Revenue Management and the Rise of the Algorithmic Economy*

Ilan Lobel†

May 18, 2020

Abstract

Revenue management has evolved over the years from its origins in the airline industry into a much broader discipline that analyzes algorithmic methods for demand and marketplace management, but many outside of the discipline are not aware of this transformation. The field's transition tracks the widespread adoption of algorithmic decision-making techniques by businesses in a wide variety of industries over the last decade. We study this evolution in the field's breadth of research, with a particular focus on revenue management papers that study online marketplaces such as e-commerce retailing, digital advertisement, and ride-hailing markets for urban transportation.

1 Introduction

The term "revenue management" (RM) has a very different meaning in academic circles than it does in colloquial speech. While the everyday use of the term is still broadly associated with airline pricing, RM is now a much bigger field of research, and one that encompasses

*This work has benefited from feedback provided by Omar Besbes, Srikanth Jagabathula, Huseyin Topaloglu, seminar participants at NYU Stern and attendees of the Management Science 65^{th} Anniversary Conference. I thank them for their valuable feedback.

†Stern School Business, New York University – ilobel@stern.nyu.edu

a far broader set of problems. For most of its history, RM was a discipline centered on a set of questions arising from a specific pair of industries, travel and hospitality, but this is no longer the case. The tremendous growth of the technology industry over the past decade has created new opportunities for RM research and practice, and RM has grown to encompass a variety of new problems related to demand management and, more broadly, market analytics. This paper documents the evolution of research in RM, with a special focus on research developments from the past decade. We highlight the transition that occurred during the 2010s as the RM field began to serve as a broader meeting point for research on algorithmic methods for demand and marketplace management.

The transformation of the discipline tracks closely the changes that occurred in the broader economy over the past decade. Over this period, vast sectors of the economy underwent a process of algorithmization, with decision-making power increasingly being delegated from humans to algorithms. The rise of the algorithmic economy has affected many different industries, and perhaps no domain exemplifies this transformation better than the advertising industry. Over the last 10–15 years, advertising was transformed from an "analog" business centered on print and television into a business in which impressions and clicks are the primary considerations. A key element of this transformation was that advertising became programatic. That is, instead of ads being sold in advance to publishers through negotiated deals, ad allocation and pricing decisions are now often delegated to algorithms that make decisions in real time.

As a whole, the 2010s has been a decade in which the technology sector has extended its reach into many different areas of business and transformed them by implementing algorithmic decision-making techniques. Urban transportation is certainly one such example, with the emergence and growth of ride-hailing companies enabled by real-time matching and pricing algorithms. Retail itself has been deeply affected by the rise of e-commerce and its accompanying algorithmization of retail. Large online retailers often do not pick fixed prices for their products, as brick-and-mortar merchants do. Instead, they design algorithms to decide how prices should respond in real time to evolving market conditions and to fluctuating inventory levels. Online retailers are also able to personalize customers' experiences by, for example, showing different sets of products to different customers based on their predicted preferences. The growing economic importance of business models in which the interaction with consumers is mediated through real-time decision-making algorithms has opened up a large new research frontier, and RM scholars have moved into this new space and transformed the field of RM in the process. The growing importance of algorithms for business decision-making has created a need for research on how to design such algorithms, as well as research on the possibilities they enable and their inherent limitations. In this process, RM emerged as the intellectual home within management science for researchers and practitioners thinking about algorithms that are used for consumer-facing business decisions.

This aim of this paper is to explore the transformation of RM over the decades, from its origins in the airline industry to the broad field it is today. It is meant to both help outsiders to the field understand what the term "revenue management" means to RM academics today and how did this meaning come to be, as well as offer some guidance on the state of the literature to early-career RM researchers. The remainder of the paper is organized as follows. In Section 2, we briefly cover the history of RM, from its origins until the mid-2000s. In Section 3, we discuss more recent work, with a focus on RM research applied to three prominent online marketplaces: online retailing, digital ad markets, and ride hailing. We conclude in Section 4 with remarks on the state of the discipline and brief thoughts on important problems for the field to tackle in the near future.

Since this paper discusses the evolution of RM research over a long span of time, the list of relevant papers is obviously far longer than what its bibliography can accommodate. We apologize to the authors whose work is relevant but not cited here. This paper is not meant as a full literature review (for a fairly comprehensive bibliography of the field, we refer the reader to the superb recent book by Gallego and Topaloglu [2019]). Readers should be aware that the choice of topics to highlight is influenced by the author's area of expertise and taste. RM is a broad field of research and there are many RM papers being written on topics beyond the ones explored in this article.

2 The History of Revenue Management

2.1 The Early Days: 1970s to Early 1990s

RM started from the need to solve a fairly straightforward business question. In the early 1970s, airline travel was a luxury product. In the United States, routes, schedules, and fares were tightly controlled by the Civil Aeronautics Board, and airlines had to petition the government if they wanted to make any changes. Deregulation was being discussed, and would eventually happen with the signing of 1978 Airline Deregulation Act. In particular, airlines had a profitable business model that consisted of selling relatively expensive tickets mainly to business customers. Suppose that airlines were permitted to sell cheaper tickets as well. Could an airline somehow increase its market size by selling cheap tickets to leisure customers while, at the same time, still selling the more expensive tickets to its traditional business customers?

Littlewood [1972] was the first paper to propose a framework for analyzing this business problem. Littlewood started from the following assumption: it is possible to discriminate between business and leisure customers because these two categories of passengers have different time preferences. Oftentimes, business customers become aware of their need to travel at the last minute, whereas price-sensitive leisure customers can plan ahead. This postulated time separation between the two customer classes is obviously not a perfect description of reality, but it is potentially a reasonable approximation. If we take time separation as a given, airlines can charge low prices well in advance of the flight departure to leisure customers and charge higher prices to business customers closer to the departure date.

However, there is a potential flaw in this business strategy. Airplanes have fixed capacities, and it is possible for an airline to sell too many seats to leisure customers and thus have insufficient seats by the time business customers arrive. Clearly, airlines need some kind of "capacity protection" to ensure that a sufficient number of seats are reserved for higher-paying business customers. Suppose the prices offered to leisure (early) and business (late) customers are p_L and p_H, respectively, with $p_L < p_H$. Let us assume that the total demand numbers from leisure and business customers are independent random variables X_L and X_H. If there are x seats remaining and leisure customers are still arriving, what should a revenue-maximizing airline do?

The correct answer is that the airline should accept a request from a leisure customer (that is, while the price is p_L) if, and only if, p_L is greater than the expected revenue from reserving that seat for a future potential business customer, i.e., $p_L > p_H \cdot \mathbb{P}(X_H \geq x)$. This answer corresponds to the solution of the newsvendor model from inventory theory (Arrow et al. [1951]), with the optimal inventory to be reserved for business customers being determined by the distribution of X_H, the business ticket price being p_H, and cost of reserving such a unit for the business market being p_L.

This simple model has been quite influential in the development of RM. By connecting RM to a classical inventory model, Littlewood [1972] served as a de facto invitation for operations management scholars to study RM. The most widely adopted airline RM models have been quantity-based approaches that are based on generalizations of Littlewood's framework. Perhaps the most famous of such generalizations are the expected marginal seat revenue EMSR-a and EMSR-b models (Belobaba [1992]), which propose mechanisms for extending Littlewood's framework from a two-class model to a general multiclass model.

2.2 The Growth of the Discipline: Mid-1990s to Mid-2000s

The mid-1990s saw RM grow into a fully formed area of research. One important development from this era was its expansion into retailing. Growing into retailing was not a matter of applying old techniques to new domains. In fact, it required a wholesale rethinking of the by-then standard (quantity-based) framework of RM and the creation of a new framework: price-based RM.

The earliest approaches to RM were quantity-based in the sense that they took prices for different product classes as given and then tried to optimize quantity allocation decisions, with Littlewood's model being a classic example of a quantity-based RM framework. However, this approach makes less sense in retailing than in the airline world. Retail is an area that shares some similarities with the basic airline RM framework, but one that also exhibits some important differences. Consider a fashion retailer that has ordered a given inventory of a certain product from a manufacturer. The retailer has a finite amount of time to sell that inventory before the season is over. The manufacturer might also require a significant lead time, and therefore, replenishing inventory might not be possible within a selling season. Under these conditions, the problem of managing fashion inventory is not substantially different from the problem of managing an inventory of airplane seats. However, opening and closing fare classes, which is the standard approach to airline RM, is far less natural in this setting. A more natural approach for retailers is simply to vary prices over time.

There is an even more fundamental reason why Littlewood's approach is not applicable to retailing: the postulated time separation of customers does not hold in this setting. If at all present, it exists in reverse, since fashion items are worth more earlier in the season than later on. Therefore, the classic reason for protecting inventory does not apply to retail settings. However, this does not mean that there is no need for RM in retail. The question of how to manage a finite inventory with a finite selling horizon is actually quite important in retail, but the primary concern is a little different in this context. Here, the main focus is on how to balance the risks of depleting the inventory too fast versus the risk of ending up with leftover stock at the end of the selling season. The need to address this trade-off is the origin of price-based RM.

To address this trade-off, Gallego and van Ryzin [1994] proposed what is perhaps the most influential model in RM, which we will refer to as the GvR model. In GvR, a firm owns an initial quantity of inventory and has a finite horizon in which to sell it. There is a stationary demand intensity function, and the higher the price chosen by the firm at a given time, the lower the Poisson arrival rate of purchases will be at that time. Gallego and van Ryzin [1994] used the Hamilton-Jacobi-Bellman equation to derive structural properties of optimal policies and to compute the optimal policy explicitly for specific demand intensity functions. They found that optimal policies take the form of prices that decrease continuously over time, except that they jump upwards when purchases occur. They also proved the asymptotic optimality of a heuristic that is based on simply ignoring the stochasticity of the problem.

Since its publication, the GvR model has served as a fundamental building block of RM, and many papers have been written to extend the GvR model and address its limitations, some of which we discuss in Section 3. An important extension was the one proposed by the same authors in Gallego and van Ryzin [1997], which they showed that their treatment of single-product RM could be extended to network RM. Network RM is the problem of selling products that are composed of multiple finite resources. Airline RM is a network problem since products sold are typically return trips (potentially including layovers) that are composed of multiple flight legs. Gallego and van Ryzin [1997] showed that a deterministic approximation of the standard network RM problem could be solved via linear programming and that this heuristic generates policies that are asymptotically optimal for the stochastic version of the problem. It is an interesting historical twist that this important technique for airline RM emerged not from the quantity-based RM world where airlines were center-of-mind, but from a modification of the GvR framework that was originally designed to capture the RM problem faced by a fashion retailer.

The study of network RM eventually led to the development of choice-based RM. The goal of choice-based RM is to bring into the model the fact that consumers typically do not face a binary buy/decline option but instead have several different purchase options available to them at any given time. Choice-based RM involves the integration of discrete choice (Ben-Akiva and Lerman [1985]) into RM models. The first paper in this literature was Talluri and van Ryzin [2004a], which considered a model in which there are multiple fare classes available to consumers but there is a single resource available that gets depleted over time. A key novelty here is that the seller's choice at each point in time is now a product set to offer instead of a price. Two influential papers in this area were Gallego et al. [2004] and Liu and van Ryzin [2008], which respectively introduced and analyzed a linear programming approach to choice-based RM. As we will see in Section 3, choice-based RM would eventually give birth to the rich world of assortment optimization.

Overall, by the early 2000s, RM was a successful academic discipline, having developed both rigorous theoretical frameworks and methods widely adopted by industry. So far, we have focused on travel and retail applications of RM, but there were other application areas that saw signficant research in dynamic pricing during this era, including queuing (Paschalidis and Tsitsiklis [2000]) and inventory management (Federgruen and Heching [1999] and Chen and Simchi-Levi [2004]). For a detailed discussion of the advances in RM up to the mid-2000s, we refer the reader to the seminal book by Talluri and van Ryzin [2004b] and to the surveys by McGill and Van Ryzin [1999] and Elmaghraby and Keskinocak [2003].

3 Revenue Management for Online Markets

Over the last decade, online markets have grown tremendously, digitizing a number of different lines of business along the way. This has created a significant business need for algorithms for real-time demand management of the kind that RM researchers have been developing for decades. In this section, we discuss three application areas that have been central to RM research over the past decade: online retailing, digital advertising and ride hailing. RM for online retail is in some ways a natural progression from the earlier research on RM for brick-and-mortar retail, though with a few important differences. The latter two application areas constitute a greater departure for RM, and the tremendous growth of RM applied to these topics would have been difficult to predict a decade ago.

3.1 Online Shopping and the Growing Importance of Retail RM

Retailing has been an important driver of RM research ever since the introduction of the GvR model in 1994. However, implementing RM algorithms has typically been quite a bit more challenging in retail settings than in the world of airlines. Putting RM algorithms in place generally requires the ability to modify prices frequently and quickly as well as maintaining an accurate real-time accounting of inventory. Both of these requirements are fairly difficult to satisfy for brick-and-mortar retailers. The rise of online shopping over the last decade has dramatically changed the picture for retail RM. Implementing an algorithm that changes a product's price in response to demand and inventory level is not nearly as difficult a task online as it is in brick-and-mortar. The rapid growth of online commerce over the last decade has enabled retail RM to take center stage within the discipline of RM.

The increasing importance of retail RM has led many authors to try to shore up some of the weak spots of GvR as a model of retail RM. An important assumption in GvR is that the selling firm knows the demand intensity function. That is, the firm knows the expected demand response from a given price. This assumption may be justifiable in an airline setting, where firms sells flights that are almost identical to each other day after day, but it is a less plausible assumption in fashion retail, where new products are likely to generate hard-to-predict demand responses. To address this question, many authors have proposed algorithms for dynamic pricing while learning the demand function. The literature includes both parametric approaches (Araman and Caldentey [2009], Broder and Rusmevichientong [2012], Harrison et al. [2012], Chen and Farias [2013], den Boer and Zwart [2013] and Besbes and Zeevi [2015]) and nonparametric ones (e.g., Besbes and Zeevi [2009] and Keskin and Zeevi [2014]). For a survey of the field of dynamic pricing with demand learning, we refer the reader to the survey by Araman and Caldentey [2011].

Another potentially problematic assumption in GvR is that of myopic consumers, that is, consumers who respond only to present prices and not to future potential price changes. This assumption has more validity in brick-and-mortar settings than in an online shopping context, where price tracking is relatively cheap and easy. Su [2007] and Aviv and Pazgal [2008] introduced this problem to the RM community, with the latter paper arguing that in certain situations strategic consumer behavior makes committing to a price a better approach than using a dynamic pricing algorithm. Over the last decade, several different approaches to this problem have been explored in the literature, ranging from papers that constrain the pricing policies in the GvR model to ensure that delaying purchases is not a winning strategy (Chen and Farias [2018] and Chen et al. [2019]) to papers that offer new models designed to analyze strategic consumer behavior (Besbes and Lobel [2015]).

Another problem that became central to RM research over the last decade is assortment optimization, which is the problem of selecting which set of products to offer consumers in order to maximize expected profits. Although this line of work emerged from choice-based RM, it has its own distinct flavor: assortment optimization relaxes inventory considerations and focuses mainly on the question of cross-product substitution when some products are unavailable. Assortment optimization is a problem that emerged from brick-and-mortar retail but that gained its current prominence due to the growth of e-commerce retailing and the immense variety of products that can be offered by retailers online. Two of the most influential papers in this area are Farias et al. [2013] and Blanchet et al. [2016]. The former proposed a model in which the consumer population is described by a distribution over rankings of products, and offered a method for estimating a choice model by assuming sparsity (where sparsity means that a model with only a few rankings is a reasonable approximation of the true distribution of rankings). The latter paper goes in a very different direction, showing how to model the consumer choice process via a simple Markov chain. The literature on this problem is vast and growing, and it includes papers such as Caro and Gallien [2007], Rusmevichientong et al. [2010], Sauré and Zeevi [2013], Davis et al. [2014], and Agrawal et al. [2019]. For more details, we refer the reader to a recent survey of the choice-based RM and assortment optimization by Strauss et al. [2018].

The problems discussed thus far are retail RM problems that became more important in the era of online shopping. There are also problems that did not exist prior to online retail that are now quite important RM questions. Foremost among them is the topic of contextualization/personalization, which is a topic at the intersection of RM and machine learning. In brick-and-mortar settings, everyone who visits a store at a given time sees exactly the same assortment of products. In online environments, stores have the ability to personalize assortments. For each individual, a different assortment can be chosen based on what is known about that individual and that particular visit to the website. Personalized assortment optimization was first studied in the literature by Golrezaei et al. [2014] and has recently picked up momentum with papers with papers such as Cheung and Simchi-Levi

[2017] and Chen et al. [2018].

A related problem that is also in the frontier between machine learning and RM is feature-based (or contextual) pricing. There are many online platforms where the number of products listed is immense and learning the value of each product individually is impractical, if not impossible. In such scenarios, the only feasible approach to demand learning is to somehow leverage the knowledge gained from offering a given product at a given price in order to try to also learn about the value of products that are distinct but have some commonality. The easiest way to do so is by learning the values of different product features. For example, consider the problem faced by a firm selling house stays, such as Airbnb or Booking.com. The firm could learn about the value of having an extra bedroom in a given neighborhood based on one listing and learn about the value of having a swimming pool based on a different listing. Contextual pricing algorithms combine these pieces of information to generate value predictions for new untested listings. Papers on this topic include Cohen et al. [2020], Javanmard and Nazerzadeh [2019], Qiang and Bayati [2016], Ban and Keskin [2018], and Lobel et al. [2018]. As a whole, the literature on contextual approaches to RM is young and growing fast. It builds closely on ideas from machine learning, such as contextual bandits (Auer et al. [2002] and Auer [2003]), and is likely to see further growth in the next few years.

This growth in research at the interface of RM and machine learning has been deeply intertwined with the growing interest in data-driven research. As Simchi-Levi [2014] argues, classical RM papers were typically problem-driven and centered on building models. However, there is growing interest in research where data is the centerpiece of the story. A good example is Ferreira et al. [2016], where the authors apply machine learning and optimization ideas to data from an online retailer to improve pricing decisions. Other papers in the same vein include Cheung et al. [2017] and Cohen et al. [2017]. For a recent survey on data-driven revenue management, we refer the reader to Mišić and Perakis [2020].

A further area of research that has been growing in the last few years as a consequence of the boom in online shopping is the study of the interplay between RM and supply chain management. This interplay is important, for example, in omnichannel retailing, where online prices affect not only online demand but also offline demand, and fulfillment decisions affect product availability across the retail network (Harsha et al. [2019]). RM ideas can also be used to manage e-fulfillment demand by shifting users to less crowded delivery slots via pricing (Yang et al. [2016]).

3.2 The Market for Digital Ads

The growth of RM into retail is perhaps unsurprising given the history of the field. The next two applications areas — online advertising and ride hailing — are not domains that the founders of RM could have foreseen, as these are industries that did not exist in their current form two decades ago (in the case of ride hailing, it did not exist even a decade ago).

Online advertising is a large and growing industry. It generates the majority of the revenues of Alphabet and Facebook, two of the largest companies in the world. Online ads look nothing like airplane seats and, at first sight, it is not clear why techniques devised for filling airplanes could have any use in the domain of digital advertising. However, there are a few key similarities. In online advertising, a publisher starts off with a batch of impressions to sell, where an impression is the placement of an ad on a webpage for a single user visit. Therefore, the notion of an exogenous inventory — a key element in airline RM — is also present here, though the inventory available has an extra element of stochasticity that is not present in the airline context (a web publisher does not know the exact number of visitors/impressions it will get over the next week or month, but only has an estimate). Furthermore, impressions are usually sold to advertisers via real-time algorithms, with allocation and pricing decisions made in the milliseconds between a person clicking on a link and the loading of the new webpage. As a whole, online advertising turned out to be an industry ready for RM ideas.

A good introduction to this literature is Balseiro et al. [2014], a paper that explored the connections between these two industries. It considered the case of a web publisher who has several contracts with advertisers, where the contracts specify quantities of impressions that need to be delivered over a given time horizon. The publisher wants to maximize the quality of the impressions delivered to the advertisers, where quality can be viewed as the probability a user will click on an ad. The publisher has a second objective, which is revenue. Beyond being delivered to the contracted advertisers, impressions can also be directed to an ad exchange where they are sold to real-time bidders. For an impression sent to the ad exchange, the publisher chooses a reserve price. The authors argued that this problem is closely related to parallel-fight network RM and that the techniques that emerge from this literature, such as bid prices (Talluri and van Ryzin [1998]), can therefore also be used here.

There are a few important distinctions between digital advertising RM and classical RM. In standard RM, the seller chooses prices (or fare classes to open). In digital advertising, the platform selling the impression (often an ad exchange) chooses an auction format. That is, it chooses the rules that determine how advertisers bid on an impression and how the impression will be allocated and priced as a function of the bids. This is a larger and more complex design space than simply choosing prices. Two popular choices in this industry are first-price and second-price auctions, often augmented by a reserve price (a minimum price for an allocation to occur). The choice of such formats is partially supported by the celebrated result of Myerson [1981], which demonstrates the optimality of these auction formats when bidders are symmetric in their valuation distributions. RM researchers have been working on how to learn optimal reserve prices when valuation distributions are not known (Kanoria and Nazerzadeh [2017] and Golrezaei et al. [2018]), a problem at the intersection of RM and machine learning. Another potential avenue for improvement would be to implement Myerson's optimal auction for settings with asymmetric bidders. However, Myerson's auction for asymmetric bidders is quite complex and challenging to implement as it requires estimating valuation densities and is therefore rarely used in practice. To tackle this problem, RM researchers have proposed a few different techniques for approximating an optimal auction, including randomization (Celis et al. [2014]) and linearization (Golrezaei et al. [2017]).

Another key distinction between classical RM and RM for online advertising is how often a customer interacts with the seller. In classical RM, the standard assumption is that every transaction involves a new customer. This assumption, however, is quite unrealistic in online advertising, where advertisers arrive to the platform desiring to buy thousands of impressions, not single ones. This allows for new phenomena to emerge. One such phenomenon is learning dynamics: buyers can learn about their valuations for impressions over time as well as learn about the strategies of other bidders in the market (Iyer et al. [2014] and Balseiro and Gur [2019]). Another important practical issue is financial constraints, such as budgets. Buyers might depart the system too early if their budgets get depleted too fast, affecting the competitive landscape of the auctions (Balseiro et al. [2015] and Balseiro et al. [2019]). A further interesting feature of this marketplace is that sellers might be able to use a mixture of real-time pricing and long-term contracts, including selling contracts that might offer buyers reduced spot market prices in the future. Perhaps surprisingly, complex contracts like this have the potential to unlock significantly more revenue and welfare than it is possible to obtain under dynamic pricing alone (Kakade et al. [2013], Balseiro et al. [2017], and Mirrokni and Nazerzadeh [2017]).

Overall, to be able to contribute to the domain of online advertising, RM has had to evolve from a real-time demand management discipline into a broader real-time marketplace management discipline. It has had to merge classical RM ideas with mechanism design techniques from economics and computer science. For a broader survey of this topic, we refer the reader to Korula et al. [2015].

3.3 Ride Hailing and Urban Transportation

Another industry that has been utterly transformed by technology in the recent past is urban transportation, which has witnessed the rise of ride-hailing companies such as Lyft and Uber. RM is at the core of how this industry functions, with the ride-hailing companies using sophisticated dynamic pricing algorithms to manage their supply and demand. Dynamic (or surge) pricing is important for ride-hailing platforms, because it enables them to offer low prices during low demand periods, while still guaranteeing high availability during high demand periods (Cachon et al. [2017]).

Dynamic pricing is the tool of choice for inducing the movement of drivers from where they are to where they are most needed. Using a natural experiment, Lu et al. [2018] showed that drivers respond to surge pricing heat maps by relocating to places where they are in higher demand. This phenomenon is analogous to the strategic consumer effect in classical RM, except that the movement of drivers is more complex since they relocate in both space and time.

Several papers have proposed models for understanding how ride-hailing platforms should use surge pricing to move drivers in space and time. Bimpikis et al. [2019] proposed a long-term network pricing model aimed at elucidating the role of traffic patterns in optimal pricing policies, and found that "balanced" traffic networks lead to better equilibrium outcomes. Besbes et al. [2020] explores a short-term model, proposing a framework for understanding how to move drivers from where they are to where they are needed via surge pricing. Afèche et al. [2017] solved an intermediate-horizon problem via driver admission control, in a way that is reminiscent of quantity-based RM. Ma et al. [2018] posed the driver movement problem as an integer program, and showed that driver-pessimal payments implement the desired driver allocation. Garg and Nazerzadeh [2019] argued that ride-hailing platforms should move from multiplicative to additive price surges in order to better manage intertemporal driver incentives.

Another interesting and novel feature of this marketplace are pick-up times. Castillo et al. [2017] pointed out that under a first-dispatch policy (closest car to rider is dispatched), low prices can cause a new form of market failure, which the authors called the wild goose chase. A wild goose chase occurs when not enough empty drivers are available, leading to drivers being sent to pick up faraway customers, increasing the system workload. They then argued that the primary role of the surge pricing algorithm is to avoid this kind of market failure. In response, Besbes et al. [2019] showed that surge pricing is not necessary to stabilize a ride-hailing system: a change in dispatch policy is an alternative solution. Overall, pick-up times are an important and relatively understudied feature of ride-hailing systems.

There is also a growing literature on matching algorithms for ride-hailing platforms. Most of these studies do not include pricing or other classical RM controls (Hu and Zhou [2016], Feng et al. [2017], Ashlagi et al. [2018], Banerjee et al. [2018], and Ozkan and Ward [2019]), but this is beginning to change; some of the latest papers in this area include pricing as one of the platform controls (Kanoria and Qian [2019] and Ozkan [2019]). For more details on RM applied to ride hailing, we refer the reader to the recent survey by Korolko et al. [2019].

4 Concluding Remarks

Over the last decade, the economy has undergone a process of algorithmization, with activities that used to be "analog," such as buying ads or hailing cars, switching to being mediated by real-time algorithms. This has created an opportunity and a need for a science of real-time demand and marketplace management, and the RM community has seized this opportunity. With today's ubiquity of computing, this process is likely to accelerate, and RM researchers should be on the lookout for new industries to which they can contribute.

An important change in the business landscape over the last few years is the growing importance of machine learning. To remain relevant, the discipline of RM needs to both learn from and integrate with machine learning. Many RM problems are dynamic programming problems in nature, and there is a subfield of machine learning called reinforcement learning that deals with dynamic programs with large state-spaces. The fact that a sizable number of papers have been written in the last few years on data-driven RM, contextual pricing, personalized assortment optimization, and contextual learning of reserve prices (see Sections 3.1 and 3.2) bodes well for future interactions between RM and machine learning.

An area that is mostly unexplored and likely to be of increasing importance is the ethics and fairness of RM systems. The same contextual pricing algorithms that can be used to learn optimal prices for large numbers of products can easily be repurposed into personalized pricing algorithms that attempt to extract all consumer surplus, a direction that is likely to cause societal backlash. With a few exceptions (e.g., Cohen et al. [2019]), the field of RM has mostly not focused on the ethics and fairness of RM systems. In a similar note, privacy concerns are an increasingly important consideration in our modern economy, and many localities are implementing laws that restrict the storage and use of data, such as the European Union's General Data Protection Regulation (GDPR). Such laws have significant implications for RM systems, especially pricing ones. RM researchers could potentially learn from computer science work that studies algorithmic decision-making under privacy constraints, such as the differential privacy framework (Dwork [2008]).

An area where RM research has historically underperformed is the study of RM under competition. Most RM papers assume monopoly settings, but almost all applications are in oligopolistic or competitive settings. There exist some work in this area (e.g., Gallego and Hu [2014]), but this literature is small relative to its practical importance. The rise of digital markets makes this issue even more pronounced as such markets make price and availability comparisons significantly easier for consumers.

Overall, RM has been an area of study at the forefront of a transformation of the business landscape, and this is not likely to change in the near future. RM has a strong track record of integrating theory and practice, by both generating tools for industry and learning from the latest business problems. *Management Science* has played a key role in the growth and development of the discipline, and with the recent creation of a *Revenue Management & Market Analytics* department, the partnership between the journal and the discipline of RM is likely to continue growing in strength.

References

Philipp Afèche, Zhe Liu, and Costis Maglaras. Ride-Hailing Networks with Strategic Drivers: The Impact of Platform Control Capabilities on Performance. *Working Paper, University of Toronto*, 2017.

Shipra Agrawal, Vashist Avadhanula, Vineet Goyal, and Assaf Zeevi. MNL-bandit: A Dynamic Learning Approach to Assortment Selection. *Operations Research*, 67(5):1453–1485, 2019.

Victor F Araman and René Caldentey. Dynamic Pricing for Nonperishable Products with Demand Learning. *Operations Research*, 57(5):1169–1188, 2009.

Victor F Araman and René Caldentey. *Revenue Management with Incomplete Demand Information*. Wiley Encyclopedia of Operations Research and Management Science, 2011.

Kenneth J Arrow, Theodore Harris, and Jacob Marschak. Optimal Inventory Policy. *Econometrica*, 19(3):250–272, 1951.

Itai Ashlagi, Maximilien Burq, Chinmoy Dutta, Patrick Jaillet, Amin Saberi, and Chris Sholley. Maximum Weight Online Matching with Deadlines. *Working Paper, Stanford University*, 2018.

Peter Auer. Using Confidence Bounds for Exploitation-Exploration Trade-Offs. *The Journal of Machine Learning Research*, 3:397–422, 2003.

Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning*, 47(2-3):235–256, 2002.

Yossi Aviv and Amit Pazgal. Optimal Pricing of Seasonal Products in the Presence of Forward-looking Consumers. *Manufacturing & Service Operations Management*, 10(3): 339–359, 2008.

Santiago R Balseiro and Yonatan Gur. Learning in Repeated Auctions with Budgets: Regret Minimization and Equilibrium. *Management Science*, 65(9):3952–3968, 2019.

Santiago R Balseiro, Jon Feldman, Vahab Mirrokni, and Shan Muthukrishnan. Yield Optimization of Display Advertising with Ad Exchange. *Management Science*, 60(12):2886–2907, 2014.

Santiago R Balseiro, Omar Besbes, and Gabriel Y Weintraub. Repeated Auctions with Budgets in Ad Exchanges: Approximations and Design. *Management Science*, 61(4): 864–884, 2015.

Santiago R Balseiro, Vahab S Mirrokni, and Renato Paes Leme. Dynamic Mechanisms with Martingale Utilities. *Management Science*, 64(11):5062–5082, 2017.

Santiago R Balseiro, Omar Besbes, and Gabriel Y Weintraub. Dynamic Mechanism Design with Budget-Constrained Buyers Under Limited Commitment. *Operations Research*, 2019.

Gah-Yi Ban and Bora Keskin. Personalized Dynamic Pricing with Machine Learning. *Working Paper, London Business School*, 2018.

Siddhartha Banerjee, Yash Kanoria, and Pengyu Qian. State Dependent Control of Closed Queueing Networks with Application to Ride-hailing. *Working Paper, Cornell University*, 2018.

Peter P Belobaba. Optimal vs. Heuristic Methods for Nested Seat Allocation. In *Presentation at ORSA/TIMS Joint National Meeting*, 1992.

Moshe E Ben-Akiva and Steven R Lerman. *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press, 1985.

Omar Besbes and Ilan Lobel. Intertemporal Price Discrimination: Structure and Computation of Optimal Policies. *Management Science*, 61(1):92–110, 2015.

Omar Besbes and Assaf Zeevi. Dynamic Pricing Without Knowing the Demand Function: Risk Bounds and Near-optimal Algorithms. *Operations Research*, 57(6):1407–1420, 2009.

Omar Besbes and Assaf Zeevi. On the (Surprising) Sufficiency of Linear Models for Dynamic Pricing with Demand Learning. *Management Science*, 61(4):723–739, 2015.

Omar Besbes, Francisco Castro, and Ilan Lobel. Spatial Capacity Planning. *Working Paper, Columbia University*, 2019.

Omar Besbes, Francisco Castro, and Ilan Lobel. Surge Pricing and its Spatial Supply Response. *Management Science*, 2020.

Kostas Bimpikis, Ozan Candogan, and Daniela Saban. Spatial Pricing in Ride-Sharing Networks. *Operations Research*, 67(3):744–769, 2019.

Jose Blanchet, Guillermo Gallego, and Vineet Goyal. A Markov Chain Approximation to Choice Modeling. *Operations Research*, 64(4):886–905, 2016.

Josef Broder and Paat Rusmevichientong. Dynamic Pricing Under a General Parametric Choice Model. *Operations Research*, 60(4):965–980, 2012.

Gerard Cachon, Kaitlin M Daniels, and Ruben Lobel. The Role of Surge Pricing on a Service Platform with Self-scheduling Capacity. *M&SOM*, 2017.

Felipe Caro and Jérémie Gallien. Dynamic Assortment with Demand Learning for Seasonal Consumer Goods. *Management Science*, 53(2):276–292, 2007.

Juan Camilo Castillo, Dan Knoepfle, and Glen Weyl. Surge Pricing Solves the Wild Goose Chase. In *Proceedings of EC*, pages 241–242. ACM, 2017.

Elisa Celis, Gregory Lewis, Markus Mobius, and Hamid Nazerzadeh. Buy-it-now or Take-a-chance: Price Discrimination Through Randomized Auctions. *Management Science*, 60 (12):2927–2948, 2014.

Xi Chen, Yining Wang, and Yuan Zhou. Dynamic Assortment Optimization with Changing Contextual Information. *Working Paper, New York University*, 2018.

Xin Chen and David Simchi-Levi. Coordinating Inventory Control and Pricing Strategies with Random Demand and Fixed Ordering Cost: The Finite Horizon Case. *Operations Research*, 52(6):887–896, 2004.

Yiwei Chen and Vivek F Farias. Simple Policies for Dynamic Pricing with Imperfect Forecasts. *Operations Research*, 61(3):612–624, 2013.

Yiwei Chen and Vivek F Farias. Robust Dynamic Pricing with Strategic Customers. *Mathematics of Operations Research*, 43(4):1119–1142, 2018.

Yiwei Chen, Vivek F Farias, and Nikolaos Trichakis. On the Efficacy of Static Prices for Revenue Management in the Face of Strategic Customers. *Management Science*, 2019.

Wang Chi Cheung and David Simchi-Levi. Thompson Sampling for Online Personalized Assortment Optimization Problems with Multinomial Logit Choice Models. *Working Paper, MIT*, 2017.

Wang Chi Cheung, David Simchi-Levi, and He Wang. Dynamic Pricing and Demand Learning with Limited Price Experimentation. *Operations Research*, 65(6):1722–1731, 2017.

Maxime Cohen, Adam N Elmachtoub, and Xiao Lei. Pricing with Fairness. *Working Paper, McGill University*, 2019.

Maxime Cohen, Ilan Lobel, and Renato Paes Leme. Feature-Based Dynamic Pricing. *Management Science*, 2020.

Maxime C Cohen, Ngai-Hang Zachary Leung, Kiran Panchamgam, Georgia Perakis, and Anthony Smith. The Impact of Linear Optimization on Oromotion Planning. *Operations Research*, 65(2):446–468, 2017.

James M Davis, Guillermo Gallego, and Huseyin Topaloglu. Assortment Optimization Under Variants of the Nested Logit Model. *Operations Research*, 62(2):250–273, 2014.

Arnoud V den Boer and Bert Zwart. Simultaneously Learning and Optimizing Using Controlled Variance Pricing. *Management Science*, 60(3):770–783, 2013.

Cynthia Dwork. Differential Privacy: A Survey of Results. In *Proceedings of International Conference on Theory and Applications of Models of Computation*, pages 1–19. Springer, 2008.

Wedad Elmaghraby and Pınar Keskinocak. Dynamic Pricing in the Presence of Inventory Considerations: Research Overview, Current Practices, and Future Firections. *Management Science*, 49(10):1287–1309, 2003.

Vivek Farias, Srikanth Jagabathula, and Devavrat Shah. A Nonparametric Approach to Modeling Choice with Limited Data. *Management Science*, 59(2):305–322, 2013.

Awi Federgruen and Aliza Heching. Combined Pricing and Inventory Control Under Uncertainty. *Operations Research*, 47(3):454–475, 1999.

Guiyun Feng, Guangwen Kong, and Zizhuo Wang. We Are on the Way: Analysis of On-Demand Ride-Hailing Systems. *Working Paper, University of Minnesota*, 2017.

Kris Johnson Ferreira, Bin Hong Alex Lee, and David Simchi-Levi. Analytics for an Online Retailer: Demand Forecasting and Price Optimization. *Manufacturing & Service Operations Management*, 18(1):69–88, 2016.

Guillermo Gallego and Ming Hu. Dynamic Pricing of Perishable Assets Under Competition. *Management Science*, 60(5):1241–1259, 2014.

Guillermo Gallego and Huseyin Topaloglu. *Revenue Management and Pricing Analytics*. Springer, 2019.

Guillermo Gallego and Garrett van Ryzin. Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons. *Management Science*, 40(8):999–1020, 1994.

Guillermo Gallego and Garrett van Ryzin. A Multiproduct Dynamic Pricing Problem and its Applications to Network Yield Management. *Operations Research*, 45(1):24–41, 1997.

Guillermo Gallego, Garud Iyengar, Robert Phillips, and Abha Dubey. Managing Flexible Products on a Network. 2004. CORC Technical Report TR-2004-01, Columbia University.

Nikhil Garg and Hamid Nazerzadeh. Driver Surge Pricing. *Working Paper, Stanford University*, 2019.

Negin Golrezaei, Hamid Nazerzadeh, and Paat Rusmevichientong. Real-Time Optimization of Personalized Assortments. *Management Science*, 60(6):1532–1551, 2014.

Negin Golrezaei, Max Lin, Vahab Mirrokni, and Hamid Nazerzadeh. Boosted Second Price Auctions: Revenue Optimization for Heterogeneous Bidders. *Working Paper, MIT*, 2017.

Negin Golrezaei, Adel Javanmard, and Vahab Mirrokni. Dynamic Incentive-aware Learning: Robust Pricing in Contextual Auctions. *Working Paper, MIT*, 2018.

Michael Harrison, Bora Keskin, and Assaf Zeevi. Bayesian Dynamic Pricing Policies: Learning and Earning Under a Binary Prior Distribution. *Management Science*, 58(3):570–586, 2012.

Pavithra Harsha, Shivaram Subramanian, and Joline Uichanco. Dynamic Pricing of Omnichannel Inventories. *Manufacturing & Service Operations Management*, 21(1):47–65, 2019.

Ming Hu and Yun Zhou. Dynamic Type Matching. *Working Paper, University of Toronto*, 2016.

Krishnamurthy Iyer, Ramesh Johari, and Mukund Sundararajan. Mean Field Equilibria of Dynamic Auctions with Learning. *Management Science*, 60(12):2949–2970, 2014.

Adel Javanmard and Hamid Nazerzadeh. Dynamic Pricing in High-dimensions. *Journal of Machine Learning Research*, 20(1):315–363, 2019.

Sham M Kakade, Ilan Lobel, and Hamid Nazerzadeh. Optimal Dynamic Mechanism Design and the Virtual-Pivot Mechanism. *Operations Research*, 61(4):837–854, 2013.

Yash Kanoria and Hamid Nazerzadeh. Dynamic Reserve Prices for Repeated Auctions: Learning from Bids. *Working Paper, Columbia University*, 2017.

Yash Kanoria and Pengyu Qian. Near Optimal Control of a Ride-Hailing Platform via Mirror Backpressure. *Working Paper, Columbia University*, 2019.

Bora Keskin and Assaf Zeevi. Dynamic Pricing with an Unknown Linear Demand Model: Asymptotically Optimal Semi-Myopic Policies. *Operations Research*, 62(5):1142–1167, 2014.

Nikita Korolko, Dawn Woodard, Chiwei Yan, and Helin Zhu. Dynamic Pricing and Matching in Ride-hailing Platforms. *Naval Research Logistics*, 2019.

Nitish Korula, Vahab Mirrokni, and Hamid Nazerzadeh. Optimizing Display Advertising Markets: Challenges and Directions. *IEEE Internet Computing*, 20(1):28–35, 2015.

Kenneth Littlewood. Forecasting and Control of Passenger Bookings. In *Airline Group International Federation of Operational Research Societies Proceedings*, pages 95–117, 1972.

Qian Liu and Garrett van Ryzin. On the Choice-Based Linear Programming Model for Network Revenue Management. *Manufacturing & Service Operations Management*, 10(2): 288–310, 2008.

Ilan Lobel, Renato Paes Leme, and Adrian Vladu. Multidimensional Binary Search for Contextual Decision-Making. *Operations Research*, 66(5):1346–1361, 2018.

Alice Lu, Peter I. Frazier, and Oren Kislev. Surge Pricing Moves Uber's Driver-Partners. In *Proceedings of EC*, 2018.

Hongyao Ma, Fei Fang, and David C Parkes. Spatio-Temporal Pricing for Ridesharing Platforms. *Working Paper, Harvard University*, 2018.

Jeffrey I McGill and Garrett J Van Ryzin. Revenue Management: Research Overview and Prospects. *Transportation Science*, 33(2):233–256, 1999.

Vahab Mirrokni and Hamid Nazerzadeh. Deals or No Deals: Contract Design for Online Advertising. In *Proceedings of WWW*, pages 7–14, 2017.

Velibor V Mišić and Georgia Perakis. Data Analytics in Operations Management: A Review. *Manufacturing & Service Operations Management*, 22(1):158–169, 2020.

Roger B Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1): 58–73, 1981.

Erhun Ozkan. Joint Pricing and Matching in Ridesharing Systems. *Working Paper, Koc University*, 2019.

Erhun Ozkan and Amy R Ward. Dynamic Matching for Real-time Ridesharing. *Stochastic Systems*, 2019.

Ioannis Paschalidis and John N Tsitsiklis. Congestion-Dependent Pricing of Network Services. *IEEE/ACM Transactions on Networking*, 8(2):171–184, 2000.

Sheng Qiang and Mohsen Bayati. Dynamic Pricing with Demand Covariates. *Working Paper, Stanford University*, 2016.

Paat Rusmevichientong, Zuo-Jun Max Shen, and David B Shmoys. Dynamic Assortment Optimization with a Multinomial Logit Choice Model and Capacity Constraint. *Operations Research*, 58(6):1666–1680, 2010.

Denis Sauré and Assaf Zeevi. Optimal Dynamic Assortment Planning with Demand Learning. *Manufacturing & Service Operations Management*, 15(3):387–404, 2013.

David Simchi-Levi. OM Research: From Problem-Driven to Data-Driven Research. *Manufacturing & Service Operations Management*, 16(1):2–10, 2014.

Arne K Strauss, Robert Klein, and Claudius Steinhardt. A Review of Choice-Based Revenue Management: Theory and Methods. *European Journal of Operational Research*, 271(2): 375–387, 2018.

Xuanming Su. Intertemporal Pricing with Strategic Customer Behavior. *Management Science*, 53(5):726, 2007.

Kalyan Talluri and Garrett van Ryzin. An Analysis of Bid-price Controls for Network Revenue Management. *Management Science*, 44(11):1577–1593, 1998.

Kalyan Talluri and Garrett van Ryzin. Revenue Management Under a General Discrete Choice Model of Consumer Behavior. *Management Science*, 50(1):15–33, 2004a.

Kalyan Talluri and Garrett van Ryzin. *The Theory and Practice of Revenue Management*. Springer Science & Business Media, 2004b.

Xinan Yang, Arne K Strauss, Christine SM Currie, and Richard Eglese. Choice-Based Demand Management and Vehicle Routing in E-Fulfillment. *Transportation Science*, 50 (2):473–488, 2016.