INFORMS TutORials in Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

# From OR to OM: Observations on the Emergence of a Discipline

John Buzacott

TUTORIALS IN
OPERATIONS RESEARCH
INFORMS 2014

 | ISBN 978-0-9843378-5-9
http://dx.doi.org/10.1287/educ.2014.0125

# From OR to OM: Observations on the Emergence of a Discipline

***John Buzacott***
Schulich School of Business, York University, Toronto, Ontario M3J IP3, Canada,
jbuzacot@schulich.yorku.ca

**Abstract** This tutorial is based on the author's experience in applying operational research or operations research (OR) in industry 50 years ago during the early 1960s. It describes some of the problems and issues, almost all of which were new at the time, long before any texts on the use of OR to understand production and inventory control. However, once academics, consultants, and software developers became aware of the issues, the body of knowledge acquired through OR became the basis for a new discipline, operations management (OM).

**Keywords** operations management; operational research; operations research; OR; operational research; history of OR; inventories; investment appraisal



## 1. Introduction

During World War II (WWII), the immediate postwar period, and Korean War rearmament, firms could sell whatever they could make. However, in the mid-1950s, markets became more competitive, and the variety of models, options, and features started to increase dramatically. Senior managers of manufacturing firms recognized that improving manufacturing performance and productivity was essential for them to remain cost competitive. So they turned to one or more of the approaches then called *systems and procedures*, *organization and methods*, *work study*, or *industrial engineering* in the hope that practitioners of these approaches would be able to help them in their search for improved productivity. On the shop floor, time and motion study, methods study, work study, and industrial engineering all proved to be very valuable in improving the performance of individual workers performing manufacturing operations. Systems and procedures or organization and methods were used to try to improve the flow of information and the coordination between different parts of the organization. Although computers were not as yet used in manufacturing, accounting machines were increasingly used for routine accounting and control tasks like payroll and the monitoring and analysis of manufacturing costs. Textbooks on production management tended to reflect this focus on data and control, as they devoted a lot of space to the design of forms to report on or control of different manufacturing functions (see Moore [7]).

The experience of consultants and others was that, particularly in small or medium-sized firms, it was sometimes possible to achieve productivity improvement of up to 30% using improved job and task design on the shop floor, or by using appropriate speeds and feeds for machining operations. However, systems and procedures or organization and methods rarely gave such improvement, although they did reduce the time a new manufacturing manager required to master his role, which perhaps led to a more rapid rotation of managers through different manufacturing management positions. Recently hired graduates often spent a year or two in a manufacturing management position before moving on to a position in marketing or sales or general management.

In the mid to late 1950s, a number of firms in the United States and the United Kingdom started to explore the use of operational research or operations research (OR) approaches for productivity improvement. Sometimes, instead of OR, the approach was called cybernetics, and sometimes systems analysis. Given the success of OR during WWII, it might have been expected that the application of OR to the manufacturing industry would have been led by people who had been active in OR during the war, but it seems that this was rarely the case. At the end of the war, most people who had been OR analysts either went back to their original scientific disciplines or stayed in the defence forces, where the Cold War and the Korean War gave ample scope for them to continue to practice military OR. Immediately after WWII, manufacturing firms did not see the relevance of OR.

More typical of the early practitioners of OR in manufacturing in the 1950s were people who had had some exposure to OR during the war as army or air force officers, or, alternatively, fresh graduates in mathematics or the physical sciences who might otherwise have become high school teachers. Formal training in OR did not begin until 1958–1959 in the United Kingdom, and perhaps a year or so earlier in the United States. Most of the early recruits learned OR on the job. This was easier than it would be today because, like the WWII practitioners, the main technique used and required was fairly elementary statistical analysis combined with basic systems thinking skills, that is, the ability to recognize and understand the components of the system being studied. It should be noted that until the early 1960s, undergraduate engineering programs did not include any probability or statistics, so few people in industry had any knowledge of the basic principles of probability and statistics.

This tutorial is based on my experience in practicing OR in industry in the early 1960s. I will describe the contribution OR made to manufacturing and distribution. The field was then new, so we encountered for the first time issues that were later on taken up by academics, consultants, and software developers. Eventually a new discipline, operations management (OM), emerged.

## 2. My Introduction to OR in Manufacturing

My first acquaintance with the practice of OR in manufacturing was in 1959. I was a recent graduate in electrical engineering working in a large electrical equipment manufacturing company. As part of the graduate training program, I spent three months at a plant making small motors, almost all the time sitting in the office of the production controller learning about how he planned, loaded, scheduled, and monitored the plant in such a way that he met customer expectations and kept costs under control. He gave me a report prepared by the OR group at the largest plant in the company. It described the use of the economic order quantity (EOQ) in inventory management. It was all new to me, and it intrigued me. I asked myself whether EOQs were relevant to the small motor plant. All motors were made to order, and there was a large variety of different motor specifications, so I decided EOQs were not appropriate for final assembly. However, EOQs did seem relevant for some of the motor parts that were used in a large number of different end products and that were bought from outside suppliers. Seeing the potential relevance of the EOQ and its basis in a mathematical model, I wondered whether OR might provide relevant insights for other problems encountered by the production controller.

My graduate training program continued, and I worked in design, development, testing, and sales as well as manufacturing. I saw the problems the firm had: long quoted delivery times that were no longer competitive with manufacturers in other countries, and even then we would deliver late; poor coordination of the manufacture of the parts and assemblies for make-to-order products so that some parts would be in inventory for many months waiting for other parts required for the product to complete manufacture; inadequate feedback to customers on the progress of their orders; high inventories of spare parts; and so on. The firm excelled in design and development, but it was mediocre in manufacturing. These observations made me wonder whether OR could contribute to improving manufacturing performance, and so I decided to learn more about it. I came across the Churchman et al. [4] book, and I looked around for master's programs in OR. In the United Kingdom, only the University of Birmingham had a program. It was 12 months in duration, 8 months of course work and 4 months for an industrial-based project. I applied, was accepted, and obtained the MSc in OR in 1962. My project was with the OR group of an oil company working on the problem of how much storage capacity was required at depots that were supplied by sea using ocean tankers. After completion of the MSc program I worked as an OR analyst in the systems group in a domestic appliance company for 15 months. The University of Birmingham offered me a generous fellowship, so I then started a PhD program.

## 3. Practicing OR in Industry

At the oil and appliance companies, almost all the problems we looked at were new, in the sense that there were no texts and very little literature that described the problems and outlined approaches for their solution. So we had to determine what issues concerned management, develop an understanding of the system or situation in which the problem arose, and arrive at approaches for "solving" the problem. The end result was sometimes a report, usually a presentation to a manager, or sometimes just a meeting or phone call where we communicated our insights and suggestions.

When I look back to that time, I can see two different approaches to the practice of OR. *Traditional operational research* was a continuation of the WWII approach developed in the United Kingdom, with the emphasis on observing an actual operational situation, trying to understand the systems involved, collecting some data, analyzing the data using elementary statistical approaches, and drawing some inferences about how the situation could be improved or the problem solved. Any models used were fairly simple, at most involving the use of fairly basic probability theory. In the United Kingdom, this approach was dominant. The other approach built on the use of more sophisticated models, particularly linear programming (LP) and its extensions. I will call this approach *modern operations research*, reflecting the origins of LP. Back then, available computing resources and mathematical understanding meant that programming models with integer variables were not solvable, although reasonable answers could sometimes be obtained using approximate or heuristic methods. Problems for which LP and extensions were relevant were primarily planning problems where the focus was not on improving present operations but on making decisions about investments in long-lived facilities. For various reasons, partly macroeconomic, partly geographic (small countries have fewer choices), and partly institutional, the latter approach was rarely used in the United Kingdom apart from the oil industry and the nationalized centrally planned industries like electricity and coal. However, it did find more application in the United States, where it fitted the needs of large firms with strong central staffs to evaluate investment proposals (see Chandler [3]).

## 4. Problems Approached Using Traditional Operational Research

Traditional operational research offered a major contribution in inventory management and in many aspects of production planning, indeed any aspect of production marked by uncertainty and variability. Systems and procedures or organization and methods could not contribute much to these areas apart from improving information flow.

For inventory management, operational research offered a deeper understanding of the short-term forecasting of demand for a large number of items. Judgment forecasts for the final demand of major product groups were acceptable, but when it came to forecasting the demand for spare parts or for individual final products in, for example, the film and photography industry, no human could produce good forecasts fast enough. So a means of producing short-term forecasts automatically became essential. In the early 1960s we found that only moving average forecasts would do. Exponential smoothing was attractive, but its problem was that it required multiplication by a decimal number, whereas moving average only required addition and perhaps division by a small integer.

Next, we found that we had to understand what determined the requirement for safety stocks. Traditionally safety stocks were set as a certain number of weeks' average demand. We recognized that the key was variability of demand in the lead time. This led to seeking to understand the difference between periodic review and continuous review of inventory and the way in which periodic review amplified effective lead time. However, we saw the problems in continuous monitoring of inventory levels, particularly in the days before computer control of inventory gave ready access to data on withdrawals from inventory.

At the oil company, we recognized that variability of demand in the lead time might also be due to variations in the lead time as well as variability in demand. When inventory was replenished by ocean tankers delivering product from an oil refinery to a coastal tank farm, weather and other delays resulted in considerable variation in lead time, three to five days early or late from the planned lead time of about two weeks. So in setting safety stocks and order points, we recognized that variability in lead time demand had to consider the variation in lead time. We used the classical formula from Feller [5] for the variance of a random sum of random variables. In 1962, this may have been one of the first applications of this formula for safety stock determination.

At the appliance company, order quantities were typically set as a certain number of weeks' demand. In seeking to apply the EOQ formula, we recognized that relatively few order quantities could be used. Much later Roundy [10] considered order quantities related by powers of 2. We chose to use order quantities where the number of weeks of demand covered by an order was either 1, 2, 6, or 26.

Inventory theory at that time tended to assume local control of inventories. However, we observed that it was more likely that we could implement improvement when inventories were controlled centrally. For example, in managing inventories of parts at service depots, the depot reported inventory at fixed periodic intervals, and the decision on replenishment quantity was made centrally at the main service warehouse. Central control meant that the focus of our OR project became the development of generic guidelines for inventory control. We recommended an ABC system, with distinction between classes being the frequency of reporting the inventory level (and thus the frequency of replenishment), plus another class of items, those that would not be stocked at depots, but rather only shipped to a depot when there was a demand. Using guidelines for determining safety stocks and target stocks for each class and the costs of holding inventory, reviewing inventory, and calculating replenishments and of stockouts, we could then determine the criteria for allocating parts to classes and present these criteria in the form of a simple graph.

The firm had had a major problem with a new model of washing machine resulting in a significant increase in demand for certain spare parts. As a result, it had decided that it should set target stocks as a multiple of forecast demand, with the forecast demand based on a moving average of recent past demand per review period. We developed an algebraic model of the calculations involved and asked ourselves, what is the relationship between variability of demand and variability of replenishment quantity? We observed that replenishments were significantly more variable than demand. When we looked at past decisions, this was particularly noticeable around public holidays or at Easter or Christmas, where there would be a week with lower demand than average followed by a week of higher than usual demand. Replenishments then showed a multiplied fluctuation. We wondered whether this would affect the required safety stock. At the time we thought it would and set recommended safety stocks quite conservatively. It was not until many years later that I developed an adequate model of the impact on required safety stocks (Buzacott [2]). The phenomenon we observed was later given the name *the bullwhip effect* (Lee et al. [6]).

One other inventory issue that we talked about but never resolved was the relationship between inventory and accounting measures such as accounts receivable and accounts payable. We had done an information systems project on the calculation of accounts receivable and accounts payable. We wondered how payment terms with suppliers affected optimal inventory levels, but we did not pursue this question.

When I worked at the appliance company, computers in industry were rare, and the company was one of the pioneers in the United Kingdom in their use. One of the earliest applications was sales invoicing. This had been done previously using accounting machines, so it was not particularly difficult to transfer the application to computer, although since all customer information was stored on magnetic tape and each run required updating the whole tape, the economics of using a computer were questionable; in fact, when there was a drop in sales as a consequence of government policy, the company developed a cheaper invoicing system that did not use computers. A computer was also installed at the plant, and apart from some inventory recording and payroll, the first major application was in material requirements planning (MRP). This required development of a bill of materials processor because there were a number of parts common to refrigerators and stoves, as well as between different models of appliances. We did not attempt time phasing, but, given the production plan over the next three months, we were able to determine the total demand for each part, taking into account current inventory in stores and in process. This simplified and helped decision making on when and how much to release to production. We recognized that lead time determination was a problem if the plant were ever to move to full MRP implementation and thought about using queueing theory but did not ever develop an approach to do this.

Demand for refrigerators was then quite seasonal, whereas demand for washing machines was more uniform, with somewhat higher sales in the spring and the fall. However, a bigger influence on demand was government macroeconomic policy, such as changes in sales tax rates or controls on installment purchase credit availability and interest rates. Although some limited overtime was feasible, the main approach used by the firm to adapt production to demand was to change the number of shifts the plant operated, varying from one shift per day to two shifts and sometimes three shifts. It was difficult to decide when to change, and often the firm ended up with excessive inventory because the decision to reduce the number of shifts worked was left too late. One policy change made was to transfer the major responsibility for the decision from production to sales. We were asked informally by the general sales manager for insight on how to make this decision. We suggested that he plot cumulative production and cumulative estimated demand so that the difference between the two curves was equal to the inventory of appliances. We then talked about how the graph could be used to demonstrate different production strategies. With only three possible production rates, it was relatively easy to gain some insight into the impact of changing production rates.

One production problem we did not work on was production scheduling. This was left to individual production supervisors. I had seen from my experience at the small electric motor plant that the challenge was in work release to production units, not in detailed sequencing. At that time it was far too difficult to collect data from the shop floor—the only way data could be collected was to have expeditors walk around the shop floor every day and see what was happening to each job. Any attempt to schedule production centrally would have failed because of the lack of reliable information about job status.

Characteristic of all the above problem areas was that the techniques used were basic probability and statistics combined with fairly simple cost models. One reason the analysis was simple was that data were limited. When looking at inventory problems, it was hard to find any records going back more than six months, so with one observation per week there were never more than about 20 observations for any part or product. So it would have been hard to use a sophisticated model. We did think a lot about what costs or other penalties were relevant to the decision, and we did think about how to present our results in a simple and understandable way.

We learned a great deal about the company's operations. We interacted with people from production, service, sales and marketing, and accounting. We also got to see the problems and concerns of all these functional areas. For example, when working on service inventory control, we came to see the limitations of the auditing process. The company always had a significant inventory loss for spare parts, which the auditor had attributed to pilferage. We chose an appropriately sized random sample of service invoices and realized that at least a third of the inventory loss was due to the servicemen not writing down on the internal invoice the cost of parts used for warranty service. This portion of the inventory loss should have appeared as a warranty cost, although this mistake did not affect the bottom line profitability of the company.

When the traditional operational research approach was augmented by some basic probability concepts, it was able to make a significant contribution to understanding problems in the day-to-day operations of the firm, particularly in manufacturing, distribution, and service. We also made contributions to market research, analyzing information from warranty registration data to try and assess the level of inventory in retail outlets.

## 5. Problems Approached Using Modern Operations Research

The traditional OR approach did not work when looking at decisions about major investments. Typically such investments would have lives of many years, and if a major mistake was made, the investment would have little continuing value. By contrast, operational mistakes like ordering too much could be recovered from in a matter of months and leave no permanent damage. Major investments had to be approved by top management and often by the board of directors. These people usually had many years of experience in the firm and so had a lot of preconceptions about what investment choice should be made. Operations research had to offer something unique to make a relevant contribution. Typically this meant that OR had to offer new techniques for choosing and evaluating investment options. Alternatively, if the investment options involved radical new technology, OR had to be able to offer significant assistance to the process of understanding and using the new technology (as it had in the early days of radar in WWII).

When I was doing my MSc project with the oil company, I became aware of the limitations of traditional operational research. The challenges of oil product distribution were not so much with day-to-day operations but in planning investments in the distribution system to meet the growing demand. There were major investment issues—which depots could be closed; whether product pipelines should be built from refineries to certain depots; what was the role of unit trains, that is, a train where all cars went from a given refinery to a given depot; how much storage was needed at ocean-tanker-supplied installations. There were many possible alternatives to be considered, and we needed tools and approaches that would enable us to evaluate and select alternatives, and we had to be able to evaluate alternatives in ways that considered their impact on the firm's balance sheet in future years. We realized that the recently developed tools of mathematical programming would be the base for identifying and evaluating options.

In arriving at investment recommendations, we had to be able first to evaluate the performance of a given set of investments and second to find the "best" set of investments. The performance of a given set of investments required us to model the operational issues, such as when to reorder, what safety stocks were needed, and how much of a given product would be shipped on a tanker or a unit train. Sometimes this step required us to use math programming; for example, given a set of depots and refineries, what is the optimal distribution plan that determines which refinery should supply which depot?

Once we had determined the cost and other performance indicators associated with an investment option, finding the "best" set of investments would have to be done using mathematical programming. However, we rapidly recognized the limitations of then-available approaches. Decisions about opening and closing depots could be represented by integer variables in our distribution planning models. Similarly, decisions about the routing and loading of ships or unit trains required integer variables. Back in 1962, neither theory nor the available software and computers were such that we could solve the mixed-integer programming problems, so we had to come up with heuristic approaches that enabled us to identify a relatively small number of "good" alternatives and then evaluate them. To determine the impact of the investment on the firm's future balance sheets, the OR group had to consider the tax consequences of various investment incentives and depreciation rules, so it pioneered the use of discounted cash flow techniques for investment appraisal by the company.

One challenging problem we met when looking at storage requirements arose because a tanker could carry more than one product. So, in determining the quantity of a product the tanker would deliver, we had to compare the cost of a tanker voyage where the tanker carried several products to one installation with the cost of a tanker voyage where the tanker carried just one product but visited several installations. Every port had a different approach for charging port costs as a function of the size of tanker and the quantity of product discharged.

There is no doubt that these investment problems were more challenging mathematically than the operational problems, although back in 1962 our solutions were technically limited. However, it was also clear that as the mathematical and computational issues were resolved, OR would be able to make a major contribution in large firms with strong central governance. However, particularly in the United Kingdom, very few firms met this requirement. They were either too small or too loosely governed for operations research to have much scope. Even in industries with large investment requirements like steel or automobiles, reliance on government assistance for investment meant that often political considerations resulted in far from optimal decisions.

## 6. PhD Research

From the appliance company, I went back to Birmingham for a PhD. At that time it was viewed as essential that a PhD dissertation be based on a real problem in industry. Using and developing techniques to solve the problem was also expected. Sometimes the use of the technique was relatively routine, for example, statistical analysis of performance data. But for dissertations with an OR emphasis, the development of new models and approaches was desired. I looked around and talked to various people in the university and in industry and came up with two problems: one was the reliability of nuclear reactor safety systems where data from the first nuclear power stations were becoming available. This led me to develop techniques for reliability evaluation of highly reliable systems with a great deal of redundancy. The other problem I looked at, the result of contact with a member of an OR group at an automobile company, was the role of inventory banks in automatic transfer lines and the way in which they could increase the efficiency of the lines. This led to the development of probability models and the use of optimal control ideas to determine how to use the banks. This research went beyond the WWII approach, as the models were no longer relatively simple. Both problems were operational issues where mathematical programming was not relevant, although I used dynamic programming for determining the optimal operation of a transfer line.

## 7. OR to OM

The impact of this early OR on manufacturing was rapid in inventory management, perhaps in part because it helped justify the introduction of computers in the factory. The early books, like Brown's *Decision Rules for Inventory Management* [1] and Plossl and Wight's [9] practitioner-oriented book, already emphasized that safety stocks should be based on the variability of lead time demand. Requirement planning with time phasing would have been impossible without computers, and it was not until the widespread use of the IBM 360 series of computers in the late 1960s and early 1970s that MRP systems became attractive, even though many firms had great difficulty in implementation. This was often because their stock records did not adequately measure actual stock. Factories were not accustomed to the discipline required to have accurate inventory records. Although aggregate planning attracted much academic interest, my impression was that industry did not use any of the OR models, rather, as with our general sales manager, industry used only the general insights. Similarly, OR models for production scheduling were rarely used apart from scheduling a single machine. Shop floor data collection systems were too inadequate until the 1980s and the widespread use of bar codes. Perhaps a more significant barrier was that OR models rarely distinguished between what Plossl and Wight called "Input Control," decisions about the work packages to be released for manufacturing each time period, and "Output Control," the detailed scheduling of jobs within the manufacturing unit.

Even though the early industrial OR was quite successful, very few OR groups in industry survived the 1970s with its high inflation and stagnant growth. It was too easy for senior managers to observe that OR rarely brought immediate bottom-line impact. However, the problems OR people in industry had identified were taken up by academics and consultants. Academics wrote papers on the problems, often changing them so that the assumptions made no longer fit any real-world situation. Consultants developed software that implemented many of the insights and understanding from OR in industry. Academics started to write texts on production management (such as Peterson and Silver [8]), which incorporated the OR models and approaches to forecasting, inventory control, production planning, and production scheduling. As texts became more similar in the domain covered and a critical number of academics made this domain the prime focus of their research, it became apparent that a new academic discipline was emerging, *operations management*. Awareness of the success of Japanese manufacturing in the 1980s led to a full integration of quality control into OM and also to academic interest in understanding the basis of Kanban control and its relative advantages and disadvantages compared with MRP. With the 1990s came technical societies such as POMS (Production and Operations Management Society) and MSOM (Manufacturing and Service Operations Management Society), with OM-focused journals and conferences. It was not long before OM incorporated other strands of research not based on the traditional OR insights, such as behavioral OM or economics-based models of contractual relationships. Also OM began to focus more on services rather than manufacturing. In services, as with manufacturing, there was pioneering OR work done in the 1950s and 1960s, particularly in queueing and appointment system design, which eventually became an important base for academic interest in service systems.

## 8. Conclusion

Personally, my early experience of practicing OR in industry has been of great value. It has enabled me to suggest possible research directions to PhD students and choose research problems that seem relevant. As methodology, techniques, and technology have developed, it has been possible to revisit some of these problem areas many years later and gain new insights.

Since there are now so few OR groups in manufacturing firms, it would be surprising if they were the source of many new issues and models. An exception may be the semiconductor industry, where the required investment in a wafer fab is so large that OR can still make a contribution that impacts the bottom line. Also the technology is still sufficiently new that it has attributes of pioneering OR. It is also possible that the biotechnology industry may require new models and approaches. The present limited detailed involvement by OR practitioners in manufacturing industries may pose risks for the future of OM. To be viable and vibrant it has to be exposed to new problems, new issues, and new technology. Academics can construct many imagined and even real problems and issues, but OR people working in industry are forced to ensure that senior managers view their work as relevant to the needs of the firm.

*Analytics* has attracted much interest lately. In many ways analytics has considerable similarity to traditional OR, but updated to reflect that nowadays firms have a great deal of data that are kept for many years, and also that there are now new statistical techniques that need more than the hand or pocket calculators used by early OR practitioners. Analytics tends to focus on operational issues, so perhaps it will identify some new problems for OM.

### Acknowledgments

The author would like to acknowledge his debt to two of his early mentors and collaborators in the practice of OR, Mike Luck and Mike Duff.

## References

[1] R. G. Brown. *Decision Rules for Inventory Management.* Holt, Rinehart and Winston, New York, 1967.

[2] J. A. Buzacott. Dynamic inventory targets revisited. *Journal of the Operational Research Society* 50(7):697–703, 1999.

[3] A. D. Chandler, Jr. *Scale and Scope: The Dynamics of Industrial Capitalism.* Belknap Press, Cambridge, MA, 1990.

[4] C. W. Churchman, R. L. Ackoff, and E. L. Arnoff. *Introduction to Operations Research.* John Wiley & Sons, New York, 1957.

[5] W. Feller. *An Introduction to Probability Theory and Its Applications*, Vol. 1. John Wiley & Sons, New York, 1950.

[6] H. L. Lee, V. Padmanabhan, and S. Whang. Information distortion in a supply chain: The bullwhip effect. *Management Science* 43(4):546–558, 1997.

[7] F. G. Moore. *Production Control.* McGraw-Hill, New York, 1951.

[8] R. Peterson and E. A. Silver. *Decision Systems for Inventory Management and Production Planning.* John Wiley & Sons, New York, 1979.

[9] G. W. Plossl and O. W. Wight. *Production and Inventory Control: Principles and Techniques.* Prentice Hall, Englewood Cliffs, NJ, 1967.

[10] R. Roundy. 98%-Effective integer-ratio lot-sizing for one warehouse multi-retailer systems. *Management Science* 31(11):1416–1430, 1985.