

Preface

Revenue management (RM) has gained attention recently as one of the most successful application areas of operations research (OR). The practice has grown from its origins as a relatively obscure practice among a handful of major airlines in the post-deregulation era in the U.S. (circa 1978) to its status today as a mainstream business practice with a growing list of industry users from Walt Disney Resorts to National Car Rental and a supporting industry of software and consulting firms. Major airlines, hotel chains, and car rental companies have large staffs of developers and analysts working on RM, and major consulting and software firms also employ large numbers of RM professionals.

There are now several major industry RM conferences each year: The Airline Group of the International Federation of Operational Research Societies (AGIFORS) sponsors an annual reservation and yield management conference that has attracted up to 200 professionals, and The International Air Travel Association (IATA) hosts an annual RM conference that has drawn up to 800 attendees in recent years. The Professional Pricing Society also hosts professional conferences that address science-based pricing methods and technologies and general pricing strategy.

Over this same period, academic and industry research on the methodology of RM has also grown rapidly. The number of published papers on RM has increased dramatically in the last ten years. INFORMS, the leading professional society of OR, has started a Pricing and RM Section, which has now hosted several annual conferences on RM, each drawing in excess of 100 researchers and professionals. And several universities now offer specialized RM courses, at both the M.B.A and Ph.D. levels.

Despite this explosion of both professional and scholarly activity, no book has comprehensively covered the field of RM. For any area in such a mature state of development and with such widespread industry usage,

such a reference is desirable. However, for RM the need is particularly acute for several reasons:

- RM is very much a professional practice and as such there is a considerable amount of “institutional” knowledge surrounding it that is relatively inaccessible to those outside the profession.
- Many of the early and even some more recent seminal ideas do not appear in published journals. Even those that have been published sometimes appear in relatively obscure sources such as AGIFORS proceedings, industry newsletters, and standard industry practice.
- The terminology, concepts, and notation have not been standardized to date, so it is often confusing for an outsider to reconcile the various contributions of the extant literature.
- There is often a considerable gap between practitioners and academics in the field. Academics are often not aware of the real world complexities faced by practitioners of RM, and industry practitioners are often not aware of the more recent advances in the academic literature.

Our aim in writing this book is to meet this need. The book seeks—as its title indicates—to cover both the *theory* and the *practice* of RM. Fundamentally, RM is an applied discipline; its value and significance ultimately derive from the business results it achieves. At the same time it has strong elements of an applied science, and the technical elements of the subject deserve rigorous treatment. Both these practical and theoretical elements of the field reinforce each other, and to a large extent this is what makes the topic exciting. It is this constructive interplay of theory and practice that we have strived to capture in this book.

Audience

We have two primary audiences in mind for this book—(1) analytically trained (or at least “analytically tolerant”) practitioners in industry and (2) academic researchers and teachers. We view our core reader as someone who has the equivalent of a master’s degree or higher in a technical subject such as engineering, operations research, statistics, or economics. However, significant portions of the text are accessible to general or business readers, particularly the introduction, Chapter 10 on industry profiles, and Chapter 11 on implementation issues. In addition, the introductions to the technical chapters provide high-level overviews of

each chapter, which are designed to provide a qualitative understanding of the main topics covered and their business context, and give the reader a sense of the essence—if not the details—of the material.

For experienced practitioners this book serves as a single-source reference for the major theory and application issues involved in RM. The key technical results in the field are organized and presented precisely and in consistent notation, so that practitioners can easily refer to relevant models, formulas, and algorithms as needed. For new employees in the RM industry our book also serves as a useful primer on the subject, allowing them to “get up to speed” on the details of the field quickly through a consistent presentation of the material. For the technically oriented user it serves as an unbiased, noncommercial source for understanding the competing methodologies available for RM and their relative strengths and weaknesses.

We view the academic audience for the book as consisting of the many researchers now working on various RM-related topics, as well as those who work in related areas (such as supply-chain management), who may want a single-source, accessible overview of the main theory and practice components of the field. Academics who teach management science or operations management courses may also find the book useful, either directly as a supplementary text or simply for the instructor’s personal use as a reference on the subject. Our experiences with colleagues outside the field has suggested that most are curious about RM but perhaps not confident enough about the theory and practice to introduce the subject in their classes. This book should help “demystify” the subject for them.

Finally, a growing number of courses have specifically focused on RM. Though not designed particularly as a textbook, the book should serve as a useful reading and reference in such courses. While we have not put in homework exercises, we did include many small, technically uncluttered examples throughout the book that illustrate the core concepts being discussed.

We forewarn the reader that the material in some places in the book has an airline bias. This is as it should be in our opinion; airline RM practice remains an important topic in its own right. In addition, a large number—indeed the vast majority—of RM practitioners and researchers working in the field today are involved directly in airline RM practices. So airline RM is deserving of rigorous and careful coverage, which is one of our goals in writing this book.

At the same time, not every industry is like the airline industry and “airlinelike” conditions are not, in our view, that necessary to apply RM ideas. Therefore, we have attempted to present RM in as generic terms as possible and included several topics and chapters that generalize

beyond the airline industry. We have tried to be somewhat forward looking in this regard, while at the same time not venturing too far into the realm of pure speculation.

Content and Style

As for the choices of material, we have aimed for an applied technical (engineering) level in our treatment of the subject. For example, we have chosen to present all problems in discrete time. This eliminates several technical complications, while still allowing us to address a wide range of problems in a simple, yet rigorous way. Moreover, continuous-time models and methods are not frequently used in practice, so the focus on discrete-time methods is well justified from a practical standpoint.

Similarly, we have not included a large number of proofs. This is both consistent with the applied orientation of the field and reflects our view that RM models and theory do not share enough in common to justify a highly formalistic, deductive approach to the subject. In a few cases we provide proofs of the theoretical results, but even these are relegated to appendices. When proofs are omitted, we provide references to the original sources and if possible give either informal arguments or intuition about the results.

In addition, the bodies of each chapter do not contain a large number of literature references. This is because we want the reader to “see the material for what it is” and not be sidetracked by a lot of discussion of the literature. Where ideas are strongly associated with specific papers and people, we, of course, point this out. Detailed references to the literature and a discussion of sources are collected in a Notes and Sources section provided at the end of each chapter. To further assist the reader, appendices containing basic results on probability theory, continuous optimization, dynamic programming, and game theory are provided to make the technical material in the book as self-contained as possible.

We tried to be comprehensive in our coverage of RM, covering both quantity- and price-based RM as well as the supporting topics of forecasting and economics. While we might have risked over-extending ourselves in this regard, we believe such a comprehensive approach is necessary to fully understand the subject. Indeed, a key contribution of the book is to unify the various forms of RM and to link them closely to each other and to the supporting fields of statistics and economics. The topics and coverage do, however, reflect our own personal choices about what is and is not important to understand RM. While we have tried to be as comprehensive, fair, and balanced as possible in arriving at these choices, undoubtedly our choices have resulted in some biases. However,

the benefit to the reader is that the text has a point of view and is not merely an uncritical inventory of all research results to date in the field.

Finally, we have also tried to come up with a notation that is generic and consistent across all the chapters. Much of this notation will not coincide with the notation found in the original papers in the field, which is by and large quite inconsistent anyway. A summary of notation is provided in Appendix A for reference. The consistency of notation and presentation, we believe, makes reading the book much easier than looking at the corresponding collection of original-source articles, and it also highlights the connections among topics.