Business Analytics in Service Operations—Lessons from Healthcare Operations

Opher Baron,
Rotman School of Management, University of Toronto
opher.baron@rotman.utoronto.ca

Abstract

We present an expanded framework for the use of business analytics in projects. To the commonly used **descriptive**, **predictive**, and **prescriptive** analytics, we add **comparative** analytics, wherein we compare the performance of systems under different interventions. This framework provides a conceptual roadmap for the implementation of business analytics projects. We then demonstrate this framework using recent operations research literature on analytics in healthcare, summarizing papers focusing on one of these aspects. Next, we discuss queue mining as an example of theory and practice illustrative of these aspects. We conclude there is room for further work by operations researchers and management scientists within business analytics projects generally and the healthcare industry more specifically. We argue future work should consider both theory and practice, especially within prescriptive analytics projects, where analysis through the lens of operations research and management science is imperative. We provide some thoughts on the current and future state of operations research and management science in business analytics.

1. Framework for Business Analytics

Recent advancements in computing and analytics suggest that both the practice and theory of service operations have been dramatically changed and will continue to shift in the near and more distant future. The availability of data and efficient tools for visualization and analysis opens several possibilities for research and practice in service operations. (Cohen 2018) presents a modern view of the impact of big data on service operations. Business analytics in service operation is very broad as it includes, for example, revenue management for services, shared economy, and post-purchase services.

To improve our understanding of the different paths towards implementing analytics, we surveyed the literature on analytics in services. As this is a vast literature, we only discuss the literature on business analytics in healthcare operations, with a focus on work related to queueing in this sector. Importantly, the trends and approaches discussed are applicable to business analytics in many services, as well as in other sectors.

A common theme of analytics is that it uses quantitative analytical tools, such as statistics, econometrics, computer science, and operations research, to analyze specific problems. Moreover, these analytic tools are used to analyze and model data. Finally, business analytics uses analytics to support decision making in businesses. Similar to the decision-centric definition of analytics in (Rose 2016), we offer the following concise definition:

- **Business Analytics**: Usage of data-driven analytical methodologies to support decisions.

1.1 Defining Descriptive, Predictive, Comparative, and Prescriptive Analytics

Quantifiable performance measures of different business objectives for different business units are often required to support business decisions. Therefore, every business analytics project benefits from familiarity with the application and its modeling. (Lustig et al. 2010) in terms of discuss descriptive, predictive, and prescriptive analytics. There are several definitions for these different aspects of analytics. Nevertheless, based on our literature review, we find these terms do not capture the complexity of how business analytics works and are thus insufficient. Specifically, business analytics should differentiate between existing systems and systems with different interventions, i.e., under different decisions. To emphasize this important distinction, we add the term *comparative analytics*. Figure 1 depicts these four aspects of business analytics projects and their intersection. We define them as follows:

- **Descriptive analytics** describes past performance of existing systems.
- **Predictive analytics** predicts future performance of systems.
- **Comparative analytics** compares performance of systems under different interventions.
- **Prescriptive analytics** prescribes interventions to improve future performance of systems.

Figure 1: Descriptive, Predictive, Comparative, and Prescriptive Business Analytics

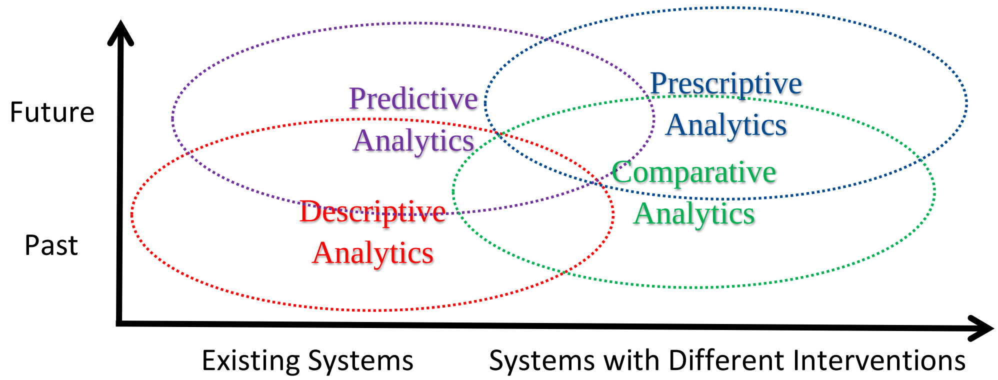

The aim of ***descriptive analytics*** is to describe and understand systems' past performance. In the lingo of statistics, descriptive analytics studies performance in-sample. While there is full certainty of what happened, there may be many different interpretations of the best way to describe it. Descriptive analytics is focused on data analysis, reporting, and visualization. A common descriptive analytics problem is the clustering of patients based on their comorbidities. Some of the main challenges in descriptive analytics are related to the interpretation of available records and the problem of dealing with missing and inaccurate data. Importantly, appropriate descriptive analytics is the building block for predictive, comparative, and prescriptive analytics. For example, in some cases, the right visualization makes the best interventions clear.

The aim of ***predictive analytics*** is to predict systems' future performance. It requires inductive reasoning to yield out-of-sample results. Predictive analytics often ignores the causal effect, or assumes it is clear. Therefore, predictive analytics captures the impact of interventions on performance by changing the values of inputs into the model. A common predictive analytics problem is demand forecasting. Predictive analytics is focused on finding trends and relationships in data that the user *believes* will continue to hold true in the future. The main challenges in predictive analytics are thus related to these beliefs. Two important concepts of predictive analytics that substantiate such beliefs are: (i) modeling the physical system - if users trust the model and the insights gained from its analysis, they are more likely to trust its predictive results, and (ii) testing how a prediction model developed on training data (in-sample) performs on prediction data (out-of-sample). Good performance on the prediction data demonstrates that the analysis has predictive power. Predictive analytics models may suffer from overfitting: a complex model that explains all training data may be useless out-of-sample. Therefore, simple models that capture only a few important effects are often preferable. Finally, appropriate predictive analytics is often a building block for prescriptive analytics.

The aim of ***comparative analytics*** is to compare the performances of different systems under different interventions. Such a comparison can be done either retrospectively or prospectively, using inductive reasoning to yield results under different interventions. Retrospective comparative analytics focuses on expressing performance that would have occurred under different interventions. This analytics manipulates existing data as if different interventions were made. Two examples of retrospective comparative analytics are expressing counterfactuals and benchmarking. Expressing counterfactuals, e.g., measuring the impact of a completed project, is an example of retrospective comparative analytics when a single business unit is analyzed. Benchmarking, e.g., comparing the performance of different branches (exposed to different intervention), is an example when several units are analyzed. The main challenges in retrospective comparative analytics are establishing causality, using the appropriate causal inference tools, and ensuring data on a specific system under different interventions do not exist. (Even in benchmarking, we only have data on each system under its own interventions.) Prospective comparative analytics focuses on predicting performance that would occur under different interventions. It can be thought of as performance analysis in the future. Thus, just like predictive analytics, it requires inductive reasoning to out-of-sample results. After establishing causality, such analysis is often achieved by combining models and data. An important tool in prospective comparative analytics is simulation, as it permits the estimation of the impact of different interventions on relevant performances without the cost of implementing such changes in practice. (Simulation is also used for retrospective comparative analytics as part of verifying and validating the simulation model with managerial judgement.) Proper simulation models are built, validated, and verified using existing data. Simulation models may contain specific machine learning (ML) components to mimic building blocks that generate data. Another approach to comparative analytics is a controlled experiment (or A/B testing in the context of user experience) that records enough data on systems exposed to different interventions. To some extent, controlled experiments transform comparative analytics questions to descriptive analytics ones. Comparative analytics also emphasizes the importance of planning data collection to improve analytics in the future. In some cases when interventions are simple, e.g., go no-go with a project, the performance evaluation supplied by comparative analytics is sufficient to prescribe interventions. Note that even when prescriptive analytics is challenging, e.g., in the context of more complex interventions, comparative analytics serves as an essential building block.

The aim of ***prescriptive analytics*** is to prescribe interventions that would improve systems' performance. The holy grail of analytics is finding such a novel intervention. Prescriptive analytics requires inductive reasoning to out-of-sample results under different interventions. Decision makers who use prescriptive analytics properly to develop recommendations no longer require a leap of faith when they implement them, as they consider and understand the important factors that affect relevant performance measures. Indeed, prescriptive analytics that is well supported by descriptive, predictive, and comparative analytics not only allows the decision maker to reduce risks from new interventions but also allows her to manage these risks. Fundamental prescriptive analytics problems have complex relationships between interventions and performance. Such complexities are especially common when these relationships cannot be measured without making specific assumptions or when they are subject to high uncertainty. In these cases, tools that model the relationships are essential to prescribe effective decisions. Prescriptive analytics can use robust optimization guided by existing data to compare and predict the potential impact of different interventions on performance. The challenges in prescriptive analytics include the challenges of descriptive, predictive, and comparative analytics.

We note that these four aspects of analytics are helpful conceptually but are typically too narrow to encompass the scope of any analytics project in practice. For example, a descriptive analytics project can be thought of as the study of history. History is interesting, but much of the motivation to study it comes from the light it sheds on future events. Similarly, descriptive analytics is interesting, but much of the motivation to perform it comes from the light it sheds on future effective management interventions. As another example, much of the focus of predictive analytics, including artificial intelligence and ML, is on expressing how changes to some variables affect dependent variables, i.e., a relevant performance measure. Once managers trust the predictive model and the causal effects behind it, they can compare the effect of different interventions on the dependent variable and may prescribe interventions.

1.2 Phases in Business Analytics Projects: From Descriptive to Prescriptive Analytics

Another important factor in the conceptualization of these four aspects of business analytics is their timing within the progression of a business analytics project, as depicted in Figure 2. Most business analytics projects progress over time, from descriptive to predictive, comparative, and prescriptive phases. Nevertheless, business analytics may focus on one of these and use tools and models tailored to pursue it.

Consider the phases in a business analytics project focusing on a congested service system. The descriptive analytics phase of the project will examine past observed performance measures of the system and address such questions as: "Is the process performing as designed? Do the main flows in practice tightly correspond to the intended flows? Are bottlenecks where we expect them to be?" In a hospital emergency department (ED), descriptive analytics can lead to observations such as: demand is highest on Mondays from 4pm-8pm, chest pain is the most common complaint, and performance measures (such as length of stay, waits, and time to physician initial assessment, all for different patient types) depend on the ED census and the nurse to patient ratio.

The predictive analytics phase will consider future performance of the system and address such questions as: "What throughput can we expect? When we expect the system to be the busiest? What would the state of the system be in the future, in terms of arrival processes and distributions of performance measures?" In an ED, predictive analytics would predict performance measures, for example, using a simultaneous equation model, as a function of the time (e.g., Monday, 4pm-8pm, winter), patient mix (e.g., percent of patients with chest pain), the ED census, the nurse to patient ratio, and number of busy beds in the wards.

Figure 2: Timing of Descriptive, Predictive, Comparative, and Prescriptive Business Analytics

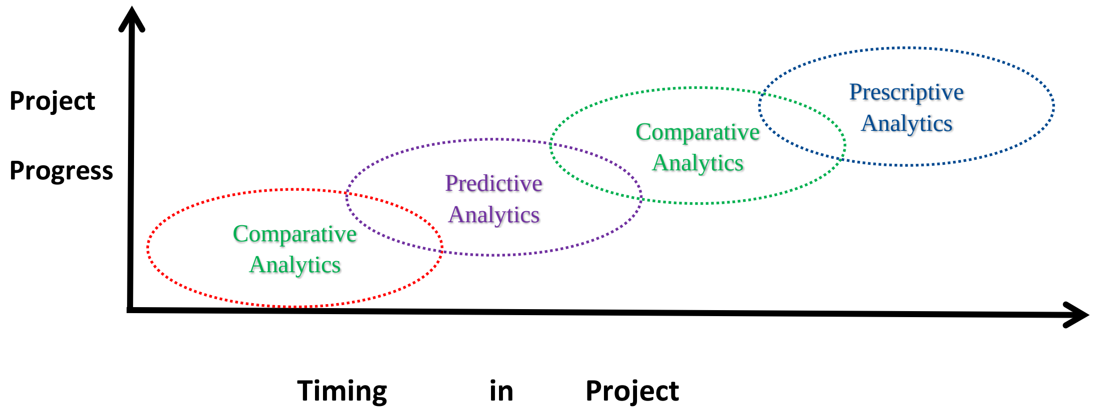

The comparative analytics phase will consider how various changes would have impacted (retrospective) or will impact (prospective) performance measures and address such retrospective questions as: "What would the performance of the system have been if a specific improvement project had not been implemented (i.e., express counterfactuals)?" Prospective analytics will address such questions as: "What impact can be achieved by adding new resources to the system? What impact can be achieved from workflow changes?" In an ED, comparative analytics would consider various new workflows and staffing changes. New workflows can include adding a dedicated track for patients with chest pain, reducing the census and waiting times, or having a physician perform triage, reducing the time to physician initial assessment. Staffing changes can include different shift schedules for nurses and doctors and/or hiring nonmedical personnel to ease the load of medical personnel during rush hours. The analysis of such changes could use a causal model or simulation (calibrated with data) to capture their respective impacts on performance measures.

At the prescriptive analytics phase, the expert will consider how to best manage future performance of the system and will address such questions as: "What is the best staffing procedure to follow? What is the best customer routing?" In an ED, prescriptive analytics would search for the optimal nurse and doctor shift schedules and provide guidelines on when to call and pay for extra medical or nonmedical personnel. The creation of this flexible staffing process requires the knowledge obtained from the descriptive, predictive, and comparative phases of the project. The prescriptive model aiming at cost minimization subject to constraints on performance measures would be informed by current data (e.g., census), an arrival prediction model, and established dependencies of performance measures in the patient mix, number of busy beds in the wards, nurse to patient ratio, and so on.

Good progress through the descriptive, predictive, and comparative analytics phases of the project will help decision makers understand the possible and preferred trade-offs among different performance measures and the cost of achieving them. Such an understanding is essential to ensure the prescriptive analytics phase of the project is effective. Without such an understanding, analysts may full into the trap of optimizing the performance, given a predictive model, a simplistic approach that is often inappropriate.

1.3 Outline

In the next sections, we stratify our discussion of analytics research in healthcare along descriptive, predictive, comparative, and prescriptive lines, with the caveat that papers focusing on, say, descriptive analytics often include predictive, comparative, and prescriptive analytics as well. We then discuss research on queue mining and show how this line of research progresses from descriptive to prescriptive analytics. We conclude with thoughts on the future contributions of operations research and management science (ORMS) to business analytics.

2. Descriptive Analytics

Descriptive analytics traditionally focused on studying well-stated existing problems. During the study, we often found alternative relevant problems, but data analysis was thought as a part of solving a given problem. In contrast, (Simchi-Levi, 2014) suggests modern data analysis allows us to define and find hidden problems and to learn from best practices.

The first step in learning from practice is being able to map practical phenomena. Some of this mapping raises the need for theoretical models and their study. Below, we discuss descriptive analytic work on cost estimation, the interaction between service and waiting times, scheduling policies, and patients' service path.

2.1 Estimating the Cost of Delay

One of the most fundamental questions in services is "What is the cost of delay?" Indeed, the cost of providing service or delaying it is often not easy to express. Access to data on decision making allows us to impute such cost. That is, rather than calculating the cost of delay, we can compare delays in practice and reverse engineer from them the delay cost that would make the practice optimal, i.e., the implied delay cost of the current practice. We believe similar approaches can be used to estimate the imputed cost and implication of different decisions.

(Ding et al. 2019) addressed this question for EDs. They investigated patient routing and prioritization and how these are affected by patients' triage levels. They considered patient level data for 20 months of visits in four EDs and reverse engineered the delay cost the hospitals attributed to patients based upon their observed prioritization. They discovered that different triage levels imply different delays, and the current practices reflect a piecewise linear and concave delay cost. Specifically, while within a triage level, service is typically administered in a first come first serve fashion (FCFS), once patients' wait is too long (e.g., above the wait time required), their priority decreases. Moreover, it does not appear that the level of complexity of the treatment (measured as total service time) impacts patients' priorities. This study does not relate specific delay cost to different patients, but quantifies the relative cost of wait for different patient classes. Understanding these relative costs allows decision makers to improve the incentives provided to service providers. The example provided highlights the inadequacy of using thresholds on waiting times as the sole measure of the performance of EDs.

2.2 Service and Wait Times

Analyzing service systems requires an understanding of arrival, service, and wait times. Standard assumptions of queueing theory are that arrival and service times are independent of each other, and the realized delay of previous and current customers. However, several studies show that service times and exceptional services impact wait times and future service times.

The study of the impact of delay in entering ICU on the service time, captured via length of stay, in (Chan, Farias, Escobar, 2016) showed that, for some patient types, there is a positive correlation between delay and service times. As longer service times increase congestion and delay, a dangerous positive feedback in ICU systems is created. These researchers developed theoretical and simulation models to show that the positive correlation between delays and ICU service time increases congestion. This finding suggests ICU managers should not ignore this correlation and the feedback it creates.

(Deo and Aditya, 2019) studied an eye care provider in south India with a transient service network. That is, all people arriving during a day are served and leave that day. (This differs than, say, an ED, where the next shift may complete work left over from a previous one.) They investigated control of resources and arrivals in view of balancing wait times with hours and effort of work. They used four months of data for visits at several branches and different doctors. Their patient level regression analysis showed that clinics work more slowly at the beginning of the day and faster later in the day. The motivation for this policy is that the beginning of the day is used to fill all buffers, such that later in the day the faster system has no idling within the network. Such a change of speed implies busier systems would have worked faster (than less busy systems) during their slow period. An important implication of this policy is that patients who arrive earlier in the day typically wait longer than patients who arrive later during the day. The researchers' data-calibrated simulation showed that such a change in work speed reduces the time to complete all work and the average length of stay compared to fixed work speed.

Another interesting finding was reported by (Bavafa and Jónasson 2020). They found that ambulance teams who treated disturbing cases, e.g., with death, slowed down their work during the rest of the shift as a method of coping. The slowdown occurred more in less standard parts of their service and depended on the team's years of experience.

2.3 Scheduling Policies

Scheduling policies that dictate customers' priority and routing represent interventions with a potentially strong impact on performance in service systems. In this section, we discuss how best practices inform theory on potential scheduling policies.

The data analysis reported in (Baron et al. 2017) for a firm providing annual healthcare checkups enables learning from practice. The annual checkup process reported by these researchers involved patients visiting several (12-20) different service centers, for such procedures as blood tests, stress tests, hearing tests, reviews with a doctor etc. The transaction level data on waiting and service times showed some “strange” practices. Specifically, in many cases, patients were waiting for a service center while it was idle, or to use the opposite terminology, a service center was waiting for a patient while she was idle. The authors dubbed this practice “strategic idling”.

Queueing theory establishes that non-idle policies are optimal in reducing the total wait and sojourn times, so this documented practice of strategic idling raises questions about possible benefits. (Baron et al. 2017) argued the psychology of queues implies that the perception of wait may differ from the objective wait. In the context of patients (customers) moving through several service centers, the perception of wait is affected by two aspects. The first is the macro measure of total wait in the system, i.e., aggregated over all service centers; the second is the micro measure of wait for any specific center. When customers experience an increase in waiting for a specific center, they become unhappy. In this case, an in-depth study of the data showed that while, as predicted by queueing theory, eliminating strategic idleness would reduce the total wait, it would significantly increase the percentage of long waits at specific centers. Indeed, (Baron et al. 2014) established that strategic idling could improve the micro level and thus perceived service level.

The suggestion that strategic idling may improve the perceived service level raises the question of how to implement it in practice. (Baron et al. 2014) presented a family of threshold-based controls to balance long waits among several sequential centers. In later work, (Baron et al. 2017) extended this family to allow such a balance in a general network of service centers.

In-depth descriptive analytics pointed to the phenomenon of strategic idling, a previously unknown practice and broadly considered to be counter-productive. To understand the motivation for strategic idling, the authors used a comparative analytics approach: they found simulation and queueing analysis of stylized models helped to quantify the benefit of strategic idling. The insights generated by the comparative analytics supported prescriptive analytics. Overall, the business analytics approach demonstrated in (Baron et al. 2014) and (Baron et al. 2017) allowed a solution to be prescribed: a systematic and effective method to implement strategic idling to improve the service quality as perceived by customers. Importantly, this work highlights the interaction among the different aspects of analytics discussed in Section 1.

In many cases, the planned scheduling policy and the realized one differ. This possibility has been investigated by both (Ibanez et al 2018) and (Chan et al 2021) in the context of doctors who read radiological (e.g., CT, X-rays) images off-site. In (Ibanez et al 2018), the doctors were assigned random images; they were expected to process these in an FCFS manner, and their compensation scheme was unclear. The researchers tested several hypotheses using transaction level data on teleradiology for over 2.75 million cases from July 2005 to December 2007. Their main conclusion was that the radiologists deviated from the FCFS scheduling over 40% of time. These deviations were mainly motivated by doctors' wish to work on shorter jobs first or to batch similar tasks. Their econometric analysis used instrumental variables to show that such deviations harm productivity. They suggested the seeking time required to deviate may cause this degrading performance. Moreover, they found that more experienced doctors deviated less and were less harmed by deviation. However, their analysis ignored the implication of deviation in work on waiting time.

In contrast, in (Chan et al 2021), the doctors chose images from a common pool, were expected to process images in accordance to priority requirements (i.e., waiting times are relevant), and were compensated based on work completed. They used about 2.65 million cases processed by over 150 radiologists at close to 100 hospitals from January 2014 to July 2017. They found the imbalance between pay and workload could cause low priority jobs to overtake high priority ones. Using instrumental variables, they showed that waits decreased (increased) in the pay-to-workload ratio for low (high) priority jobs. Their results suggest the design of payments scheme should consider a "bang for the buck", or here, dollar per minute, compensation for jobs. When incentive structures ignore this strategy, service providers may change the order in which they process jobs.

2.4 Service Path

Much effort, both in literature and in practice, has been dedicated to understanding customers' lifetime value as it relates to customer pathways and customer segmentation. In the context of healthcare, customer pathways often depend on segmentation based on diagnostic stage, such as triage.

For example, in EDs, the first segmentation occurs at the triage level. Triage decisions and their connection to load was investigated by (Chen et al. 2020). They used patient level data for over 65,000 visits. They found an increase in load caused a significant increase in triage level assignment. Patients were triaged as more severe, and the common disposition was "to be admitted." These findings contrast with those of (Richardson 1998) for an Australian hospital. An important takeaway from (Chen et al. 2020) is that an increased load in the ED creates a positive feedback that further increases the load. Admitting more patients increases bed blocking from the wards, and this increases the load in the ED. Therefore, it is important to have accurate triage that is independent of load.

In healthcare, a best treatment pathway is often provided to doctors as a recommended course of action. (Diamant et al. 2014) considered what impacts patients' clinical pathway towards bariatric surgery, a common tool in fighting obesity. The process leading to bariatric surgery includes several pre-surgery steps, where attrition may occur. Using three years of data, the authors demonstrated that specific groups of patients may be less likely to complete the entire process. Better management of the process for these groups can improve healthcare outcomes.

3. Predictive Analytics

Much of predictive analysis is focused on forecasting. As forecasts support long- and short-term business planning, forecasting is one of the fundamental activities of any business. Social platforms and the Internet provide new data sources that, when properly combined with operational data, improve both accuracy and interpretability of forecasting models. Moreover, modern algorithms, based upon e.g., ML, lead to improved forecast accuracy.

We next discuss some research on cost prediction, on forecasting the spread of infectious disease, and on demand forecasting.

3.1 Forecasting Cost

Provisioning of service is costly, so cost management is extremely important in service systems, where cost of personnel is often a critical factor. The effective provisioning of service also requires paying costs in areas most likely to benefit the service. Moreover, investigating the trade-off of interventions, as a part of prescriptive analytics, requires an understanding of their respective costs. With these questions in mind, estimating the cost of service may be critical to business analytics projects.

An example of the use of modern tools and data in the context of cost forecasting is in (Bertsimas et al. 2008). These authors used classification trees and clustering algorithms to estimate healthcare costs for 200,000 out-of-sample people with medical insurance. They demonstrated that past cost helps predict future cost, and that specific medical information is helpful to estimate the cost of high-cost patients but is not helpful for other groups. This example nicely demonstrates the difference between a predictive model and a comparative one: past costs are correlated with the future cost, so they can be used to predict it; but as past costs do not cause the future cost, they cannot be directly used to analyze interventions to impact the future cost.

3.2 Forecasting the Spread of Infectious Diseases

Forecasting in healthcare includes predictions of the spread of infectious diseases and intensity of demand for emergency treatment such as cardiac arrest requiring defibrillators. We note that these examples require a spatial demand prediction. Moreover, such predictions are essential to support prescriptive analytics; for example, decisions on the location of defibrillators require demand forecasting. The outbreak of Covid-19 demonstrates the importance of spatial models for forecasting the spread of infectious diseases, for informing decision makers on the right social distancing level to recommend in different regions, and for designing relevant travel restrictions.

One of the most common infectious diseases is influenza. The combination of data on electronic health records and on past flu activity levels was traditionally used to provide temporal and spatial predictions of flu levels. In recent years, people have begun to search for medical information on-line. (Yang et al. 2017a) showed that a dynamic multivariate regression model that combines these data with Internet (Google) searches for relevant terms can improve predictions (at a 5% significance level). It provides an accurate enough forecast, i.e., forecast that can be acted upon, for four weeks rather than two weeks ahead of an increase in flu activity. Such improved accuracy increases the allowable response time and is equivalent to a similar reduction in lead time with major operational consequences. For example, flu shots can be delivered to flu-prone areas, given to populations at risk, and thus prevent a pandemic.

The approach in (Yang et al. 2017a) was extended in (Yang et al. 2017b) to predict the activity level of dengue fever. The authors aimed at providing a disease surveillance system to predict dengue outbreaks and support interventions. They combined traditional data with Internet searches and had good accuracy; their approach worked well in areas where dengue fever levels were significant.

3.3 Demand Forecasting

Forecasting demand for healthcare services, such as arrivals to an ED or for emergency services resulting from cardiac arrest, can be improved with an increase in data availability. Moreover, the availability of adequate responses, e.g., capacity of ED or location of responders for emergency services, can be optimized using data-driven forecasts. In this section, we discuss recent analytics work on such forecasting.

(Chan et. al, 2013) considered the location of defibrillators in Toronto. They showed that the while the spatial distribution of cardiac arrest has some correlation with population density, the correlation was not perfect. Thus, they recommended considering the location of past cardiac arrests when locating defibrillators. In more recent work, (Chan, Demirtas, and Kwon 2016) developed optimization models that also considered bystanders could employ defibrillators. That is, their forecasts jointly considered demand occurrence and the possible availability of a provider (although their model for the latter is not data driven). We note that forecasting both demand and servers' availability has attracted more attention in recent years with the increase in the sharing economy and matching; see Hu (2019) and references therein.

Demand rate for many services, including EDs, has patterns of seasonality, as well as weekly and daily fluctuations. Forecasting this nonstationary demand process is important, as it affects planning and day to day operation. (Chen et al. 2021) proposed a nonhomogeneous Poisson process with a rate that is the sum of several sinusoidal waves. They explained how to estimate the number of different waves and their rates and then demonstrated their methodology using data for 170,000 arrivals over two years at an ED in an academic hospital in the US. They predicted the arrival rates at the ED for both acute and non acute patients (they ignored prediction for emergency patients as these were rare). They also demonstrated how to use their method for staffing.

4. Comparative Analytics

Comparative analytics often benefits when there are several similar business units, such as hospitals. Investigation of implemented interventions in one unit can be improved by an understanding of the impact of this intervention in other units. Proper comparative analysis of a known intervention supports its extension to other business units. Such a comparison is useful to help manage service and wait times.

In this section, we discuss some research on cost effectiveness of interventions, interventions that impact the service discipline and waiting times, problems of resource matching in healthcare, focusing on blood transfusion, and the service path and its dependence on diagnostics. The work mentioned herein includes examples of both retrospective and prospective comparative analytics.

4.1 Cost Effectiveness

One of the most important inputs into the effective management of any operation is the cost benefit analysis of different interventions. The most common interventions in services are related to changes in capacity (e.g., moving shifts, adding capacity, or changing processing times) or processes (e.g., scheduling or changing working processes, such as providing vaccinations). Evaluating the impact of such interventions is important for both retrospective and prospective interventions. We next discuss recent analytics work on the cost effectiveness of large scale immunization.

(Hutton, So, and Brandeau, 2010) performed a cost-effectiveness analysis of prospectively providing free hepatitis B vaccine to children aged 1 to 19 years in China. To overcome the absence of counter-factual data, they developed a Markov model for disease progression and infections. This model helped to estimate the impact of vaccinations on cost, quality-adjusted life years, and societal perspective. To quantify these measures, they used publicly available data published by the Chinese Ministry of Health, in peer-reviewed Chinese and English publications, and by the GAVI Alliance. Their investigation and sensitivity analysis demonstrated that offering free hepatitis B vaccination for 150 million susceptible youth in China would be beneficial; it could reduce infections by 8 million and deaths due to hepatitis B by 65,000.

Brandeau with coauthors in a series of papers (Malloy et al. 2020 a, Claypool et al. 2020) presented cost effectiveness analysis using similar principles: developing disease infection models, calibrating these using data (often publicly available), and providing detailed sensitivity analysis of different interventions. (Malloy et al. 2020 b) discussed the importance of model structure in predicting effectiveness of medical interventions.

4.2 Service Discipline and Wait Times

Comparative analytics work in the context of improving wait times examines different service disciplines and/or routing policies. One of the most useful tools for this is simulation. In this section, we discuss recent analytics work on the impact of the service discipline on service times and thus on waiting.

One of the main challenges of analysing the impact of interventions is that counterfactual data do not exist. For example, for patients who have had an intervention, there are no data on their experience in the absence of this intervention. Therefore, it is hard to establish that the intervention caused an observed effect. (Song, Tucker, and Murrell, 2015) investigated the impact on length of stay in the ED of operating a pooled versus a dedicated doctors' queue. From their patient level data, they found that operating a dedicated queue reduced the length of stay by 17%. This contrasts with queuing theory's fundamental insight that a pooled queue should be better. Their study was retrospective; they could use a difference-in-differences approach to support their results because the intervention that changed the queuing system was implemented in only one part of the ED.

(Bayati et al, 2017) considered EDs where high acuity patients received a pre-emptive priority over low acuity ones. Their empirical evidence showed low-acuity patients increased the wait times of high-acuity patients. This finding was clear in EDs with non pre-emptive priority discipline, but it contradicts the theoretical predictions of pre-emptive priority queues, see e.g., (Wang, Baron, and Scheller Wolf, 2015). (Bayati et. al 2017) estimated this increase using data from several hospitals and a quasi-randomized experiment, i.e., a prospective study, implemented in one hospital. In their experiment they slightly changed the wait time data shared with incoming patients. These data changed non acute patients' decision on whether to join the ED or leave it without being seen. This quasi-randomization corrected the existing bias in data and supported their findings. They argued low-acuity patients delay high-acuity patients due to pre-triage delay and set-ups required between pre-emptions (e.g., changing sheets).

The impact of delay announcements on hospital wait items was also studied in (Dong, Yom-Tov, and Yom-Tov 2019). Their retrospective study used three months of data on wait times posted by over 200 EDs in the US and data from a web search for ED delay information. They showed that posting wait times balanced the waits among hospitals, and this impact decreased with the distance between hospitals. Using difference-in-differences, they compared duration when EDs did not report their wait times with duration when wait times were reported. They found the balancing of wait times was impacted by information sharing; i.e., wait times were higher when no such information was provided.

Another important decision in EDs that helps control wait times is related to patients' diversion. This problem was investigated by (Xu, and Chan 2016). These authors considered how to use arrival forecasts to divert patients from an emergency department. The traditional, static policy diverts patients whenever their total number in the ED is above a threshold. They showed that, given practical forecasts' accuracy and horizon, a dynamic policy is better than a static one in reducing waiting times and patients' abandonment (leaving without being seen). Their diversion decisions combined this threshold with additional diversions of patients that initiated long busy periods based on the forecast. The authors motivated this policy by analyzing stylized models and then demonstrated its efficacy using a simulation based on data from ED records in the SEEStat database (SEE-Center 2009) for 2004. Their retrospective simulation model compared the impact of different diversion decisions on waiting time and patients' abandonment. They demonstrated that the use of practical forecasts could improve the performance of busy EDs.

The Erlang R model, where R stands for re-entrant, was introduced and analyzed by (Yom-Tov and Mandelbaum 2014). Their motivation for the study of this queueing model stemmed from healthcare applications such as radiology and EDs, where patients often return to the service provider after going through some additional tests. They showed that in steady state, the Erlang R and Erlang C (i.e., the standard M/M/C queue) models were equivalent (after appropriately changing the arrival or service times). Nevertheless, in the presence of time dependent arrivals, staffing recommendations based upon the Erlang R model led to significantly more stable performances (measured as probability of wait). They retrospectively demonstrated these results via the detailed simulation from (Marmor and Sinreich 2005), calibrated using empirical data from an ED. They enhanced their results using fluid and diffusion models and showed the Erlang R model was also useful in planning for staffing during a chemical event mass casualty drill.

The simulation models that are common in the ORMS community, such as the detailed simulation models in (Marmor and Sinreich 2005), are useful in comparative studies. Such models allow us to simulate different interventions and estimate their impact on performance measures, such as wait times. However, as highlighted in (Carter and Blake 2004) developing such simulations is often very hard. Therefore, their usage in practice is not very common. These authors highlighted the difficulties of gathering the right data to calibrate the simulation (often relevant data, e.g., on resources, are missing), the lack of scalability (a simulation model tailored to one hospital is not easily transferable to another one), and the lengthy development times (often in the order of months and years). Yet, they argued that once a proper simulation model becomes useful, it allows prospective comparative analytics and can thus serve decision makers for a relatively long period, including in support of decisions that were not initially intended to be addressed by the simulation.

We note that much of the theoretical development in performing comparative analytics using simulation, as well as casual models and AI, is done outside the ORMS community. As an example, (Zheng et al. 2018) focused on counterfactual and causal effects in view of different patient characteristics. They tested the comparative performance of their estimator with a simulation carefully calibrated to create “real” datasets. They extended the analysis of (van der Laan and Petersen 2007). This earlier work provided a statically optimal patient-specific intervention rule in the presence of missing data. (Zheng et al. 2018) proposed and explained how to implement a targeted maximum likelihood estimator of the impact of intervention on specific individuals in support of prescribing tailored medical interventions. Using the concept of targeted maximum likelihood estimator in support of comparative and prescriptive analytics may benefit ORMS professionals in their business analytics work.

An additional tool for comparative analytics is learning from best practices and best performers. Such learning can be useful for similar business units or employees. (Song et al. 2017) retrospectively investigated public versus private relative performance feedback (RPF) at the employee level in two EDs. The feedback focused on the length of stay of patients and, in both EDs, best practices were shared by management. One ED started to publicly post the RPF on is physicians; the other did not. (Song et al. 2017) compared results for both EDs. They found improved productivity in both hospitals that could likely be attributed to the RPF process and the general feedback provided. Publicly sharing the RPFs further improved productivity and reduced variance (i.e., low productivity workers improved more than medium productivity). Importantly, these improvements were not at the cost of other performance metrics (such as worker attrition).

We note that the essence of the RPF process is similar to descriptive analytics: mapping practices in order to study what practices work well, how to use them where appropriate, and how to further improve existing processes in an organization.

Another common intervention is the use of new technology. In a retrospective analysis, (Lu, Rui, and Seidmann, 2017) focused on the impact of introducing new information technology (IT) in long-term care facilities (LTF). Much of this technology is focused on reducing the workload in tracking and administering medications, i.e., reducing service times. The authors retrospectively studied the impact of implementing IT on nurses' staffing. They developed a stylized model that related the level of IT to quality and revenue in the LTF. They tested their model using regression models fitted to over 7 years of data on nurse staffing at LTF. They found that after employing new IT, staffing dropped in high-quality LTFs but went up in low-quality LTFs. These impacts were related to increased quality and changing patient mix and were driven by different substitutability and complementarity of IT with labor.

4.3 Resource Matching

In several healthcare applications, resource matching can be thought of as double-sided queue. Such applications include blood transfusion and organ transplants. They are different from standard resource matching, such as the matching of elective surgeries to operating rooms. Whereas in the first matching problems, both the supply (donations) and the demand (patients) sides have uncontrolled arrivals, in standard problems, one of these sides, typically the supply, e.g., operating rooms, is controllable. In what follows, we discuss recent analytics work on blood transfusion. We found most analytics models for organ transplant include more prescriptive content and thus discuss these in the next section.

(Sarhangian et al. 2016 and Sarhangian et al. 2017) analyzed practice and theory for the use of red blood cell transfusion. Such transfusions are required for many medical treatments, including surgeries. They considered the impact of different ordering and allocation policies on the outdate rate, shortage rate, and the age distribution of transfused blood. (Sarhangian et al. 2016) developed a simulation based on data for over 10,000 transfused units from a Canadian hospital. Their initial retrospective analysis characterized the polices used and found that if blood were allocated in a first in, first out manner, and an appropriate order-up-to level was chosen, the average age of transfused blood could be reduced by 30% without impacting the other measures. In a prospective study, (Sarhangian et al. 2016) focused on threshold-based allocation policies. They analyzed a stylized queueing theoretical model, derived performance measures, and characterized the trade-off between age and availability for such policies. They suggested the potential effect of distribution of the age of transfused blood on healthcare outcomes can be studied and argued that threshold-based policies are effective when reducing the shelf-life significantly impacts availability.

4.4 Service Path and Diagnostics

A common concept in healthcare services for chronical diseases is the clinical pathway. While there is often a recommended clinical pathway, patients may not always follow it (e.g., treating physicians recommend an alternative path, or patients skip some steps). In this context, data can help in comparing the clinical results of patients who have followed different pathways. As the pathway also depends on diagnostics, the choice of when to use different tests is relevant.

In the context of EDs, this choice may be related to the relevant congestion in the system. We therefore discuss recent analytics work on service path and diagnostics.

A data-driven robust optimization approach was used to measure the clinical pathway concordance for colon cancer patients in (Chan et al. 2019). They used 1 year of patient level data on 763 patients' pathways and survival (over 4 years) to calibrate their model and 5 years as test data (with over 4,000 patients). They retrospectively compared survival results for patients with different concordance levels and showed that higher concordance levels with the recommended pathway were correlated with better survival chances. They considered several other methods to measure the concordance level and concluded that their metric performed better than the other methods. The study suggests the benefit of the recommended pathway. Future interventions could focus on encouraging higher concordance with the recommended pathway and on adequate responses for patients whose treatment does not follow this pathway.

The service path often depends on results from diagnostic tests. While business analytics work should take the results of these tests as given, there may be a place to use analytics in the choice of testing. Focusing on the service path in EDs, (Shi et al 2020) studied the adoption of diagnostic tests by jointly considering their medical and operational benefits (measured as average length of stay). They demonstrated that some diagnostic tests are effective for a single patient from a medical perspective but may add congestion that harms the overall health delivery process and thus should not be adopted in practice (by the same token, some tests that may be less effective from a medical point of view for a single patient may be effective overall, as they reduce congestion). They studied this problem using a Markov decision process when considering the routing in the ED. They used data for over 10,000 patient visits; about 35% of these were suspected to have some level of pulmonary embolism (i.e., blood clot that blocks an artery in the lung). Diagnosing the occurrence of pulmonary embolism requires a CT scan, a congested station in the ED, and such a test may be prevented if a D-dime test suggests no pulmonary embolism is present. They prospectively compared several different paths in accordance with different levels of usage of the D-dime test and demonstrated that with proper routing, the adoption of this test may be beneficial to the patients visiting the ED (e.g., by reducing the wait time for CT for stroke patients).

5. Prescriptive Analytics

One of the potentially most impactful premises of data availability is the ability to tailor services. Tailored offerings could be based on more detailed customer segmentation, something that is feasible with big data. In the context of healthcare, finer-grained segmentation supports personalized medicine and improved service path. Operations research tools could also be used to improve diagnostics and patients' response, e.g., when the response depends on a delay in the diagnostic time.

We next discuss research on the following: scheduling policies, policies for resource matching for organ transplants, preferred service paths for patients, and improvements in the operations of diagnostic services. We note, however, that most of this work is traditional ORMS work and does not carefully depict the progression from descriptive, to predictive and comparative analytics on the way to prescriptive analytics.

5.1 Scheduling Policies

(Jiang, Abouee Mehrizi, and Van Mieghem, 2020) focused on scheduling patients waiting for MRI. They considered an MRI scheduling problem for 72 hospitals in Ontario, Canada. They used 30 million patient-level records for the period of January 2013 to December 2017. In consideration of the inherent uncertainty, and as is common in healthcare, the service level for MRI tests was defined as the tail probability of the wait time (e.g., policy makers aim at keeping the percentage of patients who wait more than a day at or below 10%). They used ML to estimate demand and service, as well as to cluster hospitals within geographic regions while considering traveling times. They showed that the impact of (the right) pooling was larger on the variance of waits (and thus tail probabilities) than on their mean and could be quite significant.

5.2 Resource Matching

The process for allocating livers to transplant patients was investigated in (Akan et al 2012). They used a fluid model for liver transplant policies while considering different objectives and performance measures, including total quality adjusted life years, deaths, and number of wasted livers. Using shadow prices, they proposed different polices based on, e.g., an objective that was a convex combination of two different measures. They used 14 years of data from the united network for organ sharing to validate, cross validate, and then test the efficacy of their suggested policies using simulation. They found their recommended policies may be better than the one used in practice for measures such as total quality adjusted life years and number of wasted livers, but may increase the total number of patient deaths while waiting for transplant.

The process for allocating a kidney to patients was considered by (Ata, Ding, and Zenios 2020). They suggested several possible allocation policies based upon a combination of fluid analysis and the achievable region approach. They used 5 years of data to estimate kidneys' quality scores; for patients, they used data on their initial health score, health decay curve, mortality rate for each health level, and quality of life from not matching (they normalized the quality of life after each successful matching to 1). They showed that quality of life measures can be improved by prioritizing healthier patients for high quality kidneys, while the less healthy receive low quality kidneys.

5.3 Service Path

Big data availability means prescriptive analytics on service paths can support personalized medicine via decision support systems. These systems are based on observed results after treatment by experts. There are two important challenges in this context. The first challenge is that errors made by experts may become the recommended course of action. The second is that following a newly constructed recommendation may result in newly constructed patient characteristics that have never been observed before. Then, automated systems that are calibrated with observed data can provide no recommendation.

Treatment for chronic diseases is perhaps the primary example of the importance and feasibility of personalized service paths. (Bertsimas et al. 2017) focused on Type 2 diabetes. The authors used detailed electronic medical data for over 10,000 patients (each with at least one year of data in the system); these data included demographic information and medical and treatment histories. For each patient, they used a K nearest (with weighted distance) neighbour regression to find a line of therapy (drug regimen) that was likely to be best for her (the Lasso and Random forest algorithms did not perform as well). A change in the line of therapy was recommended if the predicted benefit of the likely best therapy was significant. In the absence of a clinical trial, they used out-of-sample data to evaluate their recommendations. In the 31.8% of cases where the recommended therapy differed from the observed one, they estimated the medical condition by the average observed condition for neighbours who received the recommended therapy. The estimated average reduction of the mean post-treatment glycated hemoglobin based on their recommendation was 0.44%.

A main limitation of their work was their focus on controlling a unique medical measure (glycated hemoglobin); control of several measures would require a significantly larger data set. Nevertheless, their approach – performing data-driven patient segmentation—can be used to create an effective decision support system for physicians. Moreover, with enough data, novel treatment combinations can be extracted (based upon prediction of their efficacy). Such extraction would reduce the cost of future clinical trials and the time to market of the recommended treatments.

The quality of the service path is often hard to quantify, and expert opinions or surveys are often required to measure it. (Bertsimas, Czerwinski, Kane, 2013) used a combination of expert opinion and analytics to measure the quality of service provided. They used data on insurance claims for medical services and medications. The expert ranked the quality of care for 100 patients and a logistic regression model trained with these data was used to identify patients' whose quality of care should be reviewed. The out-of-sample validation of over 30 patients suggested the use of the scalable logistic regression model may be helpful in practice. Their approach provides a scalable service quality measure based on a small sample from expert opinion; this allows cases that can benefit from an expert review of the service delivery process to be flagged.

Learning from experts can also be applied to radiation therapy for cancer. Experts in this field try to balance several objectives, such as delivering a high dose to the tumor but a low dose to healthy tissue. As there is no clear way to quantify how to proceed with a specific radiation treatment, these treatments are considered knowledge-based, and an expert opinion is required. Providing automated recommendations for such treatments is thus very challenging. Inverse optimization was used by (Chan et al. 2014) to determine which mathematical formulated objectives were best aligned with expert decisions on 12 prostate cancer patients. An important contribution of their work was the identification of a relatively small number of objectives to consider in the treatment plan. Combining these objectives with patient clustering may help create a decision support tool to support treatment planners and oncologists in quickly coming up with effective treatment plans. (Boutilier et al. 2016) studied different approaches to treatment plans based on prostate cancer patients' data. Specifically, they compared (i) prediction using the best plan from a database with existing clinical plans, (ii) principal component analysis and multiple linear regression, and (iii) objective function weights using logistic regression or K-nearest neighbors. They investigated the minimum required sample size of the training set to achieve errors similar to these from a sample with 200 plans. They showed that different methods required different sample sizes. Therefore, for more complicated therapies where data are more complex, there may be a trade-off between the learning method and the quality of the generated recommendation.

Another consideration of service path was by (Ding, Gupta and Tang 2020). They considered appointment scheduling when some customers require a follow up. They investigated the impact of pre-booking appointments to a priority queue for some follow-up appointments. They focused on throughput maximization while considering that some appointments were unused, due to customers either not showing up or no longer needing them. They also considered balking due to long delay. They used over a year of data on about 700,000 appointments for a network of outpatient clinics in Minnesota where follow-up appointments were performed about 20% of the time. They observed that pre-booked, follow-up appointments were booked with the original doctor more often than non-pre-booked ones, and the latter were unused more often. They approximated the queuing system by assuming follow-up appointments see the time average of the system. They showed that while in the absence of balking, there is no benefit in pre-booking follow-ups, there is a unique level of pre-booking that maximizes the throughput (when follow-up appointments are required less than 25% of the time). They provided guidelines on how to find the optimal rate of pre-booked prioritized appointments.

5.4 Diagnostics

Several studies use data to support the careful study and modeling of operations problems in services. We focus below on studies related to diagnostics that specify if and when prescribed interventions are needed; such questions are relevant in many services.

(Deo and Milind 2015) considered how to use rapid but costly diagnostic devices for improving early infant diagnostics for HIV in Mozambique. They investigated where to locate such devices to reduce diagnostic delay. Long delays that are common under the central diagnostic system reduce the percentage of caretakers who collect the results of the diagnostic test and thus the effectiveness of treatment. An important delay in this setting is the batching delay at the lab. The lab only completes diagnostics of full batches of tests, so a test arriving after a batch is processed could wait a long time for the collection of another batch before its diagnostics is completed. The authors combined a detailed simulation study with an approximated and carefully calibrated parents' behavioral model as a part of the device location optimization problem. They showed the optimal usage of such devices can improve the diagnostics process much more than other changes to the existing system. The impact of delays and network design on this problem was also investigated in (Jonasson, Deo, and Gallien 2017). They developed a data-driven simulation that estimated the health benefit of different network designs (lab locations, assignments of clinics to labs, and lab capacities) by considering the detailed operational procedures. In particular they considered the time to diagnostic (from sample collection, transportation, in lab queueing), how the distribution of time to diagnostics impact treatment initiation (e.g., the probability of caretakers getting the lab results), and ultimately the health outcome (based on the size and age of population treated). In addition, they developed an optimization model to suggest beneficial network configurations. This tractable model approximated the impact on the distribution of delays from different designs. They then used this simulation model to estimate the results of recommended network designs. They established that a resource pooling where a single national lab is beneficial in Mozambique, but in countries with longer commute times, designs with more labs may be optimal.

6. Example: From Descriptive to Prescriptive Analytics with Queue Mining

As seen in the sections above, most analytics project are hard to define as most consider only one analytics – either prescriptive, predictive, comparative, or descriptive analytics. While different projects may focus on one, most successful projects progress through several of these phases. To demonstrate the progression of research on business analytics from descriptive, via predictive and comparative, to prescriptive analytics, we now discuss recent advances in queue mining. We believe this methodology may be helpful in many service systems, and it also nicely demonstrates some of the future directions of research on business analytics in services.

Many operational and logistics processes can be represented using a control-flow perspective. This perspective emphasizes the order of activities, decisions made along the way, and concurrent execution blocks. Control flow is typically modelled via dynamic models, such as Business Process Modelling and Notation (BPMN) or Petri nets (aka Place/Transition nets). Service processes with scarce resources and uncertainty are typically modeled via queueing theory. Queueing theory enables resource-centered analysis that is often relevant to the study and practice of service processes. We believe a synergy between approaches can be achieved, making their joint usage a very appropriate modelling framework for services. In a concise fashion, we define:

- **Queue Mining**: business analytics for capacitated systems that combines transaction level data analysis and queueing theory.

In the last two decades, the availability of detailed transaction-level data has increased. Process mining is a rapidly growing research field that aims at eliciting control-flow models from such transactional data (van der Aalst, 2011). On the queueing side, these data have often been used to calibrate theoretical queueing models (analytical, numerical, or simulation based). However, the models were first manually constructed and only later annotated with data (e.g., fitting arrival processes, estimating durations, and mining service policies).

A recent development that uses transactional data to automatically extract queueing models is queue mining (Senderovich et al. 2014, 2015) and (Senderovich et al. 2016). Queue mining is an extension of process mining to consider queueing aspects that are prominent in services. The premise of queue mining is the use of data to retrieve the realized service process.

Retrieving the realized service process based on available records still requires detailed data analysis. While much more data are available today in many service systems, they are often incomplete, and there is a main differentiation between modern manufacturing and service firms. Incomplete and inaccurate data are especially prevalent when data entry is done manually or by different organizational units. For example, the details related to the process of seeing specialist doctors (e.g., pediatricians, surgeons, etc.) in EDs are often not well documented, as these doctors are not a formal part of the ED team. There is a similar concern about the processes related to patients arriving by ambulance or to medical tests occurring outside the primary medical unit with which the patient is affiliated. As a result, many specific time stamps essential to formulate an appropriate queueing model are often missing.

Missing or incomplete data are common in most business analytics projects. Modern approaches allow missing data to be addressed by trying to impute them or by studying their impact indirectly. Within queue mining, missing data can be imputed by repairing event logs and making educated guesses for missing time stamps. (Rogge-Solti et al. 2013) applied an approach for imputing missing data to the ambulatory surgery process in a Dutch hospital. The operating theatre staff enters information manually, and there are missing data for almost 43% of patients. The authors modeled the service process as a stochastic Petri net and iteratively imputed missing data, combining available data with knowledge on the process. To reduce the computational complexity, they repaired the structure (e.g., a missing step) based on the alignments approach and the time stamps using inference in a Bayesian network. Both repairs can be done deterministically or using simulation. They further discussed how to measure the quality of the repair.

An alternative approach for making educated guesses about missing time stamps was presented in (Senderovich et al. 2016). They used event data to measure queues to identify bottlenecks and process inefficiencies. In many service systems, time stamps for finishing a service are more accurate than those entering the queue or the service (start times are often missing from data, as providers are only obliged to report ending of activities). Therefore, the authors focused on analysing these situations. Their main idea was to cluster customers, using K-means clustering, based on the inter-completion times of their services. They assumed the “shortest” cluster (with the smallest representative) contained customers who did not wait; thus, their time in the system provided an estimate of the service time. Other clusters contained customers who observed, say, a low, medium, or long queue, and the queues they observed provided estimates to queues observed by customers with similar flow times. The authors devised algorithms to impute missing data based on existing data and evaluated them based on empirical data from call centers and a hospital where several services could be performed in parallel.

There is much more to do in this line of research on descriptive analytics, such as considering incomplete data and learning operational procedures, such as prioritization, from data. The most recent advances in queue mining extend it to predictive, comparative, and prescriptive analytics. We describe these in the corresponding sections.

Queue mining supports the formulation of relevant models, helps compare the realized versus planned processes, and supports predictions and analysis of performance metrics. A comparison of the planned process and reported performance measures allows a detailed verification and validation of the resulting process model (Senderovich et al. 2016). With these steps, a well-designed queue mining can also build intuition on potential process improvement, i.e., support comparative and prescriptive analytics.

A first step in the direction of using process mining to extract simulation models was made by (Rozinat et al. 2008). Their work aimed to support prescriptive analytics for short-term interventions. (Mesabbah, Abo-Hamad, and McKeever, 2019) discussed the creation of an auto simulation model builder in the context of healthcare. They used ML to predict processing times and routing. They assumed data on waiting versus service and resources exist, which is often not the case, and thus incorporated specific resources in their simulation model. They did not report the use of real hospital data, however. In contrast, (Camargoa, Dumasa, and González-Rojasb, 2020) reported such usage with a data set containing event data for both customers and medical personnel. They also provided their source code.

The continuation of queue mining work and transforming it from descriptive to prescriptive would require further theoretical support, such as the analysis of appropriate simulation models to include practical considerations for the scheduling and control of queues. Such considerations are known to improve performance of queueing processes and include priority, resource pooling, and scheduling. The latter considerations are especially important in services that require visits to several stations. We believe further use of queue mining to automatically create simulation models using existing data would facilitate prescriptive analytics in services.

7. On Operations Research and Management Science in Business Analytics

In the preceding pages, we have presented a framework for business analytics and demonstrated its efficacy using recent research on healthcare analytics. We now make several observations on the current and potential impact of ORMS professionals on business analytics based on our survey of the literature.

7.1 The Current State of ORMS and Analytics

On the one hand, there is an increasing amount of work on business analytics in general and on healthcare in particular within the ORMS community. Several very interesting data sets have been used to demonstrate the potential benefits of analytics in healthcare, e.g., http://seeserver.iem.technion.ac.il/see-terminal/. Many papers have used econometric tools that were less commonly used by operations researchers in the past, e.g., (Song, Tucker, and Murrell, 2015). In parallel, more traditional tools, such as simulation, queueing, and optimization techniques, have been used to support service analytics. On the other hand, there is relatively little implementation of this healthcare analytics work. Much of the surveyed work focuses on descriptive or predictive analytics, with no long-term follow up of comparative analytics in support of prescriptive analytics in application. The comparative analytics examples suggest that recommending interventions based upon data is challenging, despite the increased availability. We need more methods to study the impact of interventions. Even the relatively limited work that aims at prescriptive analytics is often published before a long-term impact can be adequately measured and reported; the recommendations are not implemented in practice, possibly because the research does not rely on descriptive, predictive, and comparative phases.

We conclude that implementing business analytics in healthcare services is challenging. This conclusion is supported by others. For example, Scheinker and Brandeau (2020) shed light on the challenges involved in successfully implementing analytics projects in hospitals. Despite the expertise of ORMS professionals in prescriptive analytics in challenging settings, there is less impact than desired of prescriptive analytics on practice. We see implementation of prescriptive analytics projects in general and in healthcare in particular as an exciting and challenging opportunity to extend their impact.

7.2 On the Future of ORMS in Business Analytics

We believe ORMS tools will improve business analytics theory and practice in services such as healthcare. The thinking ORMS professionals bring to the business analytics table includes a focus on decision making and the use of several methodologies. These methodologies include modeling, optimization, stochastic analysis, and—arguably most importantly—the integration of a broad set of methodologies, including ones that traditionally belong to other disciplines (e.g., computer science, econometrics, and statistics). The inclusion of ORMS promises two main benefits. First, the focus of ORMS models on decision making is often very specific, i.e., problem driven. Second, the broad training of ORMS professionals allows us to understand business processes and relate them to performance measures in view of observed, but often incomplete data. Fundamental empirical models and their analysis start with data; we make inferences and conclusions based on these data. Such an analysis is not guided by a model of the system or by the process that generated the data. In contrast, ORMS also considers a model-based approach that is focused on each application. Note that while data often represent a single realization, useful models capture domain knowledge that applies to different realizations. Thus, without models, the inductive step from descriptive to prescriptive analytics requires a blind belief that observations from a specific realization are meaningful to other realizations (possibly from dissimilar settings). Adequate models support comparative analytics of different realizations by changing input parameters for these models. Such models assist in managing and resolving uncertainties. They also give decision makers confidence that they have adequately considered all factors and inherent uncertainties and are thus making robust and effective decisions.

One of the fundamental ORMS approaches, see e.g., (Powell, and Batt 2008), is to formulate a problem-specific model, simplify it, and capture the essence of a complicated problem using appropriate tools (e.g., mathematical analysis that leads to asymptotic analysis, heuristics, and bounds). This approach aims at generating insights that allow us to verify and validate results and assumptions made in the data analysis and also to support managers in critically thinking about the problems they face.

Several of the significant steps in analytics that supports managers in critically thinking about a problem are related to the interpretability of results. Without interpretability, the inductive steps—from descriptive to predictive, comparative, and prescriptive analytics —are large and harder to take (or to convince others to take). Seeking interpretable and insightful results is second nature to ORMS professionals. The focus of ORMS on generating insights and interpretability is therefore a requisite aspect of business analytics. Without these insights, business managers cannot make responsible decisions.

Because ML (such as deep learning) models may have low interpretability, their use in business analytics may not be advised. With time, people and managers could learn to trust machines and computers to support more decisions. For example, the car industry, over the last 80 years, has moved from manual gears to automatic gears, to having cruise control as a standard feature, to automatic breaking, and is now moving towards autonomous vehicles. Similarly, managers have moved from calculating numbers using paper, to slide rules (slipstick), to calculators, and to advanced computer software. Nevertheless, managers who make decisions on "autopilot" solely "because the algorithm recommended it" are not desirable, bankable managers. Therefore, we believe research in the ORMS community that focuses on business analytics could improve the interpretability of ML models and facilitate decisions in conjunction with their development. While interpretability is also studied by computer scientists, the model-based ORMS approach that relates this to specific decisions in application could support new frontiers. The genuine interest the ORMS community already has in business analytics may motivate us to push the theory of business analytics into the context of "ease of interpretability for the regular, savvy business manager" within ML models. Some preliminary work on improved interpretability has been done in (Bertsimas, King, 2017).

Two ORMS methodologies that would be useful to successful business analytics are optimization and stochastic analysis. For example, queue mining takes advantage of advances in information technology (process mining), and this has synergy with operations research in the study of service analytics where queueing models are appropriate. As another example, proper simulation models often require understanding of the stochasticity involved in the simulated process. Simply optimizing the performance of a predictive model is often inappropriate. Such models focus on finding correlations and may ignore causality and the impact of decision making, i.e., of interventions, on data and thus on the generated predictions. Moreover, these models are typically not designed to serve as input to optimization as a part of prescriptive analytics. The recent work of (Elmachtoub and Grigas, 2020) suggests how to jointly perform prediction and optimization in an effective fashion. This approach supports the use of ML in business analytics. Nevertheless, in many cases, performing comparative analytics is essential to support prescriptive analytics.

Another central paradigm of ORMS is model selection. It would be interesting to find effective ways to preform model selection in view of data, possibly by weighing different models. Recent advances in such model selection has been achieved outside the ORMS community. For example, the super learner approach described in (van der Laan et al. 2007) used cross validation to find an optimal weighted combination of several prediction algorithms (based on machine learning) to recommend individual treatment rules. With an appropriate definition of "optimal", a similar approach can and should be used for model selection in the ORMS community.

Finally, we believe the view of ORMS professionals supports proper implementation of descriptive, predictive, and comparative analytics. This approach would help managers with the inductive step required to implement decisions recommended by prescriptive analytics. We, therefore, urge our community to perform analytics in all its phases in support of the practice of data-driven decision making.

7.3 Publishing Research on Business Analytics

While the roots of ORMS research are data driven, publishing data-driven academic work in ORMS has been traditionally very tough. There are many objective reasons: lack of high quality data, confidentiality of problem-specific data that companies are not willing to share, the difficulty of working with public operational data (lack of availability and confidentiality), the inability to verify reproducibility of results, and more. Nevertheless, in recent years, with the advances in business analytics and the higher awareness of the ORMS community's contributions to this field, more companies are willing to share their data for research. In addition, there are several accessible data banks, e.g., https://www.kaggle.com/, and several sub communities of ORMS have data-driven research competitions. As a community, we should publish more such data sets and work with them. At the same time, journals are providing clearer guidelines on publishing research with data; this type of effort supports the progress of the ORMS academic community in business analytics and can support managers in critically thinking about a problem.

Some excellent ORMS analytics papers combine mathematical and modeling tools with empirical ones (and other excellent papers focus on one of these tools). For example, empirical analysis would support and validate assumptions in modeling (e.g., a goodness of fit for a Poisson arrival process). This sort of analytics paper would demonstrate the applicability of results and support the inductive step from descriptive analytics, i.e., past data and knowledge, to prescriptive analytics, i.e., future decisions. Writing such papers suggests the need to broaden the training of ORMS in academia and beyond and makes joint authorship more attractive. Some ORMS experts possess deep knowledge of econometrics; others know more about algorithms, and still others are experts in optimization and so on. We should acknowledge that it is impossible for one person to be an expert in everything.

Having different expertise reflects on the papers we write and the review process. The latter becomes harder: how can I review papers that use tools I am not an expert in? With this in mind, adding structure to the review process is helpful. It is important for us as a community to maintain an open mind to different research approaches. Rather than rejecting papers because they depict a different approach than our favorite one (e.g., using a distinctive set of tools), we need to judge papers based upon their contributions. In many cases, papers that follow different approaches may improve each other. Thus, rather than discounting or disqualifying contributions of a paper because we have a different approach, we should think about extending these contributions using our approach. Such a process would be much more constructive and generate synergy across approaches. Given the famous (Box 1976) quotation "all models are wrong, but some of them are useful," the judgement on accepting a paper should not be that it is "the correct" model, but that it is a useful one. Or as Chris Tang said, in personal communication (before reading this paper), "there are no perfect papers." After all, one of the foundations of the ORMS approach is inclusivity of different tools to support improved decisions. So let us remain inclusive in our refereeing efforts as well.

Acknowledgments

I thank the editor and the anonymous AE and referees for their comments that helped improving the paper. I also thank John Birge, Morris Cohen, Steve Graves, Ming Hu, Dmitry Krass, Georgia Perakis, David SImchi-Levi, and Chris Tang for discussions on business analytics from an ORMS perspective. Special thanks to Lennart Baardman, Robert Batt, So Yeon Chun, Victor Martínez de Albéniz, Arik Senderovich, John Silberholz, and Galit Yom-Tov for comments on previous versions and discussions that helped me formulate the framework for business analytics, discussed in Section 1. This work was supported by an NSERC grant of the author.

Reference

1. Ahuja V, Birge JR, Ryan C. "Approximation Methods for Determining Optimal Allocations in Response Adaptive Clinical Trials," ICORES 2014 - 16 Proceedings of the 3rd International Conference on Operations Research and Enterprise Systems, *SciTePress*, 2014; 460-65.

2. Akan M, Alagoz O, Ata B, Erenay FS, Said A. A broader view of designing the liver allocation system. Operations research. 2012 Aug;60(4):757-70.
3. Baron O, Berman O, Krass D, Wang J. Strategic idling and dynamic scheduling in an open-shop service network: case study and analysis. *Manufacturing & Service Operations Management*, 2017; 19(1), 52–71. 10.1287/msom.2016.0591
4. Baron O, Berman O, Krass D, Wang J. Using strategic idleness to improve customer service experience in service networks. *Operations Research*, 2014; 62(1), January-February, 123-140. DOI: http://dx.doi.org/10.1287/opre.2013.1236
5. Bavafa H, Jónas Oddur Jónasson Recovering from Critical Incidents: Evidence from Paramedic Performance *Manufacturing & Service Operations Management* 2020, forthcoming.
6. Bayati, Mohsen and Kwasnick, Sara and Luo, Danqi and Plambeck, Erica L., Low-Acuity Patients Delay High-Acuity Patients in an ED (December 31, 2017). SSRN: https://ssrn.com/abstract=3095039
7. Bertsimas D, Bjarnadottir M, Kane M, Kryder C, Pandey R, Vempala S, Wang G. Algorithmic Prediction of Health-Care Costs. *Operations Research*, 2008; Vol. 56, No. 6, 1382-1392.
8. Bertsimas D, Cacchiani V, Craft D, Nohadani O. A hybrid approach to beam angle optimization in intensity-modulated radiation therapy. Submitted to *Computers and Operations Research*, 2012.
9. Bertsimas D, Czerwinski D, Kane M. Measuring Quality in Diabetes Care: An Expert-based Statistical Approach. *Publication*, 2013.
10. Bertsimas D, Kallus N, Weinstein A, Zhuo Y. Personalized Diabetes Management Using Electronic Medical Records. *Diabetes Care*, 2017; 40, 210–217.
11. Bertsimas D, King A. Logistic Regression: From Art to Science. *Statistical Science*, 2017; Vol. 32, No. 3, 367–384.
12. Bertsimas D, O'Hair A, Relyea S, Silberholz J. An Analytics Approach to Designing Combination Chemotherapy Regimens for Cancer. *Management Science*, 2016.
13. Brillman JC, Doezema D, Tandberg D, et al. Triage: limitations in predicting need for emergent care and hospital admission. *Ann Emerg Med*. 1996;27(4):493-500. doi:10.1016/s0196-0644(96)70240-8
14. Boutilier JJ, Craig T, Sharpe MB, Chan TC. Sample size requirements for knowledge-based treatment planning. Medical physics. 2016; Mar 1;43(3):1212-21.
15. Box, G. E. P. (1976), Science and statistics, *Journal of the American Statistical Association*, 71 (356): 791–799, doi:10.1080/01621459.1976.10480949
16. Camargoa Marlon Dumasa Oscar González-Rojasb, Automated discovery of business process simulation models from event logs Manuel Decision Support Systems, Available online 26 March 2020, 113284 https://doi.org/10.1016/j.dss.2020.113284.
17. Carter, M.W. and Blake, J.T. “Using Simulation in an Acute Care Hospital: Easier Said Than Done” in: *Handbook of OR/MS Applications in Health Care’* Kluwer Academic Publishers International Series in Operations Research and Management Science (eds., Sainfort, F., Brandeau, M. and Pierskalla W.), 2004.

18. Chan, C. W., V. F. Farias, G. J. Escobar. The impact of delays on service times in the intensive care unit. *Management Science*. 2016; 63(7) https://doi.org/10.1287/mnsc.2016.2441.
19. Chan T. C.Y., Tim Craig, Taewoo Lee, Michael B. Sharpe, Generalized Inverse Multiobjective Optimization with Application to Cancer Therapy. *Operations Research* 2014; 62(3):680-695. https://doi.org/10.1287/opre.2014.1267.
20. Chan T. C.Y., Demirtas D., Kwon R. H., Optimizing the Deployment of Public Access Defibrillators. *Management Science*, 2016; Pages iv–vii, 3393-3672.
21. Chan, T. C.Y., Eberg, M., Forster, K., Holloway, C., Ieraci, L., Shalaby, Y., & Yousefi, N. An Inverse Optimization Approach to Measuring Clinical Pathway Concordance. (2019) arXiv preprint arXiv:1906.02636.
22. Chan, T. C.Y., Nicholas, H., Lagzi, S., Quiroga, B.F., Gonzalo, R. (2021) Negative Spillover on Service Level across Priority Classes: Evidence from a Radiology Workflow Platform (February 8). Rotman School of Management Working Paper, Forthcoming, Available at SSRN: https://ssrn.com/abstract=3781929 or http://dx.doi.org/10.2139/ssrn.3781929
23. Chan T. C.Y., Li H., Lebovic G., Tang S. K., Chan J. Y.T., Cheng H. C.K., Morrison L. J., Brooks S. C., Identifying Locations for Public Access Defibrillators Using Mathematical Optimization. *Circulation* 2013; Vol. 127, No. 17, 1801—1809.
24. Chen, N., Gurlek, R., Lee, D., and Shen, H., (2021) Can customer arrival rates be modelled by sine waves? Improving service operations with queuing data analytics. SSRN: https://ssrn.com/abstract=3125120 or http://dx.doi.org/10.2139/ssrn.3125120
25. Chen W., Linthicum B., Argon N.T, Bohrmann T., Lopiano K., Mehrotra A., Travers D., Ziya S., The effects of ED crowding on triage and hospital admission decisions, *The American Journal of Emergency Medicine*, Volume 38, Issue 4, 2020, Pages 774-779, ISSN 0735-6757, https://doi.org/10.1016/j.ajem.2019.06.039.
26. Claypool, A., Brandeau, M. L., Goldhaber-Fiebert, J. D. Prevention and control of Denge and Chikungunya in Columbia: A cost-effectiveness analysis SAGE PUBLICATIONS INC. 2020: E85–E87.
27. Cohen, M., Big Data and Service Operations, *Production and Operations Management* 27(9):1709-1723, 2018.
28. Deo, S., Aditya J. Slow first, fast later: Empirical evidence of speed-up in service episodes of finite duration. *Production and Operations Management*, 2019 28 (5), 1061-1081.
29. Deo, S., Sohoni, M.G. "Optimal decentralization of early infant diagnosis of HIV in resource-limited settings.", *Manufacturing and Service Operations Management.*, 2015 17 (2), 191-207.
30. Diamant, A., Milner, J., Cleghorn, M., Sockalingam, S., Okrainec, A., Jackson, T. D., Quereshy, F. Analysis of Patient Attrition in a Publicly Funded Bariatric Surgery Program. *Journal of the American College of Surgeons*, 2014 219(5), 1047-1055.
31. Ding, Y., Gupta, D., Tang, X. (2020). Early Reservation for Follow-up Appointments in a Slotted-Service Queue. Available at SSRN 3616565.
32. Ding Y., Park E., Nagarajan M., & Grafstein E. "Patient Prioritization in ED Triage Systems: An Empirical Study of Canadian Triage and Acuity Scale (CTAS)", *Manufacturing and Service Operations Management*, 2019 21(4), 723-741.

33. Dong J, Yom-Tov E., Yom-Tov G.B. The Impact of Delay Announcements on Hospital Network Coordination and Waiting Times. *Management Science*, 2019 65(5):1969-1994.
34. Elmachtoub N. A, Grigas P. (2021) Smart "Predict, then Optimize," *Management Science*, 10.1287/mnsc.2020.3922
35. Hu M. (Ed.) Sharing Economy: Making Supply Meet Demand, Springer, 2019 1-528.
36. Hutton D.W., So S.K., Margaret L. B., Cost-effectiveness of nationwide hepatitis B catch-up vaccination among children and adolescents in China, January 2010 https://doi.org/10.1002/hep.23310.
37. Ibanez, M. R., Clark, J. R., Huckman R.S., and Staats, B. R. (2018) Discretionary Task Ordering: Queue Management in Radiological Services, *Management Science* 64:9, 4389-4407
38. Jiang, Y., Abouee Mehrizi, H., & Van Mieghem, J. A. (2020). Geographic Resource Pooling to Balance Waiting and Traveling: Data-driven Application to MRI Hospitals in Ontario. Available at SSRN.
39. Jónasson, J. O., Deo, S., Gallien J. "Improving HIV early infant diagnosis supply chains in sub-Saharan Africa: Models and Application to Mozambique.", *Operations Research.*, 2017 65 (6), 1479-1493.
40. Lu S. F., Rui H., Seidmann A., Does Technology Substitute for Nurses? Staffing Decisions in Nursing Homes. *Management Science* (2017) Vol. 64, No. 4.
41. Lustig, I., Dietrich B., Johnson C., and Dziekan C. The analytics journey. *Analytics Magazine*, 2010 pp.11–18, November/December.
42. Malloy, G., Andrews, J., Brandeau, M. L., Goldhaber-Fiebert, J. D. Cost-effectiveness of interventions to prevent plague in Madgascar**,** Sage Publications Inc 2020: E116–E117.
43. Malloy, G., Goldhaber-Fiebert, J. D., Enns, E.A., Brandeau, M. L., Predicting the effectiveness of interventions for infectious disease control: the role of model structure**,** Sage Publications Inc 2020: E116–E117.
44. Marmor Y, Sinreich DA (2005) Emergency department operations: The basis for developing a simulation tool. *IIE Trans.* 37(3): 233–245
45. Mesabbah M, Waleed A.H, and McKeever S.. A hybrid process mining framework for automated simulation modeling for healthcare, *Proceedings of the 2019 Winter Simulation Conference N. Mustafee, K.-H.G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, eds.*
46. Powell, Stephen G., and Robert J. Batt. Modeling for Insight: A Master Class for Business Analysts. John Wiley & Sons, Inc., Hoboken, NJ, 2008.
47. Richardson, D. No Relationship between ED Activity and Triage Categorization, *Academic Emergency Medicine* 5(2), 1998, Wiley, https://dx.doi.org/10.1111/j.1553-2712.1998.tb02599.x.
48. Rogge-Solti A, Mans RS, van der Aalst WMP, Weske M. Repairing Event Logs Using Stochastic Process Models. *(Technische Berichte des Hasso-Plattner-Instituts für Softwaresystemtechnik an der Universität Potsdam; 78)* In: Potsdam:Universität Potsdam, 2013.
49. Rose Robert (2016) Defining analytics: a conceptual framework, *ORMS-Today*; June | Volume 43 | Number 3.

50. Rozinat A., Wynn M.T, van der Aalst WMP, ter Hofstede AHM, Fidge AHM, Workflow simulation for operational decision Support *Data & Knowledge Engineering*, North-Holland, 68(9), 834-850, 2009/9/1.
51. Sarhangian, V., Abouee-Mehrizi, H., Baron, O., Berman, O., Heddle, N. M. and Barty, R. Reducing the age of transfused red blood cells in hospitals: ordering and allocation policies. *Vox Sanguinis* 2016 110: 385–392. doi:10.1111/vox.12374.
52. Sarhangian, V., Abouee Mehrizi, H., Baron, O., Berman, O., Allocation policies in blood transfusion. *Manufacturing & Service Operations Management*, 2017 20(2), pp347-362 https://doi.org/10.1287/msom.2017.0650.
53. Scheinker, D., Brandeau, M.L., Implementing Analytics Projects in a Hospital: Successes, Failures, and Opportunities. *INFORMS Journal on Applied Analytics* 2020 doi:10.1287/inte.2020.1036.
54. SEE-Center Technion (2009) Seestat database http://seeserver.iem.technion.ac.il/see-terminal/.
55. Senderovich A, Leemans SJJ, Harel S, Gal A, Mandelbaum A, van der Aalst WMP. Discovering Queues from Event Logs with Varying Levels of Information. In: Reichert M., Reijers H. (eds) *Business Process Management Workshops*. BPM 2016; Lecture Notes in *Business Information Processing*, vol 256. Springer, Cham.
56. Senderovich A, Weidlich M, Gal A, Mandelbaum A. Queue Mining – Predicting Delays in Service Processes. In: Jarke M. et al. (eds) Advanced Information Systems Engineering. CAiSE 2014. Lecture Notes in Computer Science, vol 8484. *Springer, Cham*, 2014; 42-5.
57. Senderovich, Arik, Matthias Weidlich, Avigdor Gal, and Avishai Mandelbaum. Queue mining for delay prediction in multi-class service processes. Information Systems 53 (2015): 278-295.
58. Shi, P., Helm, J. E., Sebastian Heese, H., & Mitchell, A. M. (2020) An Operational Framework for the Adoption and Integration of New Diagnostic Tests. *Production and Operations Management*, forthcoming.
59. Simchi-Levi D. OM Forum—OM Research: From Problem-Driven to Data-Driven Research. *Manufacturing & Service Operations Management*, 2014; 16 (1), 2-10. https://doi.org/10.1287/msom.2013.0471.
60. Song H., Tucker A.L., Murrell K.L. The Diseconomies of Queue Pooling: An Empirical Investigation of ED Length of Stay, *Management Science*, 61(12), December 2015, Pages iv–vi, 2825-3096.
61. Song, H., Tucker A., Murrell K., Vinson D. 2017 Closing the productivity gap: Improving worker productivity through public relative performance feedback and validation of best practices. *Management Science*, 64(6).
62. van der Aalst W. *Process Mining: Discovery, Conformance and Enhancement of Business Processes.* Springer, 2011.
63. van der Laan MJ, Petersen ML. (2007) Statistical learning of origin-specific statically optimal individualized treatment rules. *Int J Biostat*;3(1):Article 6. doi: 10.2202/1557-4679.1040. PMID: 19122792; PMCID: PMC2613337.
64. van der Laan MJ, Polley EC, Hubbard AE. 2007 Super learner. *Stat Appl Genet Mol Biol*; 6: Article25. doi: 10.2202/1544-6115.1309. Epub Sep 16. PMID: 17910531.

65. Wang J., Baron, O., and Scheller Wolf, A., (2015) M/M/c queue with two priority classes. Operations Research, 63(3), pp733-749. http://dx.doi.org/10.1287/opre.2015.1375.
66. Yang S, Kou SC, Lu F, Brownstein JS, Brooke N, Santillana M. (2017a) Advances in using Internet searches to track dengue. *PLOS Computational Biology* (Internet); 13 (7), 1-14.
67. Yang S, Santillana M, Brownstein JS, Gray J, Richardson S, Kou SC. (2017b) Using electronic health records and Internet search information for accurate influenza forecasting. *BMC Infectious Diseases* (Internet); 17 (1), 332.
68. Yom-Tov, G. B., & Mandelbaum, A. (2014). Erlang-R: A time-varying queue with reentrant customers, in support of healthcare staffing. *Manufacturing & Service Operations Management*, 16(2), 283-299.
69. Zheng W, Luo Z, van der Laan MJ. (2018) Marginal Structural Models with Counterfactual Effect Modifiers. *Int J Biostat.* Jun 8;14(1):/j/ijb.2018.14.issue-1/ijb-2018-0039/ijb-2018-0039.xml. doi: 10.1515/ijb-2018-0039. PMID: 29883322; PMCID: PMC6682415.