

20TH ANNIVERSARY INVITED ARTICLE

Operations Management in the Age of the Sharing Economy: What Is Old and What Is New?

 Saif Benjaafar,^a Ming Hu^b
^aDepartment of Industrial and Systems Engineering, University of Minnesota, Minneapolis, Minnesota 55455;

^bRotman School of Management, University of Toronto, Toronto, Ontario M5S 3E6 Canada

Contact: saif@umn.edu,  <http://orcid.org/0000-0002-1949-0105> (SB); ming.hu@rotman.utoronto.ca,

 <http://orcid.org/0000-0003-0900-7631> (MH)

Received: April 8, 2019

Accepted: April 8, 2019

Published Online in Articles in Advance:
August 1, 2019

<https://doi.org/10.1287/msom.2019.0803>
Copyright: © 2019 INFORMS

Abstract. The *sharing economy*, a term we use to refer to business models built around on-demand access to products and services mediated by online platforms that match many small suppliers or service providers to many small buyers, has emerged as an important area of study in operations management. We first describe three “canonical” applications that have garnered much attention from the operations management community. We use these applications to highlight distinguishing features of sharing economy business models and to point out research questions that are new. Then we draw connections between classical operations management theory and models and those that have been used to study sharing economy applications. We do so to put in context some of the recent work on the sharing economy and to showcase the underlying modeling toolkit and identify opportunities for future research.

History: This paper has been accepted for the *Manufacturing & Service Operations Management* 20th Anniversary Special Issue.

Funding: The research of S. Benjaafar is in part funded by the United States National Science Foundation [Award SCC-1831140 (Leveraging Autonomous Shared Vehicles for Greater Community Health, Equity, Livability, and Prosperity)]. The research of M. Hu is in part supported by the Natural Sciences and Engineering Research Council of Canada [Grant RGPIN-2015-06757] and National Natural Science Foundation of China [Grant 71825007].

Keywords: sharing economy • matching markets • peer-to-peer resource sharing • on-demand service platforms • on-demand rental networks • dynamic matching • dynamic pricing

1. Introduction

The term *sharing economy* has been used in the operations management (OM) community to refer to business models built around on-demand access to products and services mediated by online platforms that match many small suppliers or service providers to many small buyers. It has also been used to refer to on-demand business models where products or services are provided by a single entity to many small buyers. A prominent example of the first type is ride-hailing platforms that match independent drivers with individual riders. An example of the second type is vehicle-sharing services that provide short-term vehicle rentals to individual users.

In this paper, we seek to relate problems that arise in the context of sharing economy business models (of either the first or second type) to classical OM problems. In doing so, we highlight problems that can be recast as instances of classical OM problems and those that cannot. We use the case of the latter to highlight unique features that distinguish certain

sharing economy applications and discuss the need for different models and approaches. We also discuss the broader managerial implications that arise from the analysis of these applications.

We primarily focus on research that is grounded in analytical models and approaches [see Hu (2019a) for some references of related empirical, experimental, and behavioral studies]. In this paper, we adopt a relatively narrow view of the sharing economy. There are other manifestations of the sharing economy that do not involve the on-demand access to a product or a service but that nevertheless involve the matching of many buyers and suppliers mediated by online platforms.¹ The interested reader is referred to the recent reviews by Chen et al. (2019d) and Hu (2019b).

The rest of this paper is organized in two parts. In the first part (Section 2), we set the stage by describing three canonical applications of the sharing economy. We use these applications to highlight important research questions and to review some of the related OM research. Because of the newness of the topic,

much of the relevant literature is from very recent years. In the second part of the paper (Section 3), we draw connections between problems that arise in sharing economy contexts and classical problems that have been widely studied by the OM and operations research communities, with a particular emphasis on the underlying model formulations. We do so to suggest opportunities for extending these models to account for the unique feature of some sharing economy applications. In Section 4, we offer concluding comments and discuss themes that may emerge as important topics of future research.

2. Three Canonical Sharing Economy Applications

In this section, we describe three canonical applications that have garnered much attention from the OM community. We use these applications to highlight important distinguishing features of sharing economy business models and to point out research questions that are new.

2.1. Peer-to-Peer Resource Sharing

Modern economies are built on a model of consumption that involves exclusive ownership and usage of resources. This model can be inefficient if the privately owned resources are poorly utilized or if their usage is intermittent. Platforms that allow peer-to-peer resource sharing enable owners to rent their assets on a short-term basis to nonowners. In a population where individuals may have different levels of usage, those with high usage levels would favor (all else being equal) being owners, whereas those with low usage levels would favor being nonowners. These choices are modulated by, among other factors, the individuals' sensitivity to the associated inconvenience, the prevailing rental price, and the likelihood of a successful match between an owner and a nonowner.

Peer-to-peer resource sharing, as just described, is a departure from traditional modes of demand and supply in the following important ways:

- *There are many buyers and sellers.* The platform, in its role as an intermediary, reduces market friction, including search costs, transaction costs, and moral hazard, enabling the participation of many small actors on both the buy side and the sell side.

- *The supply side is not distinct from the demand side.* Having more individuals choosing to be owners means fewer renters, and vice versa.

- *Supply stimulates demand, and vice versa.* Having more owners increases the likelihood of a successful rental for nonowners, whereas having more nonowners increases the likelihood of a successful rental for owners.

- *A resource unit can sustain the consumption needs from more than one consumer.* Because resources are not fully utilized, they can fulfill the usage needs from multiple consumers.

Research questions that arise include how platforms should price their services to owners and nonowners, how the matching of owners and nonowners should be carried out (particularly in settings where resources are heterogeneous), and how much effort the platform should exert in reducing market friction.

The presence of peer-to-peer sharing affects original equipment manufacturers. On the one hand, peer-to-peer sharing can curtail ownership, because individuals have the option of renting instead of owning (a *cannibalization* effect). On the other hand, it can make ownership more attractive, because the rental income could make owning more affordable (a *value enhancement* effect). Shared products experience more usage, resulting in more frequent product replacement, increasing revenue (a *usage* effect). In view of these effects, how should a platform price its products, and under what conditions would a platform benefit from the peer-to-peer sharing of its products? Should a manufacturer sidestep the sharing market by offering products for short-term rental instead of (or in addition to) selling them outright? These and related questions are explored in Abhishek et al. (2019), Blaettchen et al. (2018), Benjaafar and Pourghannad (2019), Jiang and Tian (2018), and Tian and Jiang (2018). The question of whether a manufacturer should offer its products for short-term rental (or, more generally, under a pay-per-use revenue model) is often referred to as *servicization* and is explored in Agrawal and Bellos (2017), Bellos et al. (2017), and Benjaafar et al. (2019c).

Finally, peer-to-peer sharing affects resource usage and ownership. Hence, questions regarding consumer surplus, social welfare, and environmental sustainability arise naturally. Peer-to-peer product sharing has the potential of increasing access while reducing the base of resources needed to secure this access. This could have the twin benefit of improving consumer welfare (individuals who may not otherwise afford a product have an opportunity to use it) while reducing societal costs (externalities, such as pollution, which may be associated with the production, distribution, use, and disposal of the product). It also has the potential of providing a source of income for owners. However, increased sharing may have other consequences, some of them undesirable. For example, greater access to resources could increase their usage and the associated negative externalities of such usage. Increased sharing may also lead to more ownership because products that were too expensive now become more affordable thanks to the rental income. Benjaafar et al. (2019c) examine these issues and show that peer-to-peer

sharing could lead to more ownership and more usage if the cost of ownership is sufficiently high; see Fraiberger and Sundararajan (2017) for a structural estimation model. Benjaafar et al. (2017) explore these issues in the context of car sharing, where sharing takes the form of owners offering rides to nonowners, and show that the ratio of ownership cost to usage cost determines whether sharing leads to more or less ownership and traffic.

2.2. On-Demand Service Platforms

On-demand service platforms are platforms that connect workers, acting as independent agents, with customers who require a time-sensitive service. Examples of such platforms are many and include in the United States, among others, Uber and Lyft for transportation services, Instacart and Postmates for home deliveries, and TaskRabbit and Handy for household tasks. Workers decide on how much time, if any, to devote to the platform. Workers are typically heterogeneous in their opportunity costs. So how much time workers devote to the platform can vary from one worker to another. Workers can be heterogeneous in other ways, including *spatially* in terms of the geographic area where they are willing to work, *temporally* in terms of when they can work, and *by capability* in terms of the type and quality of work they can provide. Because of the on-demand feature of the service provided, customers are typically sensitive not only to price but also to the delay they experience prior to receiving the service and to the quality of the service (often determined by the features of the service sought and the capability of the assigned worker). In settings such as ride-hailing, where the work is not overly specialized and workers do not vary greatly in capability, the platform typically sets the wages it pays to workers as well as the prices it charges customers. Higher wages induce more workers to join the platform and those who join to devote more time, whereas lower prices induce higher demand. Because workers are typically paid only when they are busy doing work, they are sensitive not only to the nominal wage they receive for work completed but also to the fraction of time they expect to be busy (i.e., their utilization).

On-demand platforms have the following distinctive features from traditional service systems:

- *Capacity affects demand, and vice versa.* Demand is stimulated by capacity (higher capacity means lower delay), and capacity is stimulated by demand (more demand means higher effective wages).

- *Capacity can be controlled only indirectly via wages and prices.* Capacity is determined by the decisions of independent workers who decide when, how much, and where to work and respond to the incentives of wages and demand.

- *Capacity and demand vary temporally and spatially.* In addition to temporal mismatches between supply and demand, there are also spatial mismatches.

Research questions that arise include questions about how platforms should price their services and how much they should pay their workers. Cachon et al. (2017) show that there can be a significant benefit to the platform from a policy that adjusts wages and prices dynamically. They show that both workers and customers can benefit from such a policy (relative to one where wages or prices are kept fixed). Taylor (2018) studies the impact of suppliers' independence, customers' congestion-driven delay disutility, and uncertainty in customers' or suppliers' valuation on the optimal wage and price. Hu and Zhou (2019) study the widely practiced commission contract under which the platform takes a fixed cut, and thus the wage is equal to a fraction of the price regardless of what price is charged. They show that under some conditions, an optimal flat-commission contract performs nearly as well as an optimal contingent contract where wage and price are set independently for each market condition. Hu et al. (2019) show that the optimal inter-temporal pricing can follow the pattern of a short-lived sharp price surge succeeded by a lower price, or take the form of a low initial price followed by a higher price. Bai et al. (2018) show that the optimal price a platform should charge is not necessarily monotonic in the market size when either the labor pool size or the waiting cost is high. This result is consistent with those obtained in Benjaafar et al. (2018), who show that the optimal price initially increases in the labor pool size and then decreases. Cohen et al. (2018) run field experiments to investigate how compensation schemes affect the repeated engagement of riders who experience a frustration. Gurvich et al. (2019) allow the platform to set a capacity cap on the number of drivers who can work, in addition to wage or price decisions. Chu et al. (2018) study how to moderate the drivers' cherry-picking behavior.

Another important research question is how service requests should be matched with service providers when there is heterogeneity in the service requests or among the service providers. Ozkan and Ward (2019) study this question in the context of ride hailing. App-based ride hailing allows the matching of a driver and a rider even if they are not colocated. A disadvantage of such matching is that it could result in long pickup times under certain dispatching policies. Feng et al. (2018) show that this is the case when the system utilization is in the midrange and the service area is large. The effect of pickup times on the efficiency of ride-hailing services is also studied in Castillo et al. (2018), who use the phrase "wild goose chase" to refer to the phenomenon of dispatching

drivers to far-away ride requests. They show that this is more likely to occur when demand is high and that the phenomenon can be mitigated by surge pricing. Besbes et al. (2018a) take into account the pickup time in the capacity planning decision of a ride-hailing system. Guo et al. (2018) study how to mitigate the gender mismatch in ride hailing, which may lead to safety issues.

Because an on-demand service platform relies on independent workers who do not enjoy the same regulatory protections granted to employees, employment through these platforms (or “gig” work) has come under scrutiny. Labor advocates have argued that the growth of these platforms has come at the expense of workers. Industry groups representing incumbent industries have also argued that these platforms compete unfairly against established businesses that are subject to regulation. Several large cities have raised concerns about the impact of growth in ride-hailing services on congestion and pollution. These various concerns have led to calls for stricter regulation. This raises important research questions regarding the extent to which growth in on-demand services harms or benefits workers, customers, and the environment. Benjaafar et al. (2018) show that growth in the labor pool size may not necessarily be harmful to workers, with initial increases in the labor pool stimulating demand, making workers busier, and increasing their effective wages. Yu et al. (2017) show that moderate forms of regulation improve social welfare. Hu and Zhou (2019) study the impact of minimum wage on the platform, drivers, and riders. In an empirical study, Gong et al. (2019) find that Uber’s entry in several Chinese cities has led to a considerable increase in new vehicle ownership. Burtch et al. (2018) find that the entry of Uber in U.S. cities has been accompanied by a decrease in entrepreneurial activity, seemingly by offering viable employment to the unemployed and underemployed. Ming et al. (2019) empirically show that surge pricing generally improves consumer and driver welfare as well as platform profits. Chevalier et al. (2018) find, using data from Uber, that drivers benefit significantly from being able to decide on when and how much to work. Finally, there is a growing body of empirical research that examines how workers make decisions about when and how much to work [see, e.g., Allon et al. (2018) and the references therein].

2.3. On-Demand Rental Networks

The preceding two applications are built around models of peer-to-peer interactions between resource owners and renters or between service providers and users. An important application of the sharing economy (on-demand access to products and services) is one that does not rely on individuals for the provisioning of

products and services and does not involve the mediation of a platform. Instead, products and services are provided by a single firm that retains ownership of the underlying resources. However, customers access these products and services as needed, including for brief periods of usage. A prominent example of this business model is vehicle sharing (e.g., the sharing of cars, bikes, and scooters), with vehicles available to customers for short-term rentals. Although the concept of renting instead of owning is not new, what is perhaps new is the possibility of renting on a short-term basis and the ease and convenience of such rentals. Other distinguishing features from traditional renting include the following:

- *No advance booking is required.* Customers do not typically provide the firm with advanced notice of when or where they plan to access the service.
- *Resources are spatially distributed.* To maximize accessibility, the firm typically distributes its inventory across multiple locations.
- *Resources rented from one location can be returned to another.* In many cases, customers are given the option of returning a resource rented from one location to another. In some cases, as in *free-floating* vehicle sharing systems, there are no designated pickup and drop-off locations.

Decisions the provider of an on-demand rental service must make can be classified based on the time scales involved and the ease with which the decision can be reversed. Over the long term, the service provider must decide on the sizing of the service, including the rental capacity, the size of the service region, and the location and size of designated pickup and drop-off points, if any. Over shorter time scales, the service provider must decide on the pricing of the service and the management of temporal and spatial mismatches between supply and demand such as the relocation of physical assets. The research on on-demand rental services, particularly as it pertains to bike and car sharing, is arguably the most mature of the three applications we discuss and has attracted substantial attention not only from the OM community but also from the operations research and engineering communities. A review of this literature can be found in Freund et al. (2019) and He et al. (2019b). Recent papers include Lu et al. (2017) and Kabra et al. (2018).

In what follows, we highlight two research questions and contrast these with related questions that arise in more traditional problems. The first concerns how best to reposition resources to mitigate the mismatch between supply and demand. A mismatch could arise because of the randomness of demand at each rental location, randomness in the location at which resources are returned, and randomness in the rental period. Hence, there is a need to periodically intervene to move inventory away from locations

with “too much” inventory into locations with “not enough.” The problem shares features of a classical inventory problem, except that the total number of units in circulation is now fixed, with some available for rent, while the rest are currently rented. Rental periods can be viewed as replenishment lead times (a unit becomes available again once a rental is complete). However, *replenishment* is not a controlled process in this case, and the location at which a rented unit will be returned is not known with certainty. Because of the multidimensionality of the problem, the heterogeneities across locations, and the various types of uncertainty that must be accounted for, characterizing an optimal policy has been elusive. Benjaafar et al. (2019b) showed recently that an optimal policy can be described by a well-specified region over the state space. Within this region, it is optimal not to reposition while outside of it; it is optimal to reposition but only such that the system moves to a new state that is on the boundary of the no-repositioning region. He et al. (2019a) provide an approximate solution approach to the problem using robust optimization.

The second question concerns the sizing of a rental network to meet a desired service level (e.g., the fraction of fulfilled demand). The dynamics of a rental network can be approximated by those of a queuing network, with each location corresponding to a multiserver queue. However, in contrast to a standard multiserver queue, upon completion, a server may not return to its original queue and instead is probabilistically routed to one of the other queues. Hence, the number of servers associated with each queue is now no longer fixed and is random instead. This introduces a new source of variability in the system. Papers that use queuing analysis to study on-demand rental networks include He et al. (2017), Benjaafar et al. (2019a), George and Xia (2011), Banerjee et al. (2017), and Braverman et al. (2019). In particular, Benjaafar et al. (2019a) provide closed-form approximations for optimal network capacity that are asymptotically exact.

3. From the Old to the New

In this section, we explore connections between classical OM theory/models and theory/models that have been used to study sharing economy applications. We do so to put in context some of the recent work on the sharing economy and to showcase the underlying modeling toolkit and identify opportunities for future research. We focus on three major pillars of OM theory: inventory theory, revenue management, and queuing theory. See Hu (2019a) for an expanded discussion.

3.1. Inventory Theory: From Controlled to Uncontrolled Supply

A principal concern of classical inventory theory is how best to orchestrate supply to meet uncertain or

time-varying demand (or to mitigate fixed costs associated with the delivery of this supply). Hence, decisions in classical inventory theory revolve around how much to order and when. Although the effective matching of supply and demand is also a primary concern in sharing economy applications, in many of these applications, supply (at least over short time scales) cannot be controlled. In fact, in some cases, it is best to view supply and demand as taking place according to processes that are beyond the reach of a coordinating platform. For example, one could view the process through which drivers become available on a ride-hailing platform as largely independent in the same way that the arrival process of customers is. As a result, the decisions of the platform now are less about orchestrating the timing of supply and more about how best to match each unit of demand with each of unit of supply (e.g., which drivers to dispatch to which riders given their respective geographic locations). The problem can be formulated as a dynamic optimization problem where the platform earns a reward that is a function of the quality of the matches it executes and of the numbers of supply units left unmatched (analogous to excess inventory in an inventory problem) and the number of customers whose request are unfulfilled (analogous to shortages in an inventory problem).

Hu and Zhou (2018) consider such a problem and show that, under some conditions, structural properties of optimal matching policies can be derived. More specifically, they show that there exist state-dependent thresholds, called *match-down-to levels*, governing the matching of a specific pair of supply and demand types. Only if the available amounts of resources exceed those levels is it optimal to match the supply and demand types down to those levels. If some pair of supply and demand types are not matched greedily, all pairs that are strictly dominated by this pair should not be matched at all, as a result of the priority structure. Chen et al. (2019a) show a similar structural property for the optimal matching under a supermodular matching reward structure with impatient demand but patient supply. Such a structural property of *priority and thresholds* is a generalization of priority structures seen in the inventory management literature (such as base-stock levels) and quantity-based revenue management literature (such as protection levels). Because of these connections, methodologies, techniques, and insights developed for one domain can be transferred to the other. For example, Hu and Zhou (2018) further show that by verifying the L^1 -concavity of the value functions of a transformed problem, the optimal total matching quantity or the optimal match-down-to levels have monotonicity properties with respect to the system state. This technique has been applied

to derive structural properties for lost-sales inventory models (Zipkin 2008).

Hence, despite the differences in how supply and demand are manifested in sharing economy applications and classical inventory problems, we expect that many of the techniques developed for inventory management continue to be useful for sharing economy applications. Similarly, we expect that computational methods such as approximate dynamic programming (ADP) continue to be useful for dealing with large-scale problems; see Benjaafar et al. (2019b) for an application of ADP to the inventory repositioning problem in an on-demand rental network.

3.2. Revenue Management: From Exogenous to Endogenous Capacity

Classical revenue management is concerned with how prices should be varied to sell a *fixed* amount of capacity over a finite horizon (see, e.g., Gallego and Van Ryzin 1994). An essential insight from revenue management is that prices should be dynamically adjusted based on the available remaining capacity, with less capacity translated to higher prices. This concept has been embraced by many sharing economy platforms, taking the form of *surge* pricing in the case of ride hailing. However, in contrast to classical revenue management, in many sharing applications, capacity is crowdsourced and is sensitive to the paid wages. Hence, there is an opportunity for two-sided dynamic pricing (wages paid to workers and prices charged to customers) based on currently prevailing supply and demand conditions. Recent research includes comparing contingent pricing, reacting to the market condition, with static pricing. Cachon et al. (2017) show that drivers and riders are generally better off with prices being contingent on varying market conditions. Banerjee et al. (2015) show that a static pricing policy can be asymptotically optimal in a thick market, results that are perhaps more analogous to those from classical revenue management (see, e.g., Gallego and Van Ryzin 1994). Chen and Hu (2019) study a version of two-sided pricing where sellers and buyers sequentially arrive to the market and can strategically time their transactions. A static pricing policy by the intermediary platform has an advantage of deterring strategic waiting behavior of sellers and buyers, which is difficult to account for by pricing models. As a result, with static pricing, all participants can be brought in as they arrive, increasing the thickness of the market. Chen et al. (2019c) consider a setting that allows the intermediary to buy in, sell out, and hold inventory.

In addition to the endogeneity of capacity, a distinguishing feature of sharing economy applications is the spatially distributed nature of supply and

demand. A matching platform may engage in dynamic two-sided pricing across multiple locations, with spillover effects among the different locations. For example, higher wages in one location draw supply away from other locations. Similarly, low demand in one location could increase supply in other locations (see Bimpikis et al. 2019). Afèche et al. (2018) show that it may be optimal for a ride-hailing platform to strategically reject demand at a low-demand location, although there is an excess supply of drivers, to induce their repositioning to the high-demand location. Besbes et al. (2018b) show that it can be optimal for a platform to use prices to create regions where driver congestion is artificially high in order to lure drivers toward more profitable locations for the platform (see also Guda and Subramanian 2019).

There is extensive research in economics that considers pricing in the context of two-sided markets, such as credit cards (bringing together card holders and retailers) or video game platforms (bringing together game developers and players; see, e.g., Rochet and Tirole 2003). In this literature, transaction volumes often take a *multiplicative* form of supply and demand volumes. For example, in the case of credit cards, transaction volumes are a function of the product of the number of participating retailers and the number of card holders. By contrast, in sharing economy applications, transaction volumes are a function of the *minimum* of supply and demand. This is because successful transactions typically involve one-to-one matching of supply and demand units. Hu (2019a) illustrates that the transaction volume taking the minimum form results in a very different pricing problem from the two-sided pricing in the economics literature involving the product of supply and demand.

The marriage of inventory theory and revenue management gave birth to a stream of research on joint pricing and inventory control (see, e.g., Federgruen and Heching 1999, Chen and Simchi-Levi 2004). There is a similar opportunity in the context of the sharing economy for research that would consider pricing and matching decisions jointly, with pricing as a mechanism to improve the volume and quality of matches and for differentiated pricing to enable greater efficiency from the executed matches. For example, some ride-hailing platforms such as Didi Chuxing now allow customers to offer additional payments to increase their chances of being matched or of being matched faster.

3.3. Queuing Systems: From Fixed to Random Number of Servers

Queuing theory is focused on the study of how congestion arises when there is finite service capacity and demand occurs continuously over time with

randomness in the interarrival times between consecutive service requests. Because many sharing economy applications are concerned with the on-demand access to products and services, queuing theory is a natural candidate as a tool for examining how congestion might arise in the context of these applications. However, as we noted in Section 2, there are several distinguishing features to these applications that may require extending the existing theory. We highlight three such features below.

3.3.1. Self-Scheduled Servers. As mentioned in Section 2.2, many on-demand service platforms rely on independent workers who decide when and how much to work. The set of workers can be modeled as a multiserver queuing system. However, the number of servers now is no longer fixed and varies over time, with some workers joining and others leaving. Ibrahim (2018) considers the optimal staffing in a service system where the number of servers is subject to variability. Dong and Ibrahim (2017) consider a system with a blended workforce composed of both full-time and contingent agents hired on a part-time basis (in response to fluctuations in demand).

3.3.2. Servers with Probabilistic Returns. As mentioned in Section 2.2, on-demand rental networks can be modeled as a set of *communicating* multiserver queues. A server upon completing a service does not necessarily return to its original queue but is instead routed probabilistically to any one of the queues. Hence, each queue can now be viewed as having a stochastic number of servers. Benjaafar et al. (2019a) show that the variability in the number of servers can have a significant impact on the performance of these systems. In particular, the sizing of such a system must account for the variability not only in arrivals and services but also in the routing of the servers.

3.3.3. Double-Ended Queues. A double-ended queue is one where the arrivals of customers and servers are independent and exogenously specified processes; see Kendall (1951) for an early reference and Afèche et al. (2014) for a more recent one. A single-location taxi stand is often modeled as a double-ended queue in which drivers and riders arrive at the taxi stand independently. Once a driver and a rider pair up, they leave the system. A double-ended queue can be a more appropriate modeling framework for a matching market where service providers do not necessarily return to the original point of service initiation after service completion. A double-ended queue could also be a useful model for peer-to-peer ride sharing (carpooling), in which, for example, commuters offer rides to others on their way to work or home.

4. Concluding Comments

We conclude this paper by highlighting a few areas for future research that are natural extensions of the early work on the sharing economy we described so far..

From Monopoly to Competition. Competition, particularly in the context of two-sided platforms, has yet to receive much attention. Because platforms compete for both supply and demand, it is not clear whether standard results would continue to hold. For example, Nikzad (2017) shows that two-sided competition can result in higher prices because of the resulting higher cost of supply. See Cohen and Zhang (2017), Bai and Tang (2018), Bernstein et al. (2019), and Hu and Liu (2019) for other examples.

From Two-Sided to Multisided Platforms. A growing number of on-demand service platforms involve more than two sides. For example, food delivery involves drivers, restaurants, and customers. Understanding how the incentives and the payoffs of the various parties are affected by multisidedness would be an important area of research.

From Workers to Machines. There is a growing consensus that machines, equipped with artificial intelligence capabilities, will replace humans in many endeavors. For ride hailing, driverless cars are likely to replace human drivers in the near future. The implications of such smart machines, especially on a large scale, would also be an important topic of study.

From Theory to Empirics. In this paper, we focused on analytical work. Because the field is nascent, perhaps this makes sense. We expect the insights uncovered by the new theory will come under scrutiny and be tested by empirical research using data drawn either from practice or via laboratory experiments.

Endnote

¹Prominent examples include funding platforms that match small borrowers and entrepreneurs with small lenders and investors such as LendingClub and Kickstarter; information platforms that provide peer-to-peer instruction, reviews, and advice such as Yelp (see, e.g., Bimpikis and Papanastasiou 2019); and peer-to-peer marketplaces that involve the buying and selling of goods and services such as Etsy, Airbnb, and Upwork (see, e.g., Allon et al. 2012, Chen et al. 2019b, Cui et al. 2019).

References

- Abhishek V, Guajardo J, Zhang Z (2019) Business models in the sharing economy: Manufacturing durable goods in the presence of peer-to-peer rental markets. Working paper, Carnegie Mellon University, Pittsburgh.
- Afèche P, Diamant A, Milner J (2014) Double-sided batch queues with abandonment: Modeling crossing networks. *Oper. Res.* 62(5): 1179–1201.

- Afèche P, Liu Z, Maglaras C (2018) Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance. Working paper, University of Toronto, Toronto.
- Agrawal VV, Bellos I (2017) The potential of servicizing as a green business model. *Management Sci.* 63(5):1545–1562.
- Allon G, Bassamboo A, Çil EB (2012) Large-scale service marketplaces: The role of the moderating firm. *Management Sci.* 58(10):1854–1872.
- Allon G, Cohen M, Sinchaisri P (2018) The impact of behavioral and economic drivers on gig economy workers. Working paper, University of Pennsylvania, Philadelphia.
- Bai J, Tang CS (2018) Can two competing on-demand service platforms be both profitable? Working paper, State University of New York at Binghamton, Binghamton.
- Bai J, So KC, Tang CS, Chen X, Wang H (2018) Coordinating supply and demand on an on-demand service platform with impatient customers. *Manufacturing Service Oper. Management*, ePub ahead of print June 28, <https://doi.org/10.1287/msom.2018.0707>.
- Banerjee S, Freund D, Lykouris T (2017) Pricing and optimization in shared vehicle systems: An approximation framework. Working paper, Cornell University, Ithaca, NY.
- Banerjee S, Riquelme C, Johari R (2015) Pricing in ride-share platforms: A queueing-theoretic approach. Working paper, Cornell University, Ithaca, NY.
- Bellos I, Ferguson M, Toktay LB (2017) The car sharing economy: Interaction of business model choice and product line design. *Manufacturing Service Oper. Management* 19(2):185–201.
- Benjaafar S, Pourghannad B (2019) Peer-to-peer trading of usage allowances. Working paper, University of Minnesota, Minneapolis.
- Benjaafar S, Bernhard H, Courcoubetis C (2017) Drivers, riders and service providers: The impact of the sharing economy on mobility. Working paper, University of Minnesota, Minneapolis.
- Benjaafar S, Liu H, Wu S (2019a) One way vs. round trip: Sizing up vehicle sharing services. Working paper, University of Minnesota, Minneapolis.
- Benjaafar S, Ding J-Y, Kong G, Taylor T (2018) Labor welfare in on-demand service platforms. Working paper, University of Minnesota, Minneapolis.
- Benjaafar S, Jiang D, Li X, Li X (2019b) Inventory repositioning in on-demand product rental networks. Working paper, University of Minnesota, Minneapolis.
- Benjaafar S, Kong G, Li X, Courcoubetis C (2019c) Peer-to-peer product sharing: Implications for ownership, usage and social welfare in the sharing economy. *Management Sci.* 65(2):477–493.
- Bernstein F, DeCroix G, Keskin NB (2019) Competition between two-sided platforms under demand and supply congestion effects. Working paper, Duke University, Durham, NC.
- Besbes O, Castro F, Lobel I (2018a) Spatial capacity planning. Working paper, Columbia University, New York.
- Besbes O, Castro F, Lobel I (2018b) Surge pricing and its spatial supply response. Working paper, Columbia University, New York.
- Bimpikis K, Papanastasiou Y (2019) Inducing exploration in service platforms. Hu M, ed. *Sharing Economy: Making Supply Meet Demand* (Springer, Cham, Switzerland), 193–216.
- Bimpikis K, Candogan O, Saban D (2019) Spatial pricing in ride-sharing networks. *Oper. Res.* 67(3):744–769.
- Blaettchen P, Hasija S, Taneri N (2018) Sharing of durable goods: Business models for original equipment manufacturers. Working paper, INSEAD, Fontainebleau, France.
- Braverman A, Dai J, Liu X, Ying L (2019) Empty-car routing in ridesharing systems. *Oper. Res.* Forthcoming.
- Burtch G, Carnahan S, Greenwood BN (2018) Can you gig it? An empirical examination of the gig economy and entrepreneurial activity. *Management Sci.* 64(12):5497–5520.
- Cachon GP, Daniels KM, Lobel R (2017) The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing Service Oper. Management* 19(3):368–384.
- Castillo JC, Knoepfle D, Weyl EG (2018) Surge pricing solves the wild goose chase. Working paper, Stanford University, Stanford, CA.
- Chen Y, Hu M (2019) Pricing and matching with forward-looking buyers and sellers. *Manufacturing Service Oper. Management*. Forthcoming.
- Chen X, Simchi-Levi D (2004) Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The finite horizon case. *Oper. Res.* 52(6):887–896.
- Chen Z, Hu M, Zhou Y (2019a) Optimal dynamic matching under impatient demand and patient supply. Working paper, University of Toronto, Toronto.
- Chen Y-J, Maglaras C, Vulcano G (2019b) Design of an aggregated marketplace under congestion effects: Asymptotic analysis and equilibrium characterization. Hu M, ed. *Sharing Economy: Making Supply Meet Demand* (Springer, Cham, Switzerland), 217–248.
- Chen N, Gallego G, Gao P, Kou S (2019c) Dynamic two-sided markets: Dealership and marketplace. Working paper, Hong Kong University of Science and Technology, Clear Water Bay.
- Chen Y-J, Dai T, Körpeoğlu CG, Körpeoğlu E, Sahin O, Tang CS, Xiao S (2019d) Innovative online platforms: Research opportunities. *Manufacturing Service Oper. Management*. Forthcoming.
- Chevalier JA, Chen MK, Oehlsen E, Rossi PE (2018) The value of flexible work: Evidence from Uber drivers. *J. Political Econom.*, ePub ahead of print December 12, <https://doi.org/10.1086/702171>.
- Chu LY, Wan Z, Zhan D (2018) Harnessing the double-edged sword via routing: Information provision on ride-hailing platforms. Working paper, University of Southern California, Los Angeles.
- Cohen M, Zhang R (2017) Competition and cooperation for two-sided platforms. Working paper, New York University, New York.
- Cohen M, Fiszer M-D, Kim BJ (2018) Frustration-based promotions: Field experiments in ride-sharing. Working paper, New York University, New York.
- Cui R, Li J, Zhang D (2019) Discrimination with incomplete information in the sharing economy: Evidence from field experiments on Airbnb. *Management Sci.* Forthcoming.
- Dong J, Ibrahim R (2017) Managing supply in the on-demand economy: Flexible workers or full-time employees? Working paper, Columbia University, New York.
- Federgruen A, Heching A (1999) Combined pricing and inventory control under uncertainty. *Oper. Res.* 47(3):454–475.
- Feng G, Kong G, Wang Z (2018) We are on the way: Analysis of on-demand ride-hailing systems. Working paper, University of Minnesota, Minneapolis.
- Fraiberger SP, Sundararajan A (2017) Peer-to-peer rental markets in the sharing economy. Working paper, Harvard University, Cambridge, MA.
- Freund D, Henderson SG, Shmoys DB (2019) Bike sharing. Hu M, ed. *Sharing Economy: Making Supply Meet Demand* (Springer, Cham, Switzerland), 435–459.
- Gallego G, Van Ryzin G (1994) Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* 40(8):999–1020.
- George DK, Xia CH (2011) Fleet-sizing and service availability for a vehicle rental system via closed queueing networks. *Eur. J. Oper. Res.* 211(1):198–207.
- Gong J, Greenwood BN, Song Y (2018) Uber might buy me a Mercedes Benz: An empirical investigation of the sharing economy and durable goods purchase. Working paper, Temple University, Philadelphia.
- Guda H, Subramanian U (2019) Your Uber is arriving: Managing on-demand workers through surge pricing, forecast communication, and worker incentives. *Management Sci.*, 65(5):1995–2014.

- Guo P, Tang CS, Tang Y, Wang Y (2018) Gender-based operational issues arising from on-demand ride-hailing platforms: Safety concerns, service systems, and pricing and wage policy. Working paper, Hong Kong Polytechnic University, Hung Hom.
- Gurvich I, Lariviere M, Moreno A (2019) Operations in the on-demand economy: Staffing services with self-scheduling capacity. Hu M, ed. *Sharing Economy: Making Supply Meet Demand* (Springer, Cham, Switzerland), 249–278.
- He L, Hu Z, Zhang M (2019a) Robust repositioning for vehicle sharing. *Manufacturing Service Oper. Management*, ePub ahead of print April 24, <https://doi.org/10.1287/msom.2018.0734>.
- He L, Mak H-Y, Rong Y (2019b) Operations management of vehicle sharing systems. Hu M, ed. *Sharing Economy: Making Supply Meet Demand* (Springer, Cham, Switzerland), 461–484.
- He L, Mak H-Y, Rong Y, Shen Z-JM (2017) Service region design for urban electric vehicle sharing systems. *Manufacturing Service Oper. Management* 19(2):309–327.
- Hu M (2019a) From the classics to new tunes: A neoclassical view on sharing economy and innovative marketplaces. *Production Oper. Management*. Forthcoming.
- Hu M, ed. (2019b) *Sharing Economy: Making Supply Meet Demand* (Springer, Cham, Switzerland).
- Hu M, Liu Y (2019) Precommitments in two-sided market competition. Working paper, University of Toronto, Toronto.
- Hu M, Zhou Y (2018) Dynamic type matching. Working paper, University of Toronto, Toronto.
- Hu M, Zhou Y (2019) Price, wage and fixed commission in on-demand matching. Working paper, University of Toronto, Toronto.
- Hu B, Hu M, Zhu H (2018) Surge pricing and two-sided temporal responses in ride-hailing. Working paper, University of Toronto, Toronto.
- Ibrahim R (2018) Managing queueing systems where capacity is random and customers are impatient. *Production Oper. Management* 27(2):234–250.
- Jiang B, Tian L (2018) Collaborative consumption: Strategic and economic implications of product sharing. *Management Sci.* 64(3): 1171–1188.
- Kabra A, Belavina E, Girotra K (2018) Bike-share systems: Accessibility and availability. Working paper, University of Maryland, College Park, College Park.
- Kendall DG (1951) Some problems in the theory of queues. *J. Royal Statist. Soc. Ser. B* 13(2):151–185.
- Lu M, Chen Z, Shen S (2017) Optimizing the profitability and quality of service in carshare systems under demand uncertainty. *Manufacturing Service Oper. Management* 20(2): 162–180.
- Ming L, Tunca TI, Xu Y, Zhu W (2019) An empirical analysis of market formation, pricing, and revenue sharing in ride-hailing services. Working paper, Chinese University of Hong Kong, Shenzhen, Guangdong, China.
- Nikzad A (2017) Thickness and competition in ride-sharing markets. Working paper, Stanford University, Stanford, CA.
- Ozkan E, Ward A (2019) Dynamic matching for real-time ridesharing. *Stochastic Systems*. Forthcoming.
- Rochet JC, Tirole J (2003) Platform competition in two-sided markets. *J. Eur. Econom. Assoc.* 1(4):990–1029.
- Taylor T (2018) On-demand service platforms. *Manufacturing Service Oper. Management* 20(4):704–720.
- Tian L, Jiang B (2018) Effects of consumer-to-consumer product sharing on distribution channel. *Production Oper. Management* 27(2):350–367.
- Yu JJ, Tang CS, Shen Z-JM, Chen X (2017) Should on-demand ride services be regulated? An analysis of Chinese government policies. Working paper, Tsinghua University, Beijing.
- Zipkin P (2008) On the structure of lost-sales inventory models. *Oper. Res.* 56(4):937–944.