, 12(1), 31-40.
- Hindle, G., and Vidgen, R. (2018). Developing a business analytics methodology: a case study in the food bank sector. *European Journal of Operational Research*, 268(3): 836-851.
- INFORMS (2108). INFORMS Ethics Guidelines. <https://www.informs.org/About-INFORMS/Governance/INFORMS-Ethics-Guidelines> (accessed 10 June 2018)
- Keeney, R. L. (1994). Using values in operations research. *Operations Research* 42, 793–813.
- Kitchin, R. (2017). Thinking critically about and researching algorithms. *Information, Communication & Society*, 20 (1): 14–29.
- Koch, T. (2000). We live in a city not in a study, *ORMS Today* 27(2), 16-17.
- Mann, G. and O'Neil, C. (2016). Hiring algorithms are not neutral. *Harvard Business Review*, 9 December. <https://hbr.org/2016/12/hiring-algorithms-are-not-neutral> (accessed 9 June 2018).
- Markkula (2018a). *A Framework for Ethical Decision Making*. Markkula Center for Applied Ethics. <https://www.scu.edu/ethics/ethics-resources/ethical-decision-making/a-framework-for-ethical-decision-making/> (accessed 10 June 2018).
- Markkula (2018b). *Making an Ethical Decision: a practical tool for thinking through tough choices*. <https://www.scu.edu/ethics/ethics-resources/ethics-app/> (accessed 10 June 2018).
- Mason, R.O., (1994). Morality and models. In: Wallace, W. (Ed.), *Ethics in Modelling*. Pergamon, Oxford, UK, pp. 183–194.

- Midgley, G. and Munlo, I. (1998). The theory and practice of boundary critique: developing housing services for older people. *Journal of the Operational Research Society*, 49, 467-478.
- Midgley, G. and Ochoa-Arias, A. E. (Eds) (2004). *Community Operational Research*. Kluwer, New York.
- Mitroff, I., & Linstone, H. (1993). *The Unbounded Mind, breaking the chains of traditional business thinking*. Oxford University Press, New York.
- Mukherjee, S., (2017). This New AI Can Detect a Deadly Cancer Early With 86% Accuracy. Fortune, 30 October. <http://fortune.com/2017/10/30/ai-early-cancer-detection/> (accessed 9 June 2018).
- Newell, S. and Mirabelli, M. (2015). Strategic Opportunities (and Challenges) of Algorithmic Decision-making: A Call for Action on the Long-term Societal Effects of 'Datafication'. *Journal of Strategic Information Systems*, 24(1): 3-14.
- Omerod, R. J. (2002). On the nature of OR: taking stock. *Journal of the Operational Research Society*, 53, 475-491.
- Omerod, R. J. and Ulrich, W. (2013). Operational Research and ethics: A literature review. *European Journal of Operational Research*, 228, 291-307.
- OR Society (2108). The OR Society - Statement of Ethical Principles. <https://www.theorsociety.com/Pages/Society/Ethical.aspx> (accessed 10 June 2018)
- Osterwalder, A. and Pigneur, Y. (2010). *Business Model Generation*. Hoboken, NJ: John Wiley and Sons, Inc.
- Patton, J., (2016). *Opportunity Canvas*. 9 September 2016. <https://jpattonassociates.com/opportunity-canvas/> (accessed 9 June 2018).
- Perkins, A. (2018). May to pledge millions to AI research assisting early cancer diagnosis. *The Guardian*, 21 May. <https://www.theguardian.com/technology/2018/may/20/may-to-pledge-millions-to-ai-research-assisting-early-cancer-diagnosis> (accessed 9 June 2018)
- Pichai, S. (2018). AI at Google: our principles. *The Keyword*. 7 June 2018. <https://blog.google/technology/ai/ai-principles/>
- Rogers, E. M., (2003). *Diffusion of Innovations*, (5th edition). Free Press, New York.
- Rosenhead, J. (1976). Some further comments on "The social responsibility of operational research". *Operational Research Quarterly*, 27, 266-277.
- Rosenhead, J., and Mingers, J., eds. (2001). *Rational analysis for a problematic world revisited*. (2nd edition). John Wiley and Sons, Chichester.
- Rosenhead, J. and Thunhurst, C. (1982). A materialist analysis of operational research. *Journal of the Operational Research Society*, 33, 111-122.
- Schwarz, T. (2005). Teaching Ethics and Computer Forensics: The Markkula Center for Applied Ethics Approach. *Proceedings of the 2nd annual conference on Information security curriculum development*, September 23-24, Kennesaw State University, Georgia, US. Pages 66-71.

- McNamara, A., Smith, J., and Murphy-Hill, E. (2018). Does ACM's Code of Ethics Change Ethical Decision Making in Software Development? *ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*, Lake Buena Vista, FL, 4-9 November 2018.
- Statt, N., and Vincent, J. (2018). Google pledges not to develop AI weapons, but says it will still work with the military. *The Verge*. 7 June.
<https://www.theverge.com/2018/6/7/17439310/google-ai-ethics-principles-warfare-weapons-military-project-maven> (accessed 10 June 2018).
- Susman, G. I. and Evered, R. D. (1978). An Assessment of the Scientific Merits of Action Research. *Administrative Science Quarterly*, 23:582-603.
- Tornatzky, L. G., and Fleischer, M., (1990). *The Process of Technology Innovation*. Lexington Books, Lexington, MA.
- Ulrich, W. (1987). Critical heuristics of social systems design. *European Journal of Operational Research*, 31, 228-276.
- Ulrich, W. (2006). Critical pragmatism: a new approach to professional and business ethics. In: Zsolnai, L. (Ed.), *Interdisciplinary Yearbook of Business Ethics, vol. I*. Peter Lang Academic Publishers, Oxford, UK, and Bern, Switzerland, pp. 53–85.
- Ulrich, W. (2007). Philosophy for professionals: towards critical pragmatism. *Journal of the Operational Research Society*, 63, 1228-1247.
- Vidgen, R. (1997). Stakeholders, soft systems and technology: separation and mediation in the analysis of information system requirements. *Information Systems Journal*, 7: 21-46.
- Wang, W., Liu, W., and Mingers, J. (2015). A systemic method for organisational stakeholder identification and analysis using Soft Systems Methodology (SSM). *European Journal of Operational Research*, 246(2): 562-574.
- Zhu, K., Kraemer, K. L., and Xu, S. (2003). Electronic business adoption by European firms: a cross-country assessment of the facilitators and inhibitors. *European Journal of Information Systems*, 12(4): 251-268.
- Zucoloto, G., and Nogueira, M., (2016). Challenges for Innovation Due to Firm Size: The Case of Brazilian Industrial Firms, pp. 356-376. In: Al-Hakim, L., Wu, X., Koronios, A., and Shou, Y. *Handbook of Research on Driving Competitive Advantage through Sustainable, Lean, and Disruptive Innovation*. IGI Global.