



The Business School
for the World®

Working Paper

2020/51/TOM

People-Centric Operations: Achievements and Future Research Directions

Guillaume Roels
INSEAD, guillaume.roels@insead.edu

Bradley R. Staats
University of North Carolina, Bradley_Staats@kenan-flagler.unc.edu

As the nature of work has become more service-oriented, knowledge-intensive, and rapidly changing, people—be they workers or customers—have become more central to operational processes and have impacted operational outcomes in novel and perhaps more fundamental ways. Research in people-centric operations (PCO) studies how people affect the performance of operational processes. In this OM Forum, we define PCO as an area of study, offer a categorization scheme to take stock of where the field has allocated its attention to date, and offer our thoughts on promising directions for future research. The future of PCO is bright: Thanks to today's availability of granular data, PCO researchers have numerous and growing opportunities to study, from both descriptive and prescriptive angles, the link between people's behavior and operational performance.

Keywords: People; Operations Management; Forum; Clustering Analysis; Natural Language Processing

Electronic copy available at: <http://ssrn.com/abstract=3719847>

Working Paper is the author's intellectual property. It is intended as a means to promote research to interested readers. Its content should not be copied or hosted on any server without written permission from publications.fb@insead.edu

Find more INSEAD papers at <https://www.insead.edu/faculty-research/research>

Copyright © 2020 INSEAD

OM FORUM

People-Centric Operations: Achievements and Future Research Directions

Guillaume Roels & Bradley R. Staats

INSEAD, Guillaume.Roels@insead.edu

Kenan-Flagler Business School, University of North Carolina, Bradley_Staats@kenan-flagler.unc.edu

Abstract

As the nature of work has become more service-oriented, knowledge-intensive, and rapidly changing, people—be they workers or customers—have become more central to operational processes and have impacted operational outcomes in novel and perhaps more fundamental ways. Research in people-centric operations (PCO) studies *how people affect the performance of operational processes*. In this OM Forum, we define PCO as an area of study, offer a categorization scheme to take stock of where the field has allocated its attention to date, and offer our thoughts on promising directions for future research. The future of PCO is bright: Thanks to today's availability of granular data, PCO researchers have numerous and growing opportunities to study, from both descriptive and prescriptive angles, the link between people's behavior and operational performance.

1. Introduction

Research in people-centric operations (PCO) studies *how people affect the performance of operational processes*. The goal of this OM Forum is to characterize this emerging field of research, identify the multiple streams of research it encompasses, and outline promising directions for future research. For many years, most research in Operations Management (OM) has treated people in operating systems as fixed, unchanging, or exogenous entities. However, this assumption has been increasingly challenged, particularly with the rise of service and knowledge-intensive businesses where workers and/or customers fundamentally impact operational outcomes. Research in PCO focuses on the operational significance of having people (workers, customers, or both) interacting with and/or within an operational system. A key distinction with other academic disciplines is that PCO does not study behavior solely for the

sake of behavior, but rather it studies how behavior changes the performance of operational processes. Sometimes, it takes a step further back to also study how the design of operational processes shapes behavior, which then affects process performance.

PCO research is rarely recognized as a mainstream field in OM research, for instance, it does not have a Special Interest Group associated with it. However, its roots go back to the foundations of the OM field – Scientific Management – and to one of the most fundamental transformations within operations in the 20th century – Lean Production. In terms of the former, individuals such as Carl Barth, Henry Gantt, Frank and Lillian Gilbreth, and Frederick Taylor, the founders and early proponents of Scientific Management (Smiddy and Naum 1954), extensively studied how people affect the performance of operational processes. Specifically, they investigated how differences in output across people could be eliminated and in turn how new steps could be introduced to improve, primarily, the efficiency of the work completed. With respect to the latter, Lean Management, which emerged from the Toyota Production System (TPS), heavily relies on people for continuous operational change (Ohno 1988). When Toyota was looking to compete against the dominant auto manufacturers of the mid-20th century they realized that they could not leverage the same economies of scale due to their smaller size. However, that did not mean that competing through operations was foolhardy (Womack, Jones et al. 1990). Rather, by understanding how people affect and can improve operational processes, they were able to succeed (MacDuffie 1995, Shah and Ward 2003). Perhaps because of the work of these OM pioneers, the productivity gains achieved in manufacturing and agriculture, which were the dominant sectors of activity prior to the 1950s, have been so large that these sectors are now much less labor-intensive and their relative gross domestic product (GDP) contribution is much smaller.

Given these past achievements, one may wonder: Is PCO still a relevant field today? We believe it is the case, and in fact more than ever, for the following reasons. First, the growth of the service sector has heightened the importance of people in the processes of value creation and delivery. For many services (e.g., tourism, retail, banking, online platforms), customer experience is a critical differentiator once basic needs are fulfilled (Pine and Gilmore 1999). As

a result, front-line employees have key roles and responsibilities, to enable innovation, experimentation, and operational excellence. Moreover, customers in service settings, unlike in manufacturing, are often active participants in the processes of production and delivery (e.g., producing and sharing content on social networks), and technology has enabled many business model innovations in that regard. Second, with the increasing use of technology and the creation of more information

(Karmarkar 2015), how individuals and firms use such information often defines success. For example, the industrialization process (e.g., automation, offshoring) that has taken place in information-intensive sectors (e.g., consulting, finance, software), far from diminishing the role of people, has enhanced it as people are essential to solve increasingly complex problems and coordinate solutions. In fact, most of the jobs today are information-related, and these jobs are associated with a significantly higher wage than physical jobs (Apte, Karmarkar, & Nath, 2015). Third, and in part because of the first two points, the operational environment continues to change rapidly and competition has intensified (D'Aveni and Gunther 1994). In such settings, initial moves are important, but organizations must adapt and learn in order to succeed. As was discovered by the early pioneers of Scientific Management and Lean Management, learning rests in large part on the interaction between people and operational processes (Staats 2018). Perhaps it is not surprising then that learning systems like Lean Management have been extended to numerous other settings outside of manufacturing (e.g., Shah, Goldstein et al. 2008, Staats and Upton 2011). For all these reasons, we believe it has never been more relevant to study PCO.

In this overview of PCO we have three interrelated objectives. First, we wish to define the area of study and highlight the areas upon which it focuses. Second, we will offer a categorization scheme for PCO. This will permit us to take stock of where the field has allocated its attention to date. Third, we will offer our thoughts on promising directions in which the field may move going forward.

2. Definition and Characterization of PCO

As noted above, we define PCO as *the study of how people affect the performance of operational processes*. Through its evolution over time, research in PCO has been punctuated with multiple review or perspective works (e.g., Boudreau, Hopp et al. 2003, Hopp and Spearman 2008, Bendoly, Croson et al. 2010, KC 2019), which have shaped our operational understanding of the field by outlining fundamental principles:

- *People, not organizations, are self-optimizing.* Although research often focuses on organizations and the actions “they” take, in reality the anthropomorphic organization is a simplifying assumption: It is people within the organization who make choices. PCO takes the perspective of individual decision-makers across different roles and offers novel insight into the operational functioning of the processes with which they interact.
- *People are different.* People have different goals, different beliefs, different skills, different rates of change (e.g., learning). In some cases assuming homogeneity may be acceptable, but differences across the world, within an operating unit, or even with a person may alter preferred strategies.
- *People change with time.* Unlike finished goods inventory that may sit in a warehouse, people are never “finished”: they add skills with use and require training in order to maintain skills. Although they can carry large burdens for finite lengths of time, they also can experience burnout. To add to the complexity these patterns do not always take the same shape over time. As a result, incorporating learning, improvement, and behavioral change in our working models is necessary to gain a deep understanding of operations.
- *People have discretion.* This discretion can manifest itself in numerous ways, but the implications of individual choice are vast. People choose whether to take a given action, individually or collectively (e.g., customer coalitions) or to even leave an operating system (attrition). If elements of the system design (e.g., responsibility and authority) are not aligned, then discretion is likely to be used unproductively and operational performance may suffer.

Taken together these four elements not only help to characterize work in PCO, but also highlight the contribution the work makes to the field.

3. Categorization of Past Research on PCO

To take stock of the PCO research to date, we reviewed the abstracts of all research articles published in M&SOM (Volumes 1-21). We chose M&SOM because of its reputation as one of the flagship OM journals; our intent was indeed to capture PCO research at the core of the OM field. We included OM Practice, commissioned papers, and datasets, but excluded forums, introductions to special issues, errata, and commentaries. Through multiple iterations and extensive sharing of our points of view, we selected articles using the following criterion: Based on its abstract, does this article study how people (as opposed to, say, widgets) affect the performance of operational processes? This resulted in a list of 98 research articles relevant to PCO, out of a pool of 679 articles. (The list of articles' DOIs appears in Appendix A.) An important limitation of our selection procedure is that we limited ourselves to papers published in M&SOM and thus ignored papers published elsewhere or before the launch of the journal in 2000.

We used this selected list of articles to build a "top-down" classification scheme of PCO research, as outlined in Section 3.1, which we applied to our pool of articles to identify the dominant research areas and trends (Section 3.2). We also used these abstracts to identify, in a "bottom-up" fashion based on keyword frequency, key clusters of research (Section 3.3).

3.1. Dimensions: Key Paradigms, Loci of Analysis, and Themes

After reading the abstracts of the 98 selected articles, we identified three dimensions of interest. The first dimension is around the research paradigm used: descriptive and prescriptive. Descriptive research can take multiple forms. For example, it not only involves exploring underlying relationships, but also yields results that map processes and relationships and then test or evaluate these underlying models. By prescriptive research we include papers that recommend new solutions and then show the implications of these solutions. In this category we would also include predictive papers that seek to forecast future outcomes. We note that to date PCO descriptive work is mostly, but not exclusively, empirical in paradigm, while prescriptive work is mostly, but not exclusively, analytical in paradigm.

The second dimension in our categorization is the locus of analysis of the study. Although PCO focuses on the impact of people, this impact can be seen at the micro (individual), meso (team), or macro (organization or market) level. To be sure, actions are always made at the individual level, but in some papers, individuals are homogeneous and their individual actions impact processes only at a more aggregate level (e.g., at the market level).

The third dimension captures the themes of the work. We note that these themes, although broad, are not exhaustive of all possible avenues future work could or even should explore. Nevertheless, it is important to start somewhere and we believe that these categories accurately capture the literature in its present state. We identified five themes, in which we were able to place all papers:

- Utility theory: rational decision making without externalities
- Strategic behavior: rational decision making that accounts for the expected actions of others.
- Behavioral biases: decision making that differs from rational models
- Learning and productivity: how individual actions dynamically impact outcomes
- Coordination and trust: how individual actions dynamically impact processes.

To illustrate the approach, we consider here several examples from seminal papers in PCO, deliberately chosen before M&SOM was launched to indicate the breadth and depth of the field:

- In the first volume of *Management Science*, Marschak (1955) proposes elements for a theory of teams, building on the nascent theory of games, with implications for organizational structure. The study is *prescriptive*; its locus of analysis is at the level of the *team*, and its theme is *coordination and trust*.
- Departing from standard queuing models, Naor (1969) explicitly models the customers' queue joining decision, trading off the net value they receive from the service with their expected wait time, and its implications on the service provider's pricing decision. The study is *prescriptive* because it leads to pricing prescriptions; its locus of analysis is at the

level of the *market*, which is an aggregation of individual decision makers; and its theme is *utility theory*.

- Sterman (1989) uses an experiment, based on the beer game, to show how individuals misperceive the consequences of past decisions. The study is *descriptive* in methodology, *individual* in locus of analysis and it studies *behavioral biases* around anchoring and adjustment. Under the same classification, OM researchers may also be familiar with the studies by Schultz et al. (1998) who experimentally show that assembly line operators' processing times are affected by the amount of buffer inventory and are therefore not independent of each other, and by Schweitzer and Cachon (2000), who explore decision biases in newsvendor ordering decisions.
- Adler and Clark (1991) use an empirical paradigm to understand how behavior shapes organizational learning curves. The paper is *descriptive* in investigating relationships, but it uses individual actions to understand the level of the *organization* around the theme of *learning and productivity*.

3.2. Dominant Research Areas and Trends

We applied the classification scheme outlined in the previous section to the selected pool of 98 papers. That is, we characterized each paper along three dimensions: research paradigm (2 levels: prescriptive vs. descriptive), locus of analysis (4 levels: individual, team, organization, market), and theme (5 levels, as defined above). In the few cases a paper overlapped across multiple levels, we used our best judgment to select the dominant one. The proposed classification of the 98 papers appears in Appendix A.

Table 1 below provides our summary findings in a 2 x 4 x 5 matrix. Beginning with the prescriptive research panel we see that there is a heavy focus on strategic behavior and utility theory at the level of the market. Looking at descriptive theory panel we find that the focus has most traditionally been at the individual level in the exploration of behavioral biases.

We note that the matrices contain numerous empty cells. Just because prior work has not filled in these cells, there is no inherent reason why they should remain empty. In fact,

Learning and Productivity in the descriptive panel shows that it is possible to fill all levels of analysis; in the prescriptive panel, this theme covers three of four loci of analysis, and future work could certainly conduct prescriptive productivity studies at the market level (e.g., how to optimize training of customers to adopt a new self-service technology). More generally, we hope that the identification of these empty cells will encourage enterprising researchers to pursue work in those areas (e.g., studying behavioral biases or strategic behavior in teams).

Table 1: Areas of PCO Research Published in M&SOM

Prescriptive Paradigm

| | | Locus of analysis | | | |
|-------|---------------------------|-------------------|------|--------------|--------|
| | | Individual | Team | Organization | Market |
| Theme | Behavioral biases | 2% | 0% | 1% | 3% |
| | Learning and productivity | 2% | 1% | 1% | 0% |
| | Coordination | 0% | 1% | 1% | 0% |
| | Strategic behavior | 0% | 0% | 0% | 16% |
| | Utility theory | 0% | 0% | 6% | 19% |

Descriptive Paradigm

| | | Locus of analysis | | | |
|-------|---------------------------|-------------------|------|--------------|--------|
| | | Individual | Team | Organization | Market |
| Theme | Behavioral biases | 17% | 0% | 2% | 4% |
| | Learning and productivity | 5% | 4% | 6% | 2% |
| | Coordination | 0% | 0% | 1% | 1% |
| | Strategic behavior | 0% | 0% | 0% | 1% |
| | Utility theory | 1% | 0% | 0% | 1% |

From Table 1, it appears that some cells, or groups of cells, are more representative of past research on PCO. To gain some insight into the evolution of certain streams of research over time, we grouped the past research into six key areas:

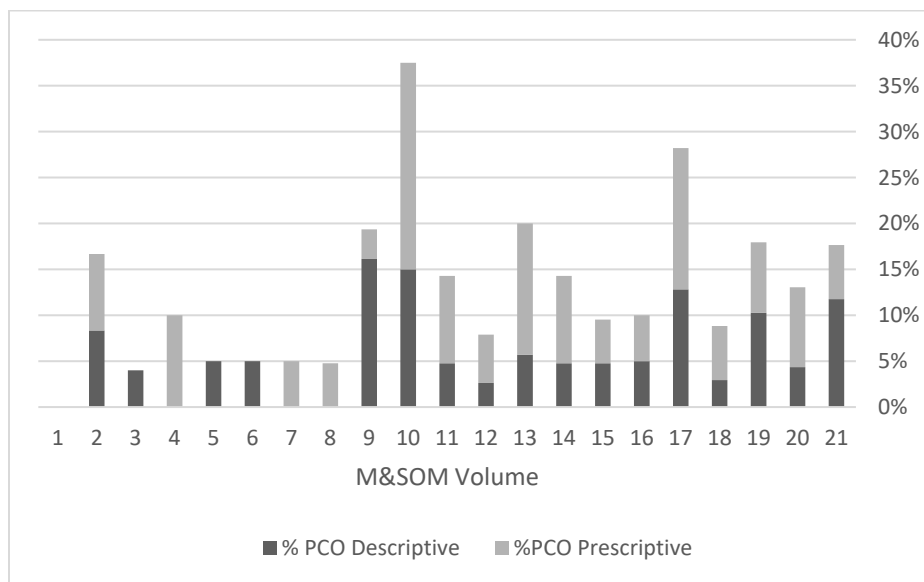
- Prescriptive; Market; Strategic behavior
- Prescriptive; Market; Utility theory
- Prescriptive; other
- Descriptive; Individual; Behavioral biases

- Descriptive; Individual/Team/Organization; Learning and productivity
- Descriptive; other

Although some of these areas are apparent in Table 1, our choice was guided by running a hierarchical clustering algorithm on our paper classification, as detailed in Appendix B.

Figure 1 depicts the evolution of the percentage of PCO articles published in M&SOM across the different volumes, broken down by research paradigm. At the outset, it is worth noting that Volume 9 contained a special issue on empirical research in OM, in which 5 out of 11 published papers were on PCO, and Volume 10 contained a special issue on behavioral OM, in which all 7 published papers were on PCO. With that in mind, we note a substantial increase in the share of PCO articles published in M&SOM, from about 5% prior to Volume 9 to about 15% after Volume 11. Prior to Volume 9, it was common to have only one PCO article published per volume (which had about 20 articles in total). In contrast after Volume 11, about 6 PCO articles are published per volume (which had on average 40 articles).

Figure 1: Evolution of the Percentage of PCO Articles Published in M&SOM



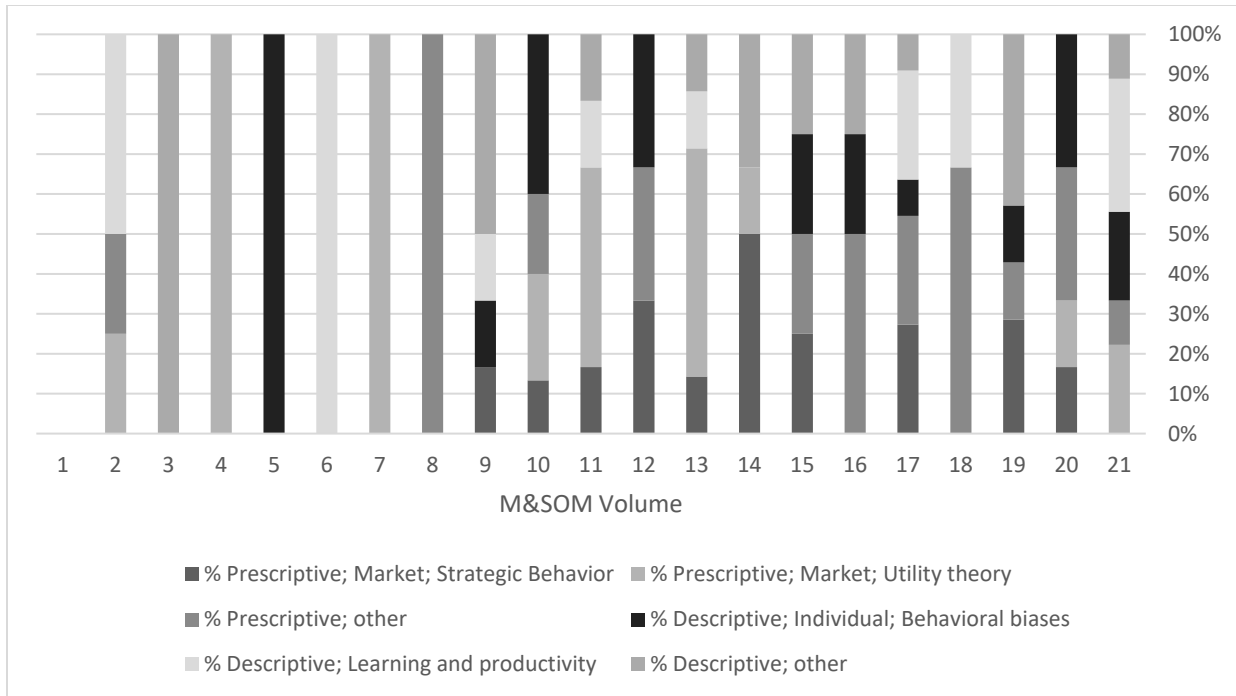
Although the special issues in Volumes 9 and 10 were certainly instrumental in creating more visibility on PCO research, we note that by design, they were focused on descriptive research. Yet, a large fraction of the subsequent PCO papers are prescriptive, so

not all growth in interest in the topic can be attributed to the special issues. In fact in most of the volumes after Volume 10, more than 50% of the PCO papers are prescriptive, with two exceptions (Volumes 19 and 21). The field of PCO research is thus strong and healthy, embracing a balance of research paradigms.

To provide a finer-grained view of the different streams of research, Figure 2 shows the evolution of the relative cluster shares of PCO research. From the figure, it is quite apparent that the shares of the different research areas have become more uniformly distributed over time. One reason is that there are more PCO papers published in M&SOM: Since volumes prior to Volume 9 contained one or two PCO papers, a cluster could easily then get a share of 100%.

Yet, even after the number of PCO papers increased (after Volume 9), it was common for a particular research area to dominate a particular volume. Like other fields in OM (Lariviere 2016), PCO research has been exposed to temporary surges of interest on particular topics as indicated by the large share of PCO research on “Descriptive; Individual; Behavioral biases” in Volume 10 and on “Prescriptive; Market; Strategic behavior” in Volume 14. Yet, it is reassuring that these areas of interest, unlike fads, which may come and go, have remained active beyond their peak and may have in fact been rejuvenated with novel research questions.

Figure 2: Relative Research Area Shares of PCO Articles Published in M&SOM



Furthermore, the more recent volumes have been more evenly distributed, allocating no more than 33% share to any particular research area. In parallel, there has been a growth in importance of the hybrid research areas, namely “Prescriptive; other” and “Descriptive; other”. In our opinion, this is another sign that the field of PCO research is healthy, with a wide distribution of interest across research areas and exploration of novel ideas and methods.

3.3. Keyword-Based Clustering Analysis

To complement and validate the top-down classification scheme outlined in Section 3.1., we performed a “bottom-up” clustering analysis based on the abstract keywords, using the nonnegative matrix factorization (NMF) method. This clustering scheme tends to distinguish papers based on their context of application (e.g., healthcare since it has many distinctive keywords such as “patient” and “physician”) and less on their research paradigms (since “prescriptive” and “descriptive” are rarely explicitly stated in the articles’ abstracts). A detailed outline of the method appears in Appendix C.

We identified six keyword-based clusters. Each cluster is characterized by a vector of weights on keywords; we reproduce below the highest-score keywords (with their scores in parentheses). In addition, we provide the paper that is the most representative of that cluster.

Table 2: Keyword-Based Clusters of PCO Research in M&SOM

| Cluster | Representative Keywords (Scores) | Representative Paper |
|----------------------|---|---|
| Product pricing | Price (0.78), consumer (0.55), strategic (0.38), seller (0.26), customer (0.22), product (0.22) | “Optimal Pricing of Seasonal Products in the Presence of Forward-Looking Consumers” (Aviv and Pazgal 2008) |
| Customer service | Customer (0.91), service (0.60), queue (0.28), capacity (0.22), server (0.19), firm (0.17) | “Joining Longer Queues: Information Externalities in Queue Choice” (Veeraraghavan and Debo 2009) |
| Process & bottleneck | Process (0.35), agent (0.28), network (0.27), time (0.26), management (0.23), capacity (0.21) | “Collaboration and Multitasking in Networks: Architectures, Bottlenecks, and Capacity” (Gurvich and Mieghem 2014) |
| Newsvendor ordering | Newsvendor (0.47), model (0.38), order (0.34), decision (0.34), inventory (0.21), quantity (0.20) | “Heterogeneity of Reference Effects in the Competitive Newsvendor Problem” (Kirshner and Ovchinnikov 2019) |
| Team knowledge | Team (0.75), knowledge (0.66), quality (0.20), task (0.18), organization (0.17), product (0.16) | “Fluid Teams and Knowledge Retrieval: Scaling Service Operations” (Huckman and Staats 2011) |
| Healthcare | Patient (0.75), appointment (0.54), clinic (0.23), wait (0.21), schedule (0.20), physician (0.16) | “Effects of Rescheduling on Patient No-Show Behavior in Outpatient Clinics” (Liu, Xie et al. 2019) |

From their set of representative keywords, the keyword-based clusters appear to be respectively about product pricing to strategic customers; customer service, and queueing systems in particular; process and bottleneck management (capacity and time), including revenue management, routing policies, time management, jobs and workforce scheduling, and fair process management; newsvendor ordering decisions; team knowledge and task quality in organizations; healthcare, studying appointment scheduling and waits. Appendix A provides the list of articles in each cluster.

Figure 3: Research Paradigm, Locus of Analysis, and Theme Explored by the Different Clusters




Figure 3 depicts the different research paradigms, loci of analysis, and themes of the papers in each cluster. The clusters on “product pricing” and “customer service” clearly study PCO at the level of market, respectively assuming strategic behavior and relying on standard utility theory. Although our top-down classification associated these fields of study with a prescriptive paradigm, the bottom-up keyword-based classification identified a few descriptive papers on these topics as well. The cluster on “process & bottleneck”, which is more hybrid, puts higher weights than the other clusters on the theme of coordination and on the

organizational locus of analysis; although it encompasses both research paradigms, it primarily adopts a prescriptive one. The “newsvendor ordering” cluster distinctively adopts a descriptive paradigm to study behavioral biases at the individual level. The “team knowledge” cluster is about learning and productivity, primarily at the organization or team level, using a descriptive paradigm. Finally, because the “healthcare” cluster is essentially defined by its context, it encompasses a variety of paradigms, loci of analysis, and themes.

Overall, this keyword-based clustering analysis validates and complements the classification scheme proposed in Section 3.1. It validates our earlier classification by showing a strong association, which can be verified with a frequency table, between the product pricing cluster and the “Prescriptive; Market; Strategic” category; between the newsvendor ordering cluster and the “Descriptive; Individual; Behavioral biases” category; between the team knowledge cluster and the “Descriptive; Individual/Team/Organization; Learning and productivity” category; and, to a smaller extent, between the customer service cluster and the “Prescriptive; Market; Utility theory”. It also complements our earlier classification by bringing contextual knowledge to the various topics explored, including healthcare, pricing decisions, newsvendor ordering decisions, as well as process and bottleneck management.

4. Future Research Directions

To conclude, we briefly reflect on future research directions for PCO. As highlighted in the introduction, the opportunities for research on PCO have become wider given the changing nature of work. Not only does the rise of service and knowledge work create more settings where PCO is relevant, their fast-changing pace calls for a deep understanding of the underlying fundamentals of PCO.

Fortunately, because these settings are typically information-rich, they provide an unprecedented opportunity to understand the complex interplay between behavior and performance. Today’s data technologies, such as RFID sensor data (Staats et al. 2016), retail

point of sale data (Perdikaki et al. 2012), or contact center records (Brown et al. 2005), are much more granular than ever. We have gone a long way from the earlier suggestion by Webb et al. (1966) to use trace data, such as indentations in carpet! As a result, today's granular level of data makes it possible to study performance at a micro level (Terwiesch 2019). Although there is an obvious benefit to descriptive research, prescriptive research can also uncover new opportunities in this information-rich environment. One opportunity could be to inform the development of new models that identify predictive and prescriptive steps to take. Another opportunity, enabled by the firms' digital transformation, would be to develop and deploy software to aid decisions (which can then lead to further descriptive research). Thus, both descriptive and prescriptive PCO research face numerous opportunities and the biggest opportunities lie in nurturing a continuous dialogue between these two research paradigms.

Following this line of discussion, it is fair to ask: Why don't we see greater alignment between prescriptive and descriptive research in the field? To be sure, this is a challenge in many research fields in OM. Researchers often, but not always, cluster by methods, rather than topic. Although understandable, there are many opportunities for learning from each other. For instance, the keyword-based clustering analysis summarized in Table 2 reveals that researchers on customer service, which tend to adopt queuing theory, could learn from empirical studies on customer efficiency;¹ and that researchers on newsvendor ordering decisions, who tend to use analytical modeling or lab experiments, could learn from empirical studies on the effect of experience on performance.² We call for PCO researchers to avoid falling into the trap of balkanization of the field. One of the purposes of this OM Forum is to show unity across research paradigms, loci of analysis, and themes and we hope that researchers belonging to a particular cluster of interest would remain open to ideas from other clusters.

¹ <https://doi.org/10.1287/msom.1060.0135>

² <https://doi.org/10.1287/msom.2017.0661>

Our review of PCO literature has identified a number of areas of research based on themes that we hope future work will explore. There is a need for empirical, analytical, and qualitative work to seek to understand more about the following topics:

- Coordination and trust –Gains from specialization have led individuals and firms to frequently become more focused on specific areas (Skinner 1974, KC and Terwiesch 2011, Staats and Gino 2012). However, complete solutions are still necessary for end-consumers and so specialized activities must be integrated requiring both coordination and trust.
- Prescriptive research on learning and productivity – Although learning and productivity is well represented on the descriptive side of PCO, and there continues to be ample opportunities for this work given the vital role that learning plays in operational performance and competitive advantage, we call for more prescriptive work. Whether in models of how the learning process impacts operations (Roels 2020) or in decision support tools that can change learning and productivity in practice, significant opportunities exist.
- Markets with co-producing customers --- Even though PCO research has endogenized customers' purchasing decision (what, how much, and when), customers' roles in many service settings are often much broader and instrumental to value creation; for instance in education, students' active participation is key to learning. There are numerous research opportunities on coordinating the employee-client dyad (Rahmani et al. 2017) and on managing customers as productive resources, with heterogeneous levels of efficiency, similar to employees (Xue and Harker 2002).
- Organization and Teams – Traditionally, most of PCO research has focused on the individual level, which is an understandable starting point. However, given the data-rich nature of most business environments, it is increasingly possible to study organizations and teams. These higher levels of organizing are important cannot be extrapolated from our understanding of individual behavior because operational performance is rarely just a summation of activities. Rather in teams and

organizations, the interplay is often complex and non-linear and this should be incorporated in more research.

- Strategic behavior of employees – Although prescriptive research on PCO has extensively studied strategic behavior of customers and there is ample ongoing opportunity for exploration in this area, we call for more research to assess the strategic behavior of employees on performance, especially from a prescriptive perspective (Armony et al. 2020). How do the strategic choices of employees alter operational outcomes and vice versa? One could argue that the field of organizational behavior arose as a result of individuals' strategic response to scientific management (Mayo 1933, Roethlisberger and Dickson 1934, Roy 1959). Organizational behavior moved on from this important area of PCO, but OM researchers could and should return to it.
- Other --- The world is full of challenges summarized by the UN Sustainable Development Goals (e.g., climate change, poverty, gender inequality, weak institutions); and at the core of each of them lie people. To tackle these challenges, we need multi-disciplinary approaches and a willingness to push the OM field further. Similarly, as technology (e.g., AI, Internet of Things) becomes more ubiquitous, there is a stronger need than ever to understand the relationship between technology, humans, and operational performance. As shown in Figure 2, PCO research has already started exploring novel combinations of the various research paradigms, loci of analysis, and themes, and we hope such exploratory efforts will be directed to tackling these emerging topics

As a final note, we echo the call of Boudreau et al. (2003) for further work studying the human resources (HR) function and its operational impact. Interestingly, the HR function at such companies as Google or Netflix is now known as “People Operations”. Despite this label it is fair to ask: How many HR professionals have a deep understanding of operations? In addition, how many OM researchers have immersed themselves in HR? This area of the firm is not only the repository of much of the data that many researchers might seek to study, but it is also where critical decisions are made about the hiring, allocation, and development of people.

Critical questions of HR, such as attrition (Emadi and Staats 2020), lend themselves to careful descriptive and prescriptive study. Although HR is not the typical domain of the operations professional, it is one in which both the practitioner and the academic can learn and provide value.

To summarize, we stress that PCO is not a new area of study in OM – its roots go back to the very founding of the field. However, we have seen a recent increase in interest as people (customers, workers) have become more central than ever to operational performance and as more granular data permits the study of past action and prescription for future action. We hope that this OM Forum not only shines a light on the significant work that has been done in the field, but through our proposed framework, it provides a path forward for ever more thoughtful work on PCO. Together we both believe that research in PCO is strong and healthy, but it has just scratched the surface of what can be done and we are excited about the future we will create together.

Acknowledgments

We would like to thank the Editor-in-Chief, Chris Tang, for encouraging us to write this OM Forum and for help us push the boundaries of the OM field further by publishing a special issue on the topic. We would also like to thank Rafa Ballestiero for special research assistance with the clustering algorithms.

Bibliography

- Adler, P. S. and K. B. Clark (1991). "Behind the learning curve: A sketch of the learning process." Management Science **37**(3): 267-281.
- Armony, M., G. Roels and H. Song (2020). "Pooling queues with strategic servers: The effects of customer ownership." Operations Research.
- Aviv, Y. and A. Pazgal (2008). "Optimal pricing of seasonal products in the presence of forward-looking consumers." Manufacturing & Service Operations Management **10**(3): 339-359.
- Bendoly, E., R. Croson, P. Goncalves and K. Schultz (2010). "Bodies of knowledge for research in behavioral operations." Production and Operations Management **19**(4): 434-452.

Boudreau, J., W. Hopp, J. O. McClain and L. J. Thomas (2003). "On the interface between operations and human resources management." Manufacturing & Service Operations Management **5**(3): 179-202.

Brown, L., N. Gans, A. Mandelbaum, A. Sakov, H. Shen, S. Zeltyn and L. Zhao (2005). "Statistical analysis of a telephone call center." Journal of the American Statistical Association **100**(469): 36-50.

D'Aveni, R. A. and R. E. Gunther (1994). Hypercompetition : managing the dynamics of strategic maneuvering. New York, The Free Press.

Emadi, S. and B. R. Staats (2020). "A Structural Estimation Approach to Study Agent Attrition." Management Science **66**(9): 4071-4095.

Gurvich, I. and J. A. V. Mieghem (2014). "Collaboration and multitasking in networks: Architectures, bottlenecks and throughput." Manufacturing & Service Operations Management **17**(1): 16-33.

Hopp, W. J. and M. L. Spearman (2008). Factory Physics. New York, NY, McGraw-Hill.

Huckman, R. S. and B. R. Staats (2011). "Fluid tasks and fluid teams: The impact of diversity in experience and team familiarity on team performance." Manufacturing & Service Operations Management **13**(3): 310-328.

Karmarkar, U. (2015). "The service and information economy: Research opportunities." Manufacturing & Service Operations Management **17**: 136-141.

KC, D. (2019). "Worker productivity in operations management." Foundations and Trends in Technology, Information, and Operations Management.

KC, D. and C. Terwiesch (2011). "The effects of focus on performance: Evidence from California hospitals." Management Science **57**(11): 1897-1912.

Kirshner, S. N. and A. Ovchinnikov (2019). "Heterogeneity of reference effects in the competitive newsvendor problem." Manufacturing & Service Operations Management **21**(3): 571-581.

Kuang, D., J. Choo and H. Park (2015). Nonnegative matrix factorization for interactive topic modeling and document clustering. Partitional Clustering Algorithms, Springer: 215-243.

Lariviere, M. A. (2016). "OM Forum—Supply chain contracting: Doughnuts to bubbles." Manufacturing & Service Operations Management **18**(3): 309-313.

Liu, J., J. Xie, K. K. Yang and Z. Zheng (2019). "Effects of rescheduling on patient no-show behavior in outpatient clinics." Manufacturing & Service Operations Management **21**(4): 780-797.

MacDuffie, J. P. (1995). "Human resource bundles and manufacturing performance." Industrial and Labor Relations Review **48**(2): 197-221.

Marschak, J. (1955). "Elements for a theory of teams." Management Science **1**(2): 127-137.

Mayo, E. (1933). The Human Problems of an Industrial Civilization. New York, Macmillan.

Naor, P. (1969). "The regulation of queue size by levying tolls." Econometrica **37**(1): 15-24.

Ohno, T. (1988). Toyota Production System: Beyond Large-Scale Production. Cambridge, MA, Productivity Press.

Perdikaki, O., S. Kesavan and J. M. Swaminathan (2012). "Effect of retail store traffic on conversion rate and sales." Manufacturing & Service Operations Management **14**(1): 145-162.

Pine, J. and J. H. Gilmore (1999). The experience economy: Work is theatre & every business a stage. Boston, Harvard Business Press.

Rahmani, M., G. Roels and U. S. Karmarkar (2017). "Collaborative work dynamics in projects with co-production." Production and Operations Management **26**(4): 686-703.

Roels, G. (2020). "High-performance practice processes." Management Science **66**(4): 1509-1526.

Roethlisberger, F. J. and W. J. Dickson (1934). Management and the Worker. Boston, Harvard Univ.

Roy, D. F. (1959). "'Banana time': Job satisfaction and informal interaction." Human Organization **18**(4): 158-168.

Schultz, K. L., D. C. Juran, J. W. Boudreau, J. O. McClain and L. J. Thomas (1998). "Modeling and worker motivation in JIT production systems." Management Science **44**(12): 1595-1607.

Schweitzer, M. E. and G. P. Cachon (2000). "Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence." Management Science **46**(3): 404-420.

Shah, R., S. M. Goldstein, B. T. Unger and T. D. Henry (2008). "Explaining anomalous high performance in a health care supply chain." Decision Sciences **39**(4): 759-789.

Shah, R. and P. T. Ward (2003). "Lean manufacturing: Context, practice bundles, and performance." Journal of Operations Management **21**(2): 129-149.

Skinner, W. (1974). "The focused factory." Harvard Business Review **52**(3): 113-121.

Smiddy, H. and L. Naum (1954). "Evolution of a 'Science of Managing' in America." Management Science **1**(1): 1-31.

Staats, B. R. (2018). Never Stop Learning. Boston, Harvard Business Review Press.

Staats, B. R., H. Dai, D. Hofmann and K. L. Milkman (2016). "Motivating process compliance through individual electronic monitoring: An empirical examination of hand hygiene in healthcare." Management Science **63**(5): 1563-1585.

Staats, B. R. and F. Gino (2012). "Specialization and variety in repetitive tasks: Evidence from a Japanese bank." Management Science **58**(6): 1141-1159.

Staats, B. R. and D. M. Upton (2011). "Lean Knowledge Work." Harvard Business Review **89**(10): 100-110.

Sterman, J. D. (1989). "Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment." Management Science **35**(3): 321-339.

Terwiesch, C. (2019). "OM Forum—Empirical Research in Operations Management: From Field Studies to Analyzing Digital Exhaust." Manufacturing & Service Operations Management **21**(4): 713-722.

Veeraraghavan, S. and L. Debo (2009). "Joining longer queues: Information externalities in queue choice." Manufacturing & Service Operations Management **11**(4): 543-562.

Webb, E. J., D. T. Campbell, R. D. Schwartz and L. Sechrest (1966). Unobtrusive Measures: Nonreactive Research in the Social Sciences. Chicago, Rand McNally.

Womack, J. P., D. T. Jones and D. Roos (1990). The Machine That Changed the World. New York, Rawson Associates.

Xue, M. and P. T. Harker (2002). "Customer efficiency: Concept and its impact on e-business management." Journal of Service Research **4**(4): 253-267.

Appendix A: List and Classification of Selected PCO Research Papers Published in M&SOM in 2000-19.

| DOI | Research Paradigm | Locus of Analysis | Research Theme | Keyword-Based Cluster |
|---|-------------------|-------------------|---------------------------|-----------------------|
| https://doi.org/10.1287/msom.1030.0030 | Descriptive | Market | Strategic behavior | process & bottleneck |
| https://doi.org/10.1287/msom.1050.0088 | Prescriptive | Market | Utility theory | product pricing |
| https://doi.org/10.1287/msom.1060.0106 | Prescriptive | Organization | Learning and Productivity | process & bottleneck |
| https://doi.org/10.1287/msom.1060.0130 | Descriptive | Individual | Behavioral biases | newsvendor ordering |
| https://doi.org/10.1287/msom.1060.0131 | Descriptive | Organization | Learning and Productivity | team knowledge |
| https://doi.org/10.1287/msom.1060.0135 | Descriptive | Market | Coordination | customer service |
| https://doi.org/10.1287/msom.1060.0145 | Prescriptive | Market | Strategic behavior | newsvendor ordering |
| https://doi.org/10.1287/msom.1060.0147 | Prescriptive | Market | Behavioral biases | customer service |
| https://doi.org/10.1287/msom.1060.0148 | Prescriptive | Organization | Utility theory | process & bottleneck |
| https://doi.org/10.1287/msom.1060.0150 | Descriptive | Individual | Learning and Productivity | process & bottleneck |
| https://doi.org/10.1287/msom.1060.0156 | Descriptive | Organization | Learning and Productivity | team knowledge |
| https://doi.org/10.1287/msom.1060.0190 | Descriptive | Individual | Behavioral biases | newsvendor ordering |
| https://doi.org/10.1287/msom.1070.0157 | Prescriptive | Market | Strategic behavior | product pricing |
| https://doi.org/10.1287/msom.1070.0169 | Prescriptive | Market | Utility theory | process & bottleneck |
| https://doi.org/10.1287/msom.1070.0177 | Prescriptive | Market | Utility theory | customer service |
| https://doi.org/10.1287/msom.1070.0180 | Prescriptive | Market | Utility theory | customer service |
| https://doi.org/10.1287/msom.1070.0183 | Prescriptive | Market | Strategic behavior | product pricing |
| https://doi.org/10.1287/msom.1070.0188 | Descriptive | Individual | Behavioral biases | newsvendor ordering |
| https://doi.org/10.1287/msom.1070.0189 | Prescriptive | Organization | Coordination | process & bottleneck |
| https://doi.org/10.1287/msom.1070.0193 | Descriptive | Market | Behavioral biases | customer service |
| https://doi.org/10.1287/msom.1070.0194 | Prescriptive | Market | Utility theory | customer service |
| https://doi.org/10.1287/msom.1070.0200 | Descriptive | Individual | Behavioral biases | newsvendor ordering |
| https://doi.org/10.1287/msom.1070.0205 | Descriptive | Individual | Behavioral biases | process & bottleneck |
| https://doi.org/10.1287/msom.1070.0210 | Prescriptive | Market | Utility theory | process & bottleneck |
| https://doi.org/10.1287/msom.1080.0228 | Descriptive | Individual | Behavioral biases | newsvendor ordering |
| https://doi.org/10.1287/msom.1080.0229 | Descriptive | Individual | Behavioral biases | newsvendor ordering |
| https://doi.org/10.1287/msom.1080.0233 | Descriptive | Organization | Learning and Productivity | team knowledge |
| https://doi.org/10.1287/msom.1080.0239 | Prescriptive | Market | Strategic behavior | customer service |
| https://doi.org/10.1287/msom.1080.0240 | Prescriptive | Market | Utility theory | product pricing |
| https://doi.org/10.1287/msom.1080.0244 | Prescriptive | Market | Utility theory | customer service |
| https://doi.org/10.1287/msom.1080.0248 | Prescriptive | Market | Strategic behavior | product pricing |
| https://doi.org/10.1287/msom.1080.0250 | Prescriptive | Organization | Behavioral biases | process & bottleneck |
| https://doi.org/10.1287/msom.1100.0294 | Descriptive | Individual | Behavioral biases | newsvendor ordering |
| https://doi.org/10.1287/msom.1100.0299 | Descriptive | Market | Utility theory | customer service |
| https://doi.org/10.1287/msom.1100.0306 | Prescriptive | Market | Strategic behavior | product pricing |
| https://doi.org/10.1287/msom.1100.0321 | Descriptive | Team | Learning and Productivity | team knowledge |
| https://doi.org/10.1287/msom.1100.0322 | Prescriptive | Market | Utility theory | customer service |

| | | | | |
|---|--------------|--------------|---------------------------|----------------------|
| https://doi.org/10.1287/msom.1110.0332 | Prescriptive | Market | Utility theory | healthcare |
| https://doi.org/10.1287/msom.1110.0338 | Prescriptive | Market | Utility theory | process & bottleneck |
| https://doi.org/10.1287/msom.1110.0344 | Prescriptive | Market | Utility theory | product pricing |
| https://doi.org/10.1287/msom.1110.0360 | Prescriptive | Market | Strategic behavior | customer service |
| https://doi.org/10.1287/msom.1120.0376 | Prescriptive | Market | Strategic behavior | product pricing |
| https://doi.org/10.1287/msom.1120.0377 | Prescriptive | Market | Strategic behavior | product pricing |
| https://doi.org/10.1287/msom.1120.0384 | Descriptive | Individual | Learning and Productivity | healthcare |
| https://doi.org/10.1287/msom.1120.0385 | Descriptive | Individual | Learning and Productivity | team knowledge |
| https://doi.org/10.1287/msom.1120.0394 | Prescriptive | Market | Utility theory | healthcare |
| https://doi.org/10.1287/msom.1120.0400 | Descriptive | Individual | Behavioral biases | customer service |
| https://doi.org/10.1287/msom.1120.0406 | Prescriptive | Market | Strategic behavior | customer service |
| https://doi.org/10.1287/msom.1120.0411 | Prescriptive | Organization | Utility theory | customer service |
| https://doi.org/10.1287/msom.2.1.32.23268 | Prescriptive | Organization | Utility theory | customer service |
| https://doi.org/10.1287/msom.2.2.144.12353 | Prescriptive | Market | Utility theory | customer service |
| https://doi.org/10.1287/msom.2.3.221.12348 | Descriptive | Organization | Coordination | process & bottleneck |
| https://doi.org/10.1287/msom.2.4.410.12339 | Descriptive | Individual | Utility theory | product pricing |
| https://doi.org/10.1287/msom.2013.0439 | Descriptive | Organization | Learning and Productivity | process & bottleneck |
| https://doi.org/10.1287/msom.2013.0453 | Descriptive | Individual | Behavioral biases | newsvendor ordering |
| https://doi.org/10.1287/msom.2013.0464 | Descriptive | Individual | Learning and Productivity | process & bottleneck |
| https://doi.org/10.1287/msom.2014.0486 | Prescriptive | Organization | Utility theory | healthcare |
| https://doi.org/10.1287/msom.2014.0493 | Prescriptive | Organization | Utility theory | process & bottleneck |
| https://doi.org/10.1287/msom.2014.0498 | Prescriptive | Team | Coordination | process & bottleneck |
| https://doi.org/10.1287/msom.2014.0499 | Prescriptive | Market | Strategic behavior | product pricing |
| https://doi.org/10.1287/msom.2014.0500 | Descriptive | Individual | Behavioral biases | newsvendor ordering |
| https://doi.org/10.1287/msom.2014.0501 | Prescriptive | Market | Strategic behavior | product pricing |
| https://doi.org/10.1287/msom.2014.0504 | Descriptive | Market | Behavioral biases | product pricing |
| https://doi.org/10.1287/msom.2014.0507 | Prescriptive | Team | Learning and Productivity | team knowledge |
| https://doi.org/10.1287/msom.2015.0522 | Descriptive | Organization | Learning and Productivity | team knowledge |
| https://doi.org/10.1287/msom.2015.0523 | Descriptive | Market | Learning and Productivity | customer service |
| https://doi.org/10.1287/msom.2015.0527 | Prescriptive | Market | Strategic behavior | product pricing |
| https://doi.org/10.1287/msom.2015.0532 | Prescriptive | Market | Behavioral biases | product pricing |
| https://doi.org/10.1287/msom.2015.0537 | Descriptive | Organization | Behavioral biases | product pricing |
| https://doi.org/10.1287/msom.2015.0547 | Descriptive | Market | Learning and Productivity | customer service |
| https://doi.org/10.1287/msom.2015.0568 | Prescriptive | Individual | Learning and Productivity | process & bottleneck |
| https://doi.org/10.1287/msom.2016.0587 | Prescriptive | Individual | Learning and Productivity | process & bottleneck |
| https://doi.org/10.1287/msom.2016.0597 | Descriptive | Team | Learning and Productivity | team knowledge |
| https://doi.org/10.1287/msom.2016.0612 | Prescriptive | Individual | Behavioral biases | product pricing |
| https://doi.org/10.1287/msom.2016.0613 | Prescriptive | Market | Strategic behavior | product pricing |
| https://doi.org/10.1287/msom.2016.0615 | Descriptive | Organization | Learning and Productivity | customer service |
| https://doi.org/10.1287/msom.2017.0620 | Descriptive | Individual | Behavioral biases | newsvendor ordering |
| https://doi.org/10.1287/msom.2017.0634 | Prescriptive | Market | Strategic behavior | product pricing |

| | | | | |
|---|--------------|--------------|---------------------------|----------------------|
| https://doi.org/10.1287/msom.2017.0642 | Prescriptive | Market | Utility theory | customer service |
| https://doi.org/10.1287/msom.2017.0651 | Prescriptive | Market | Strategic behavior | product pricing |
| https://doi.org/10.1287/msom.2017.0654 | Descriptive | Individual | Behavioral biases | newsvendor ordering |
| https://doi.org/10.1287/msom.2017.0661 | Descriptive | Individual | Learning and Productivity | newsvendor ordering |
| https://doi.org/10.1287/msom.2017.0673 | Prescriptive | Market | Utility theory | product pricing |
| https://doi.org/10.1287/msom.2017.0682 | Prescriptive | Organization | Utility theory | process & bottleneck |
| https://doi.org/10.1287/msom.2017.0685 | Descriptive | Individual | Behavioral biases | product pricing |
| https://doi.org/10.1287/msom.2017.0688 | Prescriptive | Individual | Behavioral biases | product pricing |
| https://doi.org/10.1287/msom.2017.0701 | Descriptive | Individual | Behavioral biases | newsvendor ordering |
| https://doi.org/10.1287/msom.2017.0704 | Descriptive | Team | Learning and Productivity | team knowledge |
| https://doi.org/10.1287/msom.2018.0707 | Prescriptive | Market | Utility theory | customer service |
| https://doi.org/10.1287/msom.2018.0708 | Descriptive | Individual | Behavioral biases | newsvendor ordering |
| https://doi.org/10.1287/msom.2018.0719 | Descriptive | Organization | Behavioral biases | healthcare |
| https://doi.org/10.1287/msom.2018.0721 | Prescriptive | Market | Behavioral biases | process & bottleneck |
| https://doi.org/10.1287/msom.2018.0724 | Descriptive | Market | Behavioral biases | healthcare |
| https://doi.org/10.1287/msom.2018.0728 | Descriptive | Market | Behavioral biases | customer service |
| https://doi.org/10.1287/msom.3.1.25.9996 | Descriptive | Team | Learning and Productivity | team knowledge |
| https://doi.org/10.1287/msom.4.3.208.7753 | Prescriptive | Market | Utility theory | customer service |
| https://doi.org/10.1287/msom.4.3.228.7755 | Prescriptive | Market | Utility theory | product pricing |
| https://doi.org/10.1287/msom.5.3.179.16032 | Descriptive | Individual | Behavioral biases | process & bottleneck |

Appendix B: Hierarchical Clustering Algorithm Based on the Classification Scheme of Section

3.1

We associated each paper with a vector of dimension 11, each entry representing a factor in categorization scheme outlined in Section 3.1: 2 for research paradigm, 4 for locus of analysis, and 5 for theme. For instance, a prescriptive paper will have a 1 in the “prescriptive” entry and a 0 in the “descriptive” entry. Each vector has thus three entries equal to 1 and eight zero entries. We then ran a hierarchical clustering algorithm on those vectors using Ward’s minimum variance method. After analyzing the silhouette graphs, we opted for six clusters as it leads to a high average silhouette score while providing enough density in each cluster. Analysis of the dendrogram revealed that the first branching occurred on the dimension of the method.

Appendix C: Keyword-Based Nonnegative Matrix Factorization (NMF) Clustering Analysis

We defined a “document” to be a combination of title and abstract. We cleaned the documents (i.e., removing punctuation and stop words) to keep the most relevant words and lemmatized them (e.g., so that “price” and “pricing” are considered as one keyword); we hereon refer to the remaining stem words as keywords. We only considered keywords that appeared in at least two documents but no more than 99% of the documents.

For each document, we associated a numerical score with each keyword based on its relative frequency, namely the TF-IDF (term frequency–inverse document frequency) score. According to this scheme, a keyword has a high score if it appears frequently within a few abstracts. We then represented each document as a vector of numerical scores associated with all keywords; that is, if a document has many occurrences of “pricing”, but the others do not, its vector would have a high score associated with the “pricing” keyword. The pool of documents can thus be represented as a matrix V , of dimension $n \times m$, in which n is the number of documents under consideration (here, $n = 98$) and m is the number of keywords ($m = 926$).


The NMF method consists in approximating V as a product of nonnegative matrices of smaller dimension, i.e., to find $W \geq 0$, of dimension $n \times k$, and $H \geq 0$, of dimension $k \times m$, with $k \ll n, m$, which minimize $\|V - WH\|_F^2 + 0.1 \|W\|_F^2 + 0.1 \|H\|_F^2$; here, the last two terms are regularization terms which prevent matrices W and H from growing too large. Matrix H is effectively a features matrix, representing each of the k features as a vector of scores on each of the m keywords. Matrix W is a coefficients matrix, representing each of the m documents as a linear combination of the k features.

We ran the NMF algorithm for $k = 5, \dots, 10$. For each number of clusters k , document i , for any $i = 1, \dots, n$, is selected to belong to cluster j , for $j = 1, \dots, k$, if $W_{ij} = \arg \max_{l=1, \dots, k} W_{il}$. Following the method outlined in Kuang et al. (2015), we computed the dispersion coefficient at each iteration, which measures the stability of the clustering algorithm through random

sampling of documents. In addition, we computed the silhouette scores, which is a measure of the quality of the classification scheme.

We decided to set $k = 6$ given that this iteration was associated with both a local maximum of the coefficient of dispersion and that it had the highest average silhouette score. Analysis of the silhouette graph displayed in Figure 3 reveals that the silhouette plots have similar thickness and that 5 out of 6 had scores that peaked above the average silhouette score and had few points with negative scores. Although considering $k = 5$ would have resulted in peaks above the average silhouette score for all five clusters, we felt there was something distinctive about the additional cluster (Cluster "Process and bottleneck" on the figure) and settled for $k = 6$.

Figure 3: Silhouette of the 6 Clusters Obtained with the NMF Algorithm



In Table 2, for each cluster $j = 1, \dots, k$, the scores associated with the keywords correspond to the largest entries of the row $[H_{ji}]_{i=1, \dots, m}$ and the most representative paper is identified as the document that has the highest coefficient W_{ji} over all $i = 1, \dots, n$.