POMS PRODUCTION AND OPERATIONS MANAGEMENT

PRODUCTION AND OPERATIONS MANAGEMENT POMS

Vol. 0, No. 0, xxxx–xxxx 2020, pp. 1–18
ISSN 1059-1478 | EISSN 1937-5956 | 20 | 00 | 0001

DOI 10.1111/poms.13330

From the Classics to New Tunes: A Neoclassical View on Sharing Economy and Innovative Marketplaces

Ming Hu*

Rotman School of Management, University of Toronto, Toronto, Ontario M5S 3E6, Canada, ming.hu@rotman.utoronto.ca

Operations management has the tradition of coming from and going back to real-life applications. It deals with the management of the process of matching supply with demand. The emerging business process in a sharing economy or an innovative marketplace calls for active management at the operational level. We take a neoclassical perspective by drawing inspiration from the classic models in operations management and economics. We aim at building connections and identifying differences between those traditional models and the new applications in sharing economy and innovative marketplaces. We also point out potential future research directions.

Key words: sharing economy; innovative marketplace
History: Revised: August 2020, Accepted: November 2020 by J. George Shanthikumar, after one revision.

1. Introduction

Operations management deals with managing the process of matching supply with demand. The discipline has the tradition of coming from and going back to real-life applications. Each one of the founding pillars, such as inventory management and revenue management, establishes itself as a field not only because it has beautiful theories but also because the theories were motivated by practice and can be readily implemented. On the one hand, the widely practiced classic inventory theory typically treats demand as given and focuses on minimizing the cost associated with inventory replenishment on the supply side. On the other hand, the traditional revenue management models, initially practiced by the airlines and hotels, assume a fixed supply side and focus on maximizing the revenue of selling a limited amount of capacity by regulating the demand.

Nowadays, facilitated by the proliferating adoption of Internet-connected sensors and devices, the age-old idea of resource sharing (e.g., library book sharing) is reviving and developing into a crop of innovative business models, which are often referred to as sharing economy and innovative marketplaces. The notion, in general, may refer to a market model that allows sharing of access to goods and services, or an online platform that enables individuals or small entities as buyers and sellers to "transact." The promise of this emerging industry lies in increasing the utilization of resources/assets and improving the efficiency of transactions. These business models require active management of regulating supply and demand at the same time, taking into account the incentives and decisions of the sellers/service providers and buyers/customers. For example, Uber, as an intermediary platform, crowdsources services from independent drivers to fulfill trip requests by riders. The platform determines wage for drivers and price for riders, and dispatches a driver to serve a rider after they find it incentive compatible to enter the matching pool. Unlike Uber, Airbnb is an online marketplace in which homeowners take the wheel by charging prices for their short-term rentals to guests. The platform, as a broker, can influence the matching between hosts and guests but cannot prescribe hosts' decisions.

What does this article do and not do?

1. This article takes a neoclassical perspective by drawing inspiration from the classic models in operations management. It aims at building bridges between the traditional models and new areas, which allow us to apply tools developed for the old to the new. The models presented here are parsimonious and not the closest to reality. But they are built upon the classic operations/economics models and tuned for those innovative applications. They can be a starting point to get one step closer to reality.
2. This article serves as a tutorial with basics, which one may use to start exploring this exciting field. It also points out some future research opportunities (see also Benjaafar and Hu 2020, Chen et al. 2020c) As our main purpose is not to comprehensively survey the area which is rapidly growing (see Hu 2019 for a collection of earlier works), not all of the related papers are covered.

Please Cite this article in press as: Hu, M. From the Classics to New Tunes: A Neoclassical View on Sharing Economy and Innovative Marketplaces. *Production and Operations Management* (2020), https://doi.org/10.1111/poms.13330

3. In this article, we will use Uber and Airbnb, two arguably most successful examples of sharing economy and innovative marketplaces, to motivate a set of prototypical models. There are a variety of unique aspects associated with other innovative business models and applications, such as capacity planning and rebalancing in bike sharing (Freund et al. 2019) and service zone design in electric vehicle sharing (He et al. 2017), which we do not cover here.
4. This article focuses on analytical models to be aligned with the special issue's theme. There is a growing literature of empirical, experimental, and behavioral studies, see, for example, Cohen et al. (2020), Allon et al. (2018), Ming et al. (2020),and Jiang et al. (2020) on the ride-sharing market such as Uber and Lyft, Li et al. (2019), Cui et al. (2020a), and Cui et al. (2020b) on the short-rental market such as Airbnb, Kabra et al. (2020) on bike sharing and Bimpikis et al. (2020) on an online marketplace.

2. From Inventory Theory to Dynamic Matching

Motivated by Uber's matching of drivers and riders, we study a parsimonious model of matching which connects back to the classic inventory theory. In each period, supply and demand of various "types" arrive in random quantities. In a sharing economy, as supply is crowdsourced, there exists uncertainty on the supply side in the same fashion as the demand side. In ride-hailing, types can be geolocations, as drivers and riders view the transportation as more of a homogeneous service. Then, the matching between a pair spatially closer to each other generates a higher reward because of a shorter pickup time for the driver and shorter waiting time for the rider. For parsimony, we assume away pricing decisions and focus on the centralized matching decisions. With unmatched drivers and riders fully or partially stay in the system, the platform aims at maximizing the total expected rewards by optimizing the matchings. This problem lies in the center of many sharing economy platforms, which often use crowdsourced supply and match it dynamically with customer demand in a centralized fashion, as supply and demand randomly arrive at the platform.

To stay close to the classic periodic-review inventory system, we assume there is a finite horizon of T periods. At the beginning of each period t, m types of supply and n types of demand arrive in *random* quantities, which could follow a Bernoulli distribution (taking only a value of 0 or 1) as a special case. We use i to index a supply type and j to index a demand type. There is an exogenously given reward r_{ij}^t for matching one unit of type i supply and one unit of type j demand in period t. For example, r_{ij}^t may be a constant reward minus the traveling cost from location i to location j in period t. We can write the rewards in a matrix form as $\mathbf{R}^t=(r_{ij}^t)$. To obtain a parsimonious formulation, one can account for waiting costs of those supply and demand types that are not immediately matched by incorporating those costs into the matching rewards. The state for a given period t constitutes the supply and demand levels of various types before matching but after the arrival of random supply $\mathbf{S}^t$ and demand $\mathbf{D}^t$ of various types for that period, with $\mathbf{S}^t$ and $\mathbf{D}^t$ of a vector in the dimension of m and n, respectively. We denote, as the system state, the supply vector by $\mathbf{x}=(x_1,\ldots,x_m)$ and the demand vector by $\mathbf{y}=(y_1,\ldots,y_n)$, where x_i and y_j are the quantities of type i supply and type j demand available to be matched. On observing the state $(\mathbf{x},\mathbf{y})$, the firm decides on the quantity q_{ij} of type i supply to be matched with type j demand. For conciseness, we write the decision variables of matching quantities in a matrix form as $\mathbf{Q}=(q_{ij})$.

Given a feasible matching quantity matrix $\mathbf{Q}$, the post-matching levels of type i supply and type j demand are given by $u_i=x_i-\sum_{j'=1}^{m}q_{ij'}$ and $v_j=y_j-\sum_{i'=1}^{n}q_{i'j}$, respectively. The post-matching supply and demand vector is denoted by $\mathbf{u}$ and $\mathbf{v}$, respectively. The unmatched supply and demand at the end of a period carry over to the next period with a fraction of $\alpha\in[0,1]$ and $\beta\in[0,1]$, respectively. The carry-over rates can be made to be time dependent. The platform's goal is to determine a matching policy $\mathbf{Q}^*=(q_{ij}^*)$ that maximizes the expected total matching rewards over the finite horizon. Denote by $V_t(\mathbf{x},\mathbf{y})$ the optimal expected total rewards given that it is in period t and the current state is $(\mathbf{x},\mathbf{y})$. We formulate the following stochastic dynamic program:

$$V_t(\mathbf{x},\mathbf{y})=\max_{\mathbf{Q}\in\{\mathbf{Q}\geq\mathbf{0}|\mathbf{u}\geq\mathbf{0},\mathbf{v}\geq\mathbf{0}\}}\mathbf{R}^t\circ\mathbf{Q}+\gamma EV_{t+1}(\alpha\mathbf{u}+\mathbf{S}^t,\beta\mathbf{v}+\mathbf{D}^t), \tag{DM}$$

where "$\circ$" gives the sum of elements of the entry-wise product of two matrices and $\gamma\in[0,1]$ is a discount factor. The boundary conditions can be $V_{T+1}(\mathbf{x},\mathbf{y})=0$ for all $(\mathbf{x},\mathbf{y})$, without loss of generality.

Under some conditions, structural properties of the optimal matching policies for Problem (DM) can be derived. Hu and Zhou (2020) establish the so-called "modified Monge condition" that specifies a dominance relation between two pairs of supply and demand types. The modified Monge conditions are sufficient, and necessary in a robust sense, for the optimal matching policy to satisfy the following priority properties in Problem (DM): For any two pairs of supply and demand types with one strictly

Please Cite this article in press as: Hu, M. From the Classics to New Tunes: A Neoclassical View on Sharing Economy and Innovative Marketplaces. *Production and Operations Management* (2020), https://doi.org/10.1111/poms.13330

dominating the other, it is optimal to prioritize the matching of the dominating pair over the dominated pair. The modified Monge condition generalizes the condition of a Monge sequence, discovered by Gaspard Monge, a French mathematician, in 1781, which guarantees a static and balanced transportation problem to be solved by a greedy algorithm (see Table 1 for detailed comparisons). As a result of the priority properties, the optimal matching policy boils down to a match-down-to structure when a specific pair of supply and demand types is considered, along with the priority hierarchy. That is, 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 as much as possible, *all* pairs that are strictly dominated by this pair should not be matched at all, due to the priority structure.

This structural property of "priority and thresholds" is a generalization of priority structures seen in the balanced and deterministic transportation problems, and the threshold-type policies seen in the inventory management (such as base-stock levels) and quantity-based revenue management (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 (2020) further show, by verifying the $L^{\natural}$-concavity of the value functions of a transformed problem, the optimal total matching quantity or the optimal match-down-to levels can have monotonicity properties with respect to the system state. This technique has been applied for deriving structural properties for lost-sales inventory models (Zipkin 2008) and perishable-inventory models (Chen et al. 2014). For another example, in view of the difference between backlogs vs. lost-sales inventory systems, Chen et al. (2019c) focus on a lost-sales analogy of the above dynamic matching problem that allows backlogs. The authors show that under a supermodular reward structure, the optimal dynamic matching policy can still have the structure of priority and thresholds for both backlogs and lost-sales systems, but the specific priority may be entirely different in the two systems.

Table 1 Comparisons Between the Monge Sequence and Modified Monge Condition

Monge sequence	Modified monge condition
Static, deterministic, and balanced	Dynamic, stochastic, and unbalanced
Transportation problem	Matching problem
Defined for a sequence of all pairs	Defined for specific pairs
Sufficient and necessary	Sufficient, and robustly necessary
Greedy algorithm:	Priority and thresholds:
(1) Priority sequence	(1) Priority hierarchy
(2) Match as much as possible	(2) Match-down-to policy

Research opportunities.

1. Competitive analysis. In the above formulation, supply and demand have known arrival probabilities. The platform may want to solve a robust optimization version of the problem with unknown arrival patterns. The objective then becomes to design online algorithms (of making decisions on the fly) to protect the firm against the worst-case scenario and have theoretical performance guarantees over all possible instances (see, e.g., Karp et al. 1990 for an early work and Ma and Simchi-Levi 2020, Truong and Wang 2019 for the latest developments).
2. Computational approach to stochastic dynamic programming. The structure of priority and thresholds holds for certain reward structures. But even for those cases, computing the thresholds can be cumbersome. It is desirable to have a general computational approach that can work for any reward structure. Given the success of applying approximate dynamic programming and Lagrangian relaxation approaches to dynamic pricing problems in giving rise to structural properties and efficient algorithms (see, e.g., Adelman 2007, Balseiro et al. 2020) it leaves a lot to be desired for applying similar approaches to the dynamic matching problem.
3. Centralized vs. decentralized dynamic matching. In some marketplace, for example, the online labor market, sequentially arriving participants on both supply and demand sides make self-interested (i.e., decentralized) matching decisions. It can be essential for the platform that operates such a marketplace to quantify the gap between centralized vs. decentralized matching in a dynamic setting and identify ways to close the gap (see, e.g., Baccara et al. 2020).

3. From One-Sided to Two-Sided Pricing

First, consider the most parsimonious one-sided pricing problem. A seller faces a market of customers with size M and random customer valuation V. On the supply side, the marginal cost of the product is a constant of w. The seller decides on the price p to solve the one-sided pricing problem:

Please Cite this article in press as: Hu, M. From the Classics to New Tunes: A Neoclassical View on Sharing Economy and Innovative Marketplaces. *Production and Operations Management* (2020), https://doi.org/10.1111/poms.13330

$\max_p (p-w)d(p)$, where $d(p)=M \cdot P(V \geq p)$ is the downward-sloping demand curve. This formulation serves as the foundation of classic finite-horizon revenue management problems, in which the marginal cost dynamically changes, contingent on the seller's remaining capacity and the time until the end of the sales horizon.

Now we extend this one-sided pricing problem to a two-sided one. As the purpose is to reveal connections among problems, we slightly abuse the notation and it would be clear from the context that the same notation may refer to different variables in different problems. Suppose an intermediary platform crowdsources a homogeneous good or service from a pool of independent sellers or contractors with size N and random opportunity cost C and sell it to a market of buyers with size M and random customer valuation V. See Figure 1 for an illustration of such a market. If there is unlimited supply, the total amount of customers who are willing to pay for the service at price p is $d(p)=M \cdot P(V \geq p)$. Given a posted wage w, the total amount of independent suppliers or contractors, who are willing to provide the good or service, is $s(w)=N \cdot P(C \leq w)$. This is the number of suppliers who would show up if they were guaranteed to be matched with a customer. The platform decides on the wage w offered to the sellers and the price p charged to customers to solve the following two-sided pricing problem:

$$\max_{w,p} (p-w) \min\{s(w), d(p)\}. \tag{P}$$

One can show that for a given wage w and price p, even if the suppliers and customers strategically anticipate their matching likelihood, the matching quantity is still equal to $\min\{s(w),d(p)\}$.

Problem (P) is different from the capacitated pricing problem $\max_p (p-w)\min\{Q, d(p)\}$, to which the deterministic version of the classic revenue management problem boils down (see Gallego and Van Ryzin 1994). The difference comes from in the capacitated pricing problem, the capacity Q and the marginal operating cost w are independent and exogenously given, while they are linked and endogenized in Problem (P).

In economics, research on two-sided markets has studied platforms such as credit cards, video game consoles, and organ allocation/exchange. Rochet and Tirole (2003) consider a general model of competition between two platforms with the transaction volume in the *multiplicative* form of supply and demand. Such a form of the transaction volume is appropriate for a two-sided market platform with a *long-term* objective. For instance, the credit card company cares about the potential transaction volume in proportion to $s(w) \cdot d(p)$ if there are $s(w)$ merchants on the supply side accepting the credit card for payment and $d(p)$ customers on the demand side using the credit card. Motivated by the ride-hailing market, Problem (P) studies an on-demand matching platform's minute-by-minute pricing decisions that adapt to the changing market conditions. As a result, the transaction volume takes the form being *minimum* of supply and demand quantities. To see this, if in a *short run*, there are a number of $s(w)$ drivers and $d(p)$ riders within a neighborhood, the transaction volume is close to $\min\{s(w),d(p)\}$. In view of the classic newsvendor problem with pre-committed supply, taking the minimum of supply and demand is the most natural form from an operational perspective.

The transaction volume taking the form of $\min\{s(w),d(p)\}$ results in a pricing problem of a totally different nature from the one-sided pricing problem $\max_p (p-w)d(p)$, or the problem $\max_{w,p} (p-w)s(w)d(p)$ with a long-term objective. To see this, we illustrate below how the optimal price changes as a function of exogenously given wage. Under very mild regularity conditions, Hu and Zhou (2018) show that the optimal price $p^*(w) = \arg\max_p (p-w)\min\{s(w), d(p)\}$, as a function of exogenously given wage w, is U-shaped. That is, it

Figure 1 Two-Sided Market

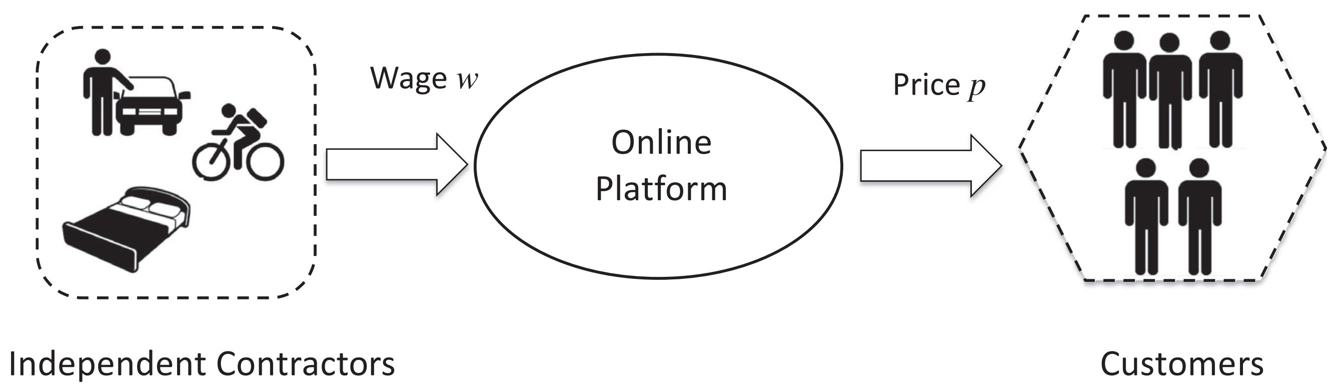

Please Cite this article in press as: Hu, M. From the Classics to New Tunes: A Neoclassical View on Sharing Economy and Innovative Marketplaces. *Production and Operations Management* (2020), https://doi.org/10.1111/poms.13330

is first decreasing in w and then, increasing in w. See Figure 2 for an illustration.

The intuition of the U-shaped optimal price function $p^*(w)$ is as follows. When the wage is meager, say, close to zero, just a few drivers are willing to come out and work. The platform sets a price that is exorbitantly high so that only those riders who have very high valuations get a ride. When the exogenous wage increases from being very low, more drivers come out to work, and the platform lowers the price so that more riders get the service to catch up with the increasing amount of available drivers. When the exogenous wage is high enough, there are too many drivers willing to come out and expect to get some work. If the platform wants to give every driver a job, the price charged to a rider would be too low, even possibly lower than the exogenous wage. In this case, the platform is better off to charge the optimal one-sided price with the wage as being given. In other words, when the exogenous wage w is below a threshold $w^* \leq \bar{w}$, the supply is limited due to the relatively low wage, and thus the optimal price is the *market clearing price*, that is, the price such that $d(p)=s(w)$, which would decrease as the wage rises. When the exogenous wage w is above the threshold, the supply is ample due to the relatively high wage, and thus the optimal price is the unconstrained *revenue maximizing price*$\arg\max_p (p-w)d(p)$, which would increase as the wage rises.

This U-shape property of $p^*(w)$ is in stark contrast with the traditional supply chain settings. Consider a supply chain where a retailer procures from its supplier who may, in turn, procure from further upstream suppliers. Any cost surge in the supply chain, for example, a wage or cost increase at some firm along the supply chain, would always push up the

Figure 2 U-Shaped Optimal Price as a Function of Exogenous Wage [Color figure can be viewed at wileyonlinelibrary.com]

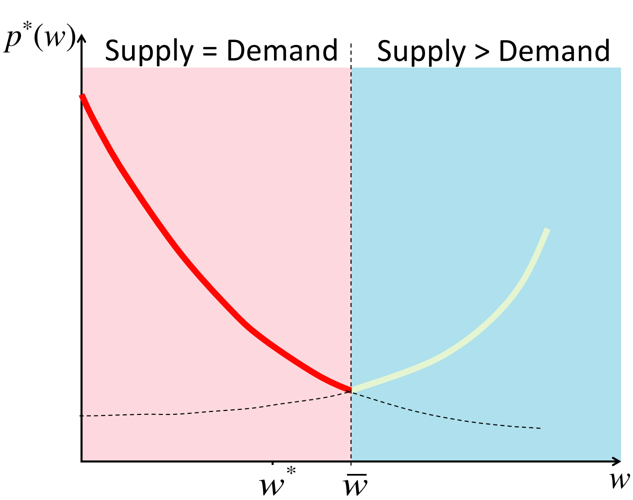

wholesale price w and lead to an increase in the optimal retail price $\arg\max_p (p-w)d(p)$. Moreover, the U-shape property of $p^*(w)$ is also in stark contrast to the classic economics literature on two-sided pricing. As mentioned, Rochet and Tirole (2003) assume a multiplicative form of the transaction volume. That is, the platform solves the problem $\max_{w,p}(p-w)s(w)d(p)$. It is easy to see that the objective function $(p-w)s(w)d(p)$ is log-supermodular in w and p. Then by Topkis's theorem, the optimal price $\arg\max_p (p-w)s(w)d(p)$ is always increasing in w.

That range in which the optimal price may decrease in the exogenous wage is certainly relevant. Note that the optimal wage and price $(w^*,p^*)=\arg\max_{w,p}(p-w)\min\{s(w),d(p)\}$ would satisfy $w^* \leq \bar{w}$, because $s(w^*)=d(p^*)$ must hold as there is no uncertainty and the platform can fully control wage and price.

The above discussion reveals that Problem (P) with the short-term objective and the transaction volume taking the form of the minimum of supply and demand is natural for operations management researchers but is fundamentally different from the traditional unconstrained supply chain setting (i.e., the one-sided pricing) and the two-sided pricing in the economics literature. It provides us with a justification for theoretical novelty in examining the two-sided pricing problem at the operational level.

The parsimonious form of Problem (P) can serve as a workhorse model on top of which many additional features can be incorporated. For example, Hu and Zhou (2018) focus on the widely practiced, flat, across-the-board commission contracts, under which the platform takes a fixed cut and the wage is equal to a fraction of the price, regardless of what price is charged, that is, $w=\gamma p$, where γ is the payout rate to the drivers and $(1-\gamma)$ is the commission rate the platform takes. The platform precommits to a commission rate before it sets a contingent price depending on the randomly realized market conditions. The authors study the performance of an optimal commission rate contract benchmarked with the case that the platform can freely choose wage and price for any market condition without any constraint.

Research opportunities.

1. Dynamic two-sided pricing. In practice, there exists uncertainty in the timing, number, and valuation of arrivals of sellers/contractors and buyers. The platform can dynamically vary wage and price to balance supply and demand. Unmatched sellers and buyers may hang around in the market waiting to be matched. Contingent pricing, reacting to the market condition, obviously can have an edge over a fixed price. Cachon et al. (2017) show that drivers

and riders are generally better off with varying prices contingent on the realized supply and demand conditions. Chen et al. (2020b) consider a dynamic version of Problem (P) allowing the intermediary to buy in, sell out, and hold inventory. Moreover, static pricing may also be preferred. Banerjee et al. (2015) show that a static pricing policy can be asymptotically optimal in a thick market. Chen and Hu (2020b) further study a dynamic version of Problem (P) in which sellers/contractors and buyers sequentially arrive at 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 hard 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. There are still many unanswered questions. For example, it is interesting to characterize the structural property of the optimal dynamic two-sided pricing policy with and without strategic behavior of sellers and buyers.

2. Spatial two-sided pricing. Problem (P) is a single-location model, assuming all supply and demand are located nearby in a small region. If the spatial dimension of ride-hailing is considered, it is critical to take into account drivers' incentives of moving around which involve the time and effort in repositioning themselves, picking up riders, and going to the riders' destination where it may be hard to find a good ride. Bimpikis et al. (2019) explore the spatial price equilibrium, in a multilocation model allowing drivers to decide whether to work and where to position themselves in a network of interconnected locations. The authors highlight the impact of the origin-destination demand (im)balancedness across the network of locations on the platform's wage, price, profits, and consumer surpluses. Besbes et al. (2019) consider a continuously dispersed linear city where the drivers can reposition themselves. Under a fixed commission rate, the platform sets location-specific prices along the linear city, and then, the riders' requests along the city realize. The drivers then relocate themselves in a simultaneous-move game based on prices, demands, and driving costs. Garg and Nazerzadeh (2020) show that drivers would cherry-pick trip requests under the multiplicative surge (i.e., the payout scales with the trip length), and the issue can be addressed by the additive surge (i.e., a surge component in the payout independent of the trip length). Spatial pricing is still an underexplored area with many opportunities. For example, it is highly desirable to consider a stochastic dynamic spatio-temporal model that allows the travel time to be proportional to the travel distance between locations (see Ata et al. 2019 for an econometric study).
3. Joint pricing and matching. The marriage of supply chain and revenue management gives birth to the stream of research on joint pricing and inventory control (see, e.g., Chen and Simchi-Levi 2004, Federgruen and Heching 1999). It will be a fruitful direction to consider the pricing and matching decisions jointly. On the one hand, the platform can use wage and price to regulate supply and demand of different types, for example, at different geolocations, and on the other hand, the platform then can decide on how to match those supply and demand that accept the given wage and price. The coupling of pricing and matching makes the problem challenging yet fascinating.

4. From Queueing to Resource Sharing

In the previous two sections, we draw connections between sharing economy and the classic inventory (supply) and revenue (demand) management. Another critical methodology in operations management is queueing theory that can be applied to study service systems. One unit of resources can be shared among randomly arriving customers for a random amount of time. It is not a new idea to view the resources as the servers and apply a queueing model to capture the dynamics of supply–demand mismatch at the operational level, for example, a customer who does not find an available resource upon arrival needs to wait (see Cachon and Feldman 2011 for Netflix DVD sharing), or leaves the system without getting served (which is referred to as a "loss" system). What can be novel is to consider crowdsourced servers' incentives of participation, which determine the supply quantity, and in turn, interact with demand, potentially under the moderation by a platform. For example, Benjaafar et al. (2019) apply a multi-server loss queueing system to study peer-to-peer product sharing.

Pricing control. As one type of resource sharing, ride-hailing can naturally be viewed from a queueing perspective. For example, one can revisit the two-sided pricing problem (1) by adopting a queueing formulation. The benefit of doing that is to capture the experienced delay by riders (demand) or drivers (supply), a critical feature at the operational level that is

Please Cite this article in press as: Hu, M. From the Classics to New Tunes: A Neoclassical View on Sharing Economy and Innovative Marketplaces. *Production and Operations Management* (2020), https://doi.org/10.1111/poms.13330

absent in the simple formulation (P). Now consider the following model. Suppose the riders are delay sensitive with a waiting cost c per unit of time. They make joining or balking decisions based on the expected wait time $E(W)$ before getting served. This is analogous to the treatment in the unobservable queue, and an alternative treatment can be similar to the observable queue (see Hassin and Haviv 2003, Chap. 2 and 3). The effective arrival rate of customers requesting the service should satisfy:

$$\lambda = \Lambda \cdot P(V - p - cE(W) \geq 0), \tag{1}$$

where Λ is the arrival rate of potential riders. Suppose each rider takes $1/\mu$ amount of time to serve, that is, μ is the service rate. Given there are $s(w)$ number of drivers (servers) available on the street, Taylor (2018) and Bai et al. (2019) adopt the M/M/k queue formula as a natural approximation to express the expected wait time $E(W)$ experienced by riders in terms of the arrival rate λ, the service rate μ and the number of servers k=$s(w)$, which together with Eq. (1) can capture the equilibrium outcome. Here, w should be understood as the expected payoff of drivers from signing up and joining the workforce. Another coarser approximation is to use the M/M/1 queue formula to express the expected riders' delay $E(W)$ with arrival rate λ and service rate $s(w)\mu$, see, for example, Benjaafar et al. (2020) and Guo et al. (2020), which then lends tractability to explore other phenomena under the approximation. In equilibrium, the effective rate of riders requesting services can be expressed in terms of wage and price, denoted by $\lambda^*(w,p)$, and the number of drivers who find it incentive compatible to come out and work can also be expressed in terms of wage and price, denoted by $s^*(w,p)$. Analogous to Problem (P), the platform's problem is to solve $\max_{w,p}(p-w)\lambda^*(w,p)$, subject to $s^*(w,p)\mu < \lambda^*(w,p)$ which is the stability condition. The underlying assumption of this parsimonious M/M/k formulation is that there are a fixed number of drivers in the system of a single location.

In contrast, Banerjee et al. (2015) study an "open" queueing network with two queues (where the "open" network refers to that drivers can come to and exit from the system), an M/M/1 queue to model idle drivers at one location and an M/M/∞ to model busy drivers within the region (see also He et al. 2017 in which the same modeling approach is adopted for electric vehicle sharing). The authors focus on comparing static pricing and state-dependent contingent pricing policies.

Routing and other controls. All queueing models mentioned above are single-location models. To account for cars' spatial movement and the platform's routing controls, there is a stream of research adopting the framework of a "closed" loss queueing network. Specifically, there is a finite number of locations/regions and a finite number of drivers/cars (which can be controlled by wage or routing decisions). Each region has its own stream of arriving riders. When a rider arrives at a location, if there is an idle driver, the rider and driver can be matched and travel together to another location with a given transition probability and a certain travel time. The fixed number of cars move between locations but never leave the system (referred to as the "closed network"), and riders get lost if finding no idle driver in the desired location (referred to as the "loss system"). In this stream, Afèche et al. (2018a) consider two locations and focus on the performance impact of repositioning control (or drivers' self-repositioning) and demand-side admission control (with a fixed price p). Braverman et al. (2019) consider multiple locations and focus on empty-car routing control. Chu et al. (2018) study a single-location model in which drivers can cherry-pick riders, and focus on information, routing, and priority controls by the platform.

Moreover, Özkan and Ward (2020) consider an "open" loss queueing network in which the drivers' arrival and departure processes at each location are exogenous, and focus on the dynamic matching decision. In a general setting, Gurvich and Ward (2014) study the dynamic control of matching queues at the operational level. All these routing papers focus on the fluid deterministic counterpart of the original stochastic system, due to the apparent complexity of the problem.

Self-scheduling servers. One feature of a sharing economy is supply uncertainty because resources of goods and services are crowdsourced. Classical queueing models assume a fixed number of servers. This may reflect the reality that even though drivers have the freedom of dictating when and how long to work, as soon as they start working, the setup cost keeps them to stay around not just for one ride/task but for many. As a result, at a more operational level, the number of available drivers may be somewhat constant over certain time window. But ideally one wants to build queueing models that can handle random service capacity which results from self-scheduling. For example, Ibrahim (2018) studies the controls of staffing, compensation, and delay announcements, to effectively control a queueing system with a random number of servers and impatient customers.

Research opportunities. There is still a lot to be done in the directions mentioned above. For example, most of those above multilocation queueing models assume loss demand, which is desirable to be extended. Moreover, here are some broad directions and problems to explore:

Please Cite this article in press as: Hu, M. From the Classics to New Tunes: A Neoclassical View on Sharing Economy and Innovative Marketplaces. *Production and Operations Management* (2020), https://doi.org/10.1111/poms.13330

1. Double-ended queue. As an alternative queueing model to the M/M/k queue, the double-ended queue has been used to model a single-location taxi stand, at which drivers and riders arrive independently, see Kendall (1951). After a driver and a rider pair up, they leave the system. The double-ended queue can be a more appropriate modeling framework for a matching market where service providers do not come back after a matching, for example, shoppers at a grocery store help make a delivery for their neighbors, and commuters pool their car with riders on the way to work or home. This double-ended queue framework is a continuous-time version of the dynamic matching model introduced in section 2 with one supply type and one demand type, and may be readily extended to account for two-sided pricing or/and multiple locations (types).
2. Carpooling. When a car is shared among riders, the capacity of the "server" (i.e., the car) is literally increased. However, such a capacity increase is endogenized by the riders' choices. It would be interesting to extend some of the above queueing models to account for the option of carpooling by riders, see, for example, Jacob and Roet-Green (2019), Hu et al. (2020b).
3. Free floating system. Those multilocation spatial models mentioned above are one step closer to reality. However, they still do not fully capture the reality of free-floating cars. The gap is analogous to the difference between docked vs. dockless bike sharing. A spatial model with finer granularity is desired and can fit data better. In practice, from the drivers' point of view, a renewal "cycle" starts with dropping off the previous rider, and follows with waiting for a rider request, then on the way to pick up the rider after receiving a dispatch order and then, driving to the destination with the rider in the car. From the riders' point of view, the wait not only includes the amount of time earlier riders still occupy the drivers but also consists of the pickup time for the rider and some or all of the previous riders. The pickup time, proportional to the travel distance, makes the dynamics complicated and hard to analyze because it depends on the number of available drivers when a rider request arrives at the system and also on the routing policy (see, e.g., Feng et al. 2020). The system dynamics can be similar to the classic stochastic and dynamic vehicle-routing problem but not the same. Bertsimas and Van Ryzin (1991, 1993) provide potential directions on how to better understand the system dynamics. As the latest development in this direction, Besbes et al. (2018) modify an M/M/k queue with a state-dependent service rate that takes into account the pickup time under the match-to-the-closest dispatch rule. On top of the performance evaluation for a given number of drivers and a specific routing policy (see, e.g., Hu 2020), one can optimize the routing policy or study the optimal one- or two-sided pricing problem (see, e.g., Chen and Hu 2020a).
4. Three-sided matching. In grocery/food delivery, platforms use technology to stitch together couriers (drivers), grocery stores/restaurants, and customers. The problem is more complicated than ride-hailing in which only two parties are involved. Other than more parties involved, the complexity may come from: When couriers arrive at a designated place to pick up an order, the order may not be ready, as there is a separate process of preparing it for pickup; Couriers may pickup multiple orders at the same or different places and send them to the same or different destinations (see, e.g., Gorbushin et al. 2020); Other than pricing and routing decisions, the platform can optimize the product offerings by highlighting various options to moderate demand.

5. From Newsvendor to Resource Sharing

The above-mentioned queueing formulations have the advantage of capturing asymmetric decision timescales of the drivers and riders in ride-hailing. More specifically, drivers tend to evaluate more long-term payoffs when deciding on whether to provide the service while riders make a purchase decision in a much shorter timescale. The fixed-server or closed-network queueing formulation assumes a fixed number/amount of drivers or service capacity to serve riders who arrive at a transient system state.

An alternative formulation with fixed capacity can be analogous to the classic newsvendor problem. Cachon et al. (2017) consider a strategic-level decision period in which drivers decide whether to sign up for a ride-hailing platform taking into account future expected payoffs and the platform can impose a cap on the driver pool. Then in the short run, for example, at a day-to-day level, drivers decide on whether to come out and work depending on their realized opportunity cost for that period, and the platform may set wage or price contingently. The problem is somewhat analogous to a newsvendor problem, because of the endogenized driver pool size. In addition to the cap on the driver pool at the strategic level, Gurvich et al. (2019) also allow the platform to set a cap on the number of drivers who can work, in addition to wage or price decisions, in the short run. Here, the short-run cap imposed by the firm is literally a newsvendor quantity.

Similarly, one can build a simple newsvendor-type extension of Problem (P). At the strategic level, the drivers anticipate the expected earning $\bar{w}$, and thus the number of drivers who sign up for work is $s(\bar{w})$. In the short run, given the fixed number $s(\bar{w})$ of drivers at work which is analogous to the newsvendor quantity, the platform can set a price contingently, in view of the surge pricing practice. That is, at the operational level, the platform is to solve for the contingent wage w_ε^* and price p_ε^* such that they achieve

$$\max_{w,p}(p-w)\min\{s(\bar{w}),d_\varepsilon(p)\}, \qquad \text{(N)}$$

where $d_\varepsilon(p)$is a realization of a random demand curve $D(p)$ which can be obtained by perturbing the deterministic demand curve $d(p)$ additively or multiplicatively with a random variable $\mathscr{E}$ (whose realization is denoted by ϵ). One could impose constraints on the relationship between p and w in Problem (N) such as w=γp where γ is a fixed payout rate. Given the contingent wage w_ε^*, the fraction of drivers who get a job, out of those making themselves available, is $s(w_\varepsilon^*)/s(\bar{w})$. To make sure the drivers earn what is promised on expectation, as a constraint,

$$\bar{w}=E\left(w_{\mathscr{E}}^*\cdot\frac{s(w_{\mathscr{E}}^*)}{s(\bar{w})}\right).$$

If the platform sets a committed price p at the strategic level, that is, $p_\varepsilon^*=p$ does not depend on the random demand realization ϵ, the problem is then close to the classic newsvendor selling to price-sensitive customers (see Petruzzi and Dada 1999). The difference, however, is that here the "newsvendor" in sharing economy achieves its "newsvendor" quantity $s(\bar{w})$ through announcing its committed pricing decision p, where $\bar{w}$ is the expected effective earning of drivers under price p.

Research opportunities.

1. Pricing and commission policies. Among others, the newsvendor-type formulations provide opportunities to study various pricing strategy in the short run such as contingent vs. committed pricing strategy, and crowdsourcing contracts at the strategic level such as contingent commission rate vs. committed minimum wage contracts.
2. The employee mode. One can consider a variant of the classic newsvendor selling to price-sensitive customers with the price contingently determined on the uncertainty realization and the "newsvendor" quantity is obtained as a function of effective earnings. In this model, the drivers are hired as employees and made incentive compatible to sign up for the job in anticipation of expected earnings. Such an employee mode can become more practically relevant than the independent contractor mode (see Problem (P)), as California Assembly Bill 5, in effect starting Jan 1, 2020, compels gig-economy platforms to treat their worker as employees as opposed to independent contractors. It can be of significant practical value to quantify the impact on the stakeholders such as the drivers, riders, and the platform, when the system switches from the independent contractor mode of Problem (P) to the employee mode analogous to Problem (N).

6. From Strategic Customers to Strategic Buyers and Sellers

In all formulations mentioned above, when customers make a purchase decision, they compare available options at the moment upon their arrival, that is, their decision making is myopic. In the operations management literature, a large body of literature studies so-called strategic, or more accurately, forward-looking, customer behavior, in which customers consider purchase opportunities over time (see, e.g., Shen and Su 2007). This literature is motivated by potential strategic waiting behavior by customers in anticipate of a future price drop. In ride-sharing and other marketplaces, not only buyers but also sellers may demonstrate strategic behavior. Moreover, they may be not only inter-temporally strategic but also spatially strategic.

First, let us consider strategic customer behavior on the demand side. Hu et al. (2020a) consider a two-period model, commonly adopted in the strategic customer literature, to capture riders' strategic waiting in response to surge pricing in ride-hailing. At the beginning of period 1, a total volume of N=1 (infinitesimal) drivers are available in a small region, and a total volume of M>1 (infinitesimal) riders appear in the same region, constituting a demand surge. The riders' incremental valuations of getting rides from the platform, net of their next best options are heterogeneous following a distribution V. Each unit of rider requires one unit of the drivers to provide service. Once matched, both rider and driver leave the region indefinitely.

In period 1, the platform sets price p_1 for that period. Given p_1, some riders will request rides in this period, some may strategically wait out period 1 and request rides only in period 2, and the rest never request rides. The riders requesting rides in period 1 and available drivers are randomly matched. In period 2, new drivers with an opportunity cost c who were nearby and strategically decided in period 1 to "chase the surge" arrive in the surge region to join any remaining drivers, and the platform sets p_2. The remaining riders from period 1 then decide whether to request rides, and those who do are again randomly matched to available drivers. (For simplicity, one can assume that no new riders appear in period 2.) Riders discount valuations of getting rides in period 2. This model can explain why surge pricing, a short-lived sharp price surge followed by a lower price, can even benefit consumer surplus: The platform strategically inflates the initial price to make riders voluntarily wait out the initial surge period, so as to attract drivers to the surge region and then, serve riders with a lower price. The surge pricing has a signaling effect to drivers that the higher the surged price, the more unserved riders out there, without any demand information communicated to drivers.

On the supply side, Afèche et al. (2018b) study the drivers' forward-looking behavior in deciding whether to reposition toward the hotspot after an unexpected demand shock of uncertain duration occurs, given their location-dependent repositioning delay and payoff risk. The paper compares the performances of various dynamic policies of setting driver wages and rider prices by the platform, taking into account the interplay of three timescales, rider patience, demand shock duration, and drivers' repositioning delays. As mentioned earlier, Chu et al. (2018) study the cherry-picking behavior by drivers, who may skip low-payoff riders and wait for high-payoff riders. Moreover, Bimpikis et al. (2019), Afèche et al. (2018a), and Besbes et al. (2019) focus on the drivers' strategic spatial repositioning, which have been mentioned above; see also Guda and Subramanian (2019) in the marketing literature.

Research opportunities.

1. Both sellers and buyers are strategic. Chen and Hu (2020b) allow both sellers and buyers of a homogeneous good or service to be inter-temporally forward-looking, and show a two-sided pricing and matching policy that induces their myopic behavior is asymptotically optimal. There is a lot that can be further explored in this area. For instance, one wants to build more realistic models that account for contingent pricing and joint strategic drivers' and riders' waiting and repositioning behavior.
2. Mechanism design. Drivers and riders have their own private information, and it is a critical problem of how the platform can design mechanisms to tease out such information so to enable more efficient matching (see Ma et al. 2019 for a mechanism design problem in ride-hailing under complete information).

7. From One-Sided to Two-Sided Competition

The most classic competition models in economics are Bertrand (price) and Cournot (quantity) competition models. Here we briefly review a version of those models for differentiated products, which we will extend to account for the competition between two ride-hailing platforms such as Uber and Lyft. Consider a symmetric duopoly market with firms 1 and 2 each selling one product by setting prices p_1 and p_2, respectively. The two products are substitutable. The demand system that governs the market has a linear structure: provided that demand quantities are positive,

$$d_1(p_1,p_2)=a-p_1+\alpha p_2,\quad d_2(p_1,p_2)=a-p_2+\alpha p_1, \tag{2}$$

where $a>0$ is the potential market size and $\alpha \in [0,1)$ measures the substitutability between the two products. Such a linear demand structure can be obtained from a Hotelling line model or a representative consumer maximizing a quadratic utility function. Both firms have the same marginal cost to procure, produce and distribute their products. In the Bertrand competition, both firms simultaneously set *prices* to maximize their profit. In the Cournot competition, both firms simultaneously make decisions on the targeted *sales quantity* to maximize their profit, and the prices are resolved as the market clearing prices such that the targets are achieved in the competitive market. The celebrated Kreps and Scheinkman equivalency says that Bertrand price competition with precommitted capacity yields the same equilibrium outcome as Cournot competition (see Kreps and Scheinkman 1983). That is, consider a two-stage sequential game in which the firms maximize their profit. In the first stage, both firms simultaneously decide on the *capacity* of their production and distribution. In the second stage, given the capacity level they build in the first stage, both firms simultaneously set *prices*. The outcome of this two-stage game is the same as that of Cournot competition. The intuition behind the Kreps and Scheinkman equivalency is that precommitment in capacity can curb throat-cutting price competition

Please Cite this article in press as: Hu, M. From the Classics to New Tunes: A Neoclassical View on Sharing Economy and Innovative Marketplaces. *Production and Operations Management* (2020), https://doi.org/10.1111/poms.13330

because firms cannot sell beyond their capacity level.

Now one can ask similar questions for a competitive market where two platforms fight on two fronts. That is, on the one hand, platforms compete in offering wages to the drivers and on the other hand, they compete in setting prices posted to the riders, see Figure 3. On the demand side, we can adopt the above linear demand system Eq. (2). On the supply side, we can also assume a linear structure: provided that supply quantities are positive,

$$s_1(w_1,w_2) = w_1 - \beta w_2, \quad s_2(w_1,w_2) = w_2 - \beta w_1,$$

where $\beta \in [0,1)$ measures the substitutability between the two platforms in the eyes of drivers. Similar to section 3, we consider a short-term objective of the platforms. At the operational level, given the competitor's decisions, the objective of each platform i, i=1,2 is to solve

$$\max_{w_i,p_i} (p_i - w_i) \min\{s_i(w_i,w_{3-i}), d_i(p_i,p_{3-i})\}. \tag{P2}$$

We assume platforms simultaneously make their wage and price decisions. This game (P^2) is a natural extension of the two-sided pricing problem (P) to two-sided competition and that of the one-sided Bertrand competition to two-sided wage and price competition, for simplicity, referred to as two-sided price competition.

The counterpart of the Cournot competition in the two-sided market is that firms simultaneously decide on the matching quantity and then, on both supply and demand sides, market clearing wages and prices are resolved to achieve the targeted matching quantity, which is referred to as two-sided quantity competition. Hu and Liu (2019) show that analogous to the comparison between the classic Cournot and Bertrand competition, two-sided quantity competition leads to more alleviated market outcomes with higher market prices than two-sided pricing competition. Moreover, a two-sided counterpart of the Kreps and Scheinkman equivalency remains to hold. That is, consider a two-stage sequential game in which platforms maximize their profit. In the first stage, both platforms simultaneously decide on the capacity level which limits their matching quantity. In the second stage, given the capacity level they choose in the first stage, both firms simultaneously set wages and prices. The authors show that the outcome of this two-stage game is the same as that of two-sided quantity competition. These results demonstrate the connection of the two-sided competition with the classic one-sided competition, and show that the insights from the one-sided competition remain robust for the two-sided competition.

Here, come the new tunes. Because now the platforms compete on both supply and demand sides, motivated by the competition between Uber and Lyft, one can consider two-stage games with other precommitment devices than capacity. First, the platforms can commit to wages and then, compete on prices. Second, the platforms can commit to prices and then, compete on wages. Third, the platforms can commit

Figure 3 Two-Sided Competition

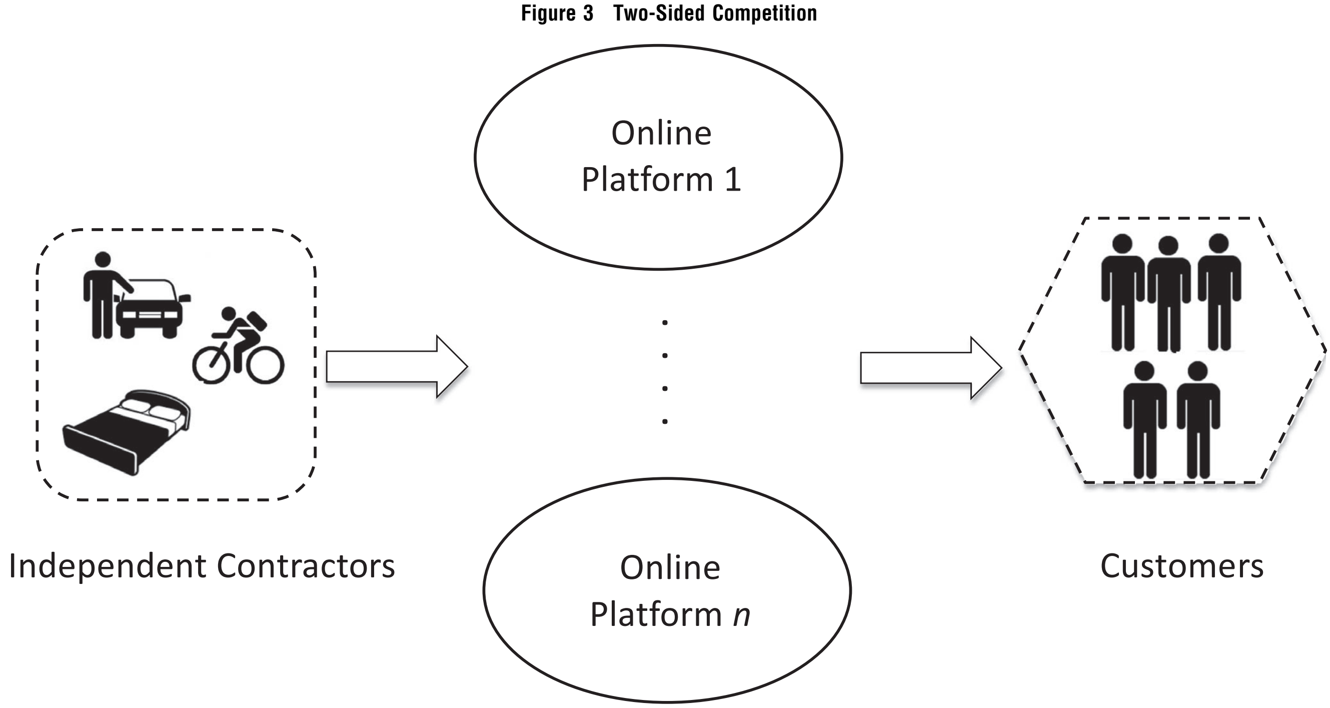

to the commissions, or equivalently, the payout ratios, that is, the ratio between wage and price, and then, compete on prices. By comparing these competition modes, Hu and Liu (2019) show the following new insights among others. First, precommitment effectiveness can vary depending on the competitiveness of two sides of the market. In one-sided price competition, it is always more beneficial to firms if they commit on the supply side. In two-sided price competition, platforms need to investigate which side, supply or demand side, of the market, is more competitive. Precommitment on the more competitive side actually intensifies the competition and performs worse than no commitment at all. Second, in a deterministic setting, the firm has an incentive to set wage and price so that supply is precisely matched with demand, which is referred to as "two-sided balancing." This two-sided balancing incentive (the same reason that drives the downward shape of the optimal price as a function of the exogenous wage in the two-sided pricing problem (P), see section 3) serves as an intrinsic constraint that alleviates competition on the more competitive side. These insights prevail even with market uncertainty in an extension where certain precommitment occurs in the first stage, before a contingent decision is made in the second stage, reacting to the realized market condition.

In a two-sided competition model, Bernstein et al. (2019) build supply and demand structures on top of the linear systems by subtracting a disutility term that depends on supply and demand utilizations to capture the congestion effects when the two sides are unbalanced. The authors focus on comparing single-homing vs. multi-homing behavior by the drivers (i.e., drivers have their dedicated platform, vs. drivers can serve both platforms). Cohen and Zhang (2019) adopt a MultiNomial Logit choice model for the demand side (another commonly used demand structure, see Federgruen and Hu 2017) and a general supply curve for each platform assuming single-homing. Their focus is to investigate the impact of a new service jointly offered by competing platforms where one platform's dedicated drivers can provide shared rides that are available to riders from either platform.

Research opportunities. Two-sided competition has not been much explored in the operations management literature. One direction is to take a one-platform two-sided model as a base (see, e.g., Bai and Tang 2020) and extend it to competition, and another is to extend a one-sided competition model (see, e.g., Federgruen and Hu 2015) to two-sided competition. In either path, the new two-sided competition model needs to be practically grounded and generate novel results/insights to make an impact. Moreover, a fruitful route can also be to study features arising from sharing economy such as competition on the payout rate and bonus payment to drivers (see, e.g., Chen et al. 2020a).

8. From Supply Chain to Marketplace

The classic supply chain model involves three parties with a supplier (or manufacturer) selling through a retailer to customers, see Figure 4. Nowadays a typical marketplace involves three parties as well, with a group of independent sellers selling on a platform to customers, see Figure 5. The key differences are (i) suppliers in a marketplace are usually in a large quantity, each with limited influence on the whole market; and (ii) unlike a retailer, the platform that operates a marketplace may have limited control on a transaction between a seller and a customer.

The supply chain contract literature focuses on how the supplier can offer a contract to the retailer to maximize the supply chain surplus and then, split the benefit that has grown to the largest (see, e.g., Cachon 2003). Unlike the traditional supply chain, the independent sellers in some marketplace have the full right to set their own prices. But the platform can influence the sellers' decisions by offering a crowdsourcing contract, or a financing scheme to sellers (see, e.g., Dong et al. 2018), or influence the matching process between buyers and sellers, for example, by manipulating information disclosed to an interested

Figure 4 Traditional Supply Chain

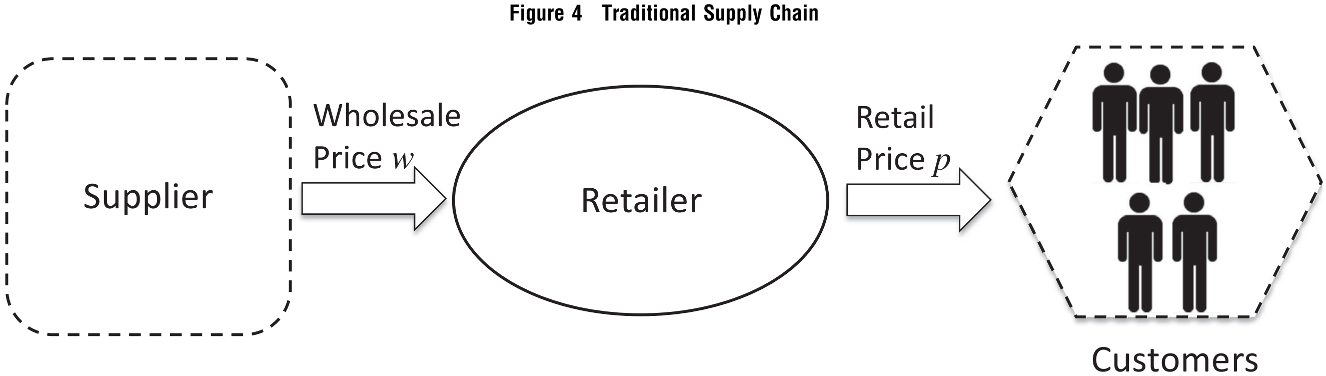

Please Cite this article in press as: Hu, M. From the Classics to New Tunes: A Neoclassical View on Sharing Economy and Innovative Marketplaces. *Production and Operations Management* (2020), https://doi.org/10.1111/poms.13330

Figure 5 Platform Economy

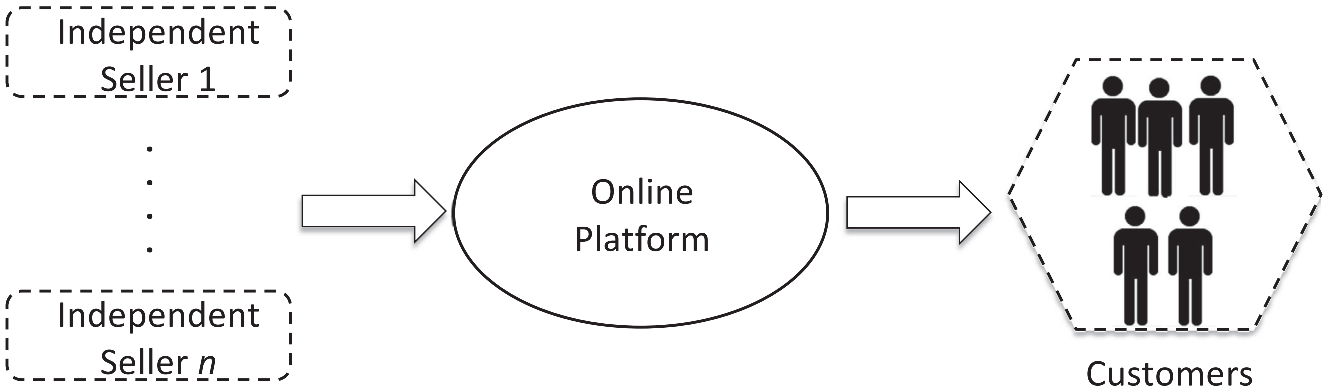

party (see, e.g., Liu et al. 2020). To use Airbnb as an example, hosts on the platform have full freedom to decide the price they would love to charge for their listings. There are individual property owners, or corporate players who control multiple assets (see Zhu et al. 2019). For any type of hosts, the platform can optimize the commission rate it proportionally takes at the strategic level or customize the ranking of listings made available to a guest.

As mentioned earlier, platforms often charge a commission for each transaction. That is, for whatever price posted to the customers, regardless of whether that pricing decision is made by the platform (e.g., in the case of Uber) or by the independent sellers (e.g., in the case of Airbnb), the platform takes a fixed fraction out of the total price paid by customers. This commission contract structure is meant to qualify a platform for the legal status of being a broker. This commission contract literally is a revenue-sharing contract, which can coordinate a traditional supply chain with one supplier and one retailer. However, it is easy to see that such a commission contract signed between an independent seller and the platform would most likely not be in the best interest of the platform, because the seller now makes the pricing decision. Under competition, a seller may be forced to set a price lower than what is desired by the platform.

The commission contract is signed and applied for a relatively long period of time, while prices may vary at the transaction level. For example, a commission rate is fixed for Uber drivers over at least a couple of months, but prices Uber sets for each ride vary depending on the market conditions. Hu and Zhou (2018) show that by carefully selecting a commission rate, even though the platform may lose some flexibility in moderating supply and demand (as now the wage is derived from the pricing decision under the commission contract), the profit loss may not be significant, and there exists a provable performance guarantee. Hu and Liu (2019) show that the commission contract as a constraint governing the wage and price may, in fact, benefit platforms under competition (compared to no commission contract).

Research opportunities.

1. Crowdsourcing contract. Parallel to supply chain contracting, it is interesting to study crowdsourcing contracts of various forms, which affect the incentive of the independent suppliers/vendors. For example, Chen et al. (2019b) propose a compensation-while-idling mechanism, offered to independent self-interested suppliers, to achieve a centralized solution for a marketplace. Balseiro et al. (2019) study the mechanism design problem of an intermediary who offers a contract to an advertiser in an online advertising marketplace. For another example, Netflix and Spotify crowdsource movies and music from numerous vendors and artists and bundle them to sell at the price of one (see Hu and Wang 2019). In the case of Spotify, the collected subscription fees for the bundle are allocated according to the realized contribution, among all, by each crowdsourced product. In the case of Netflix, the platform negotiates payments for each product ex-ante before releasing it on the platform.
2. Social utility. The sharing economy is based on the sharing activities between participants. As the social structure is embedded as an indispensable foundation, the sharing economy has blurred the line between economic and social transactions. For example, Cui et al. (2020b) show that on Airbnb which facilitates shared living between a guest and a host, the guest may draw a social utility from staying with the host. The existence of social utility can create an incentive misalignment between the platform and a host. Think about the joy of offering hospitality to travelers around the world, which is the idea behind couchsurfing and

Please Cite this article in press as: Hu, M. From the Classics to New Tunes: A Neoclassical View on Sharing Economy and Innovative Marketplaces. *Production and Operations Management* (2020), https://doi.org/10.1111/poms.13330

leads to the birth of Airbnb. With this joy, the host has an incentive to set a price lower than without social utility and lower than what the platform desires. For another example, ride-hailing and food delivery platforms often allow tipping to service providers, which encourages social interactions and is not taxed by the platforms. With the presence of social utility in sharing economy justified, it is then valuable to study how such presence would affect the platform's decisions.

3. Platform intermediation. In practice, many marketplace operators are making efforts to move toward a more centrally controlled platform to gain efficiency or/and profitability. For instance, eBay now provides price suggestions to the sellers and buyers in the form of "the product is trending" at this particular price to facilitate transactions. Other than priming the sellers or using a crowdsourcing contract to influence their decisions, the platform that runs the marketplace can adopt other moderating tools depending on the context. For example, Allon et al. (2012) study the impact of facilitating the matching between service providers and customers, and enabling communication among service providers, in an online service marketplace. Arnosti et al. (2020) study a matching market in which participants face search and screening costs when seeking a match. The authors focus on the platform's moderating controls, such as limiting the number of applications that an individual can send or making it more costly to apply. Bimpikis and Papanastasiou (2019) study information disclosure mechanisms by a platform. Liu et al. (2018) investigate the commission or joint commission and price control by an on-demand healthcare service platform. It is worthwhile to delve into the operations of a specific platform to study how the platform could improve profitability or efficiency.

9. From Optimal Transport to Matching Supply with Demand

The origin of matching supply with demand is about the optimal transportation and resource allocation. There can be a fundamental, yet unexplored, connection between the optimal transport problem, a classic problem intensively studied in mathematics (see, e.g., Villani 2009), and the operational problem of matching supply with demand. The optimal transport problem is to find the most cost-effective way to move mass, for example, to find a way that transforms a given probability density function into another that minimizes the cost of transport. This problem has accumulated a large body of deep theories that have resulted in a couple of Fields medalists and may deepen our understanding of matching supply with demand. The optimal transport problem treats the original status and the targeted level as continuous distributions over a continuous space. To see its connection with operations, imagine a heat map indicating the distribution of autonomous cars. A central planner could move around those cars to serve riders. Traditional operations management models often assume the system state and space to be discrete to stay closer to reality, for example, there is an integral number of echelons/stages in an inventory system, and there are discrete units of capacity to sell in a revenue management setting. We do see promising evidence that relaxing those discrete assumptions may lead to neater results, see, for example, Song and Zipkin (2013), not to mention the trivial example that the Economic Order Quantity formula gives a continuous quantity as the optimal solution.

Research opportunities. In applying existing theories from the optimal transport problem to operational problems, there are still significant gaps, which may call for developing new theories.

1. Temporal dimension. The classic optimal transport problem only has a spatial dimension. In the physical world, moving resources is not only costly but also time consuming. It is highly desirable to incorporate the time dimension into the optimal transport problem.
2. Uncertainty. The optimal transport problem does not have uncertainty. In operational problems, there exists uncertainty on both sides of supply and demand.
3. Moderation. In a matching market, both supply (corresponding to the original status) and demand (corresponding to the target level) can be moderated by incentives such as wage and price before the matching.
4. Decentralized incentive. The optimal transport is a centralized solution dictating how resources should be moved. In ride-hailing, resources such as drivers, unlike autonomous cars, have self-interested incentives to be fulfilled. Besbes et al. (2019) study a decentralized version of the one- or two-dimensional optimal transport problem in which drivers self-interestedly respond to the platform's (surge) pricing decisions.

It would be highly desirable to extend some of the primary forms of solution characterizations of the optimal transport problem to account for the above mentioned operational features in a general form.

10. From Efficiency or Profit Maximization to Alternative Objectives and Other Social and Environmental Issues

In almost all the above discussions, the default objective is to maximize the system's operational efficiency or the firm's profitability. As there are at least three parties of stakeholders involved in a sharing economy, alternative objectives could be a weighted sum of the surpluses of buyers, sellers, and the intermediary firm. For example, more legislation such as California Assembly Bill 5 is signed to require ride-hailing platforms to put more weight on the drivers' welfare in their decision making. As explained in section 3, an increase in the drivers' welfare can entice more drivers to come out and work and force the platform to lower prices, which benefit riders at the same time.

Moreover, we are also faced with a brand-new set of unintended or unaccounted-for social and environmental consequences that arise from the emergence of sharing economy. For example, Chen et al. (2019a) show that more customers having access to a food delivery service may hurt the food delivery platform itself and the society. Thus it can be practically impactful to figure out how we can match service providers and products with users more efficiently while minimizing some of the problems and waste caused by these technologies. Agrawal et al. (2019) point out challenges and research opportunities in studying environmental sustainability in the circular economy. Here, we provide some additional research ideas that one may explore in a specific sharing economy or innovative marketplace.

Research opportunities.

1. Alternative objectives. All the previously discussed formulations may be re-examined under an alternative objective. One can also compare the optimal policies under different objectives.
2. Ride-hailing. The great thing about ride-sharing apps is that riders only have to wait a few minutes for a ride, but that convenience comes with a price. Drivers are waiting idle on city streets for requests or have to come from a long way to pickup customers. This phenomenon of a so-called "wild goose chase" (see Castillo et al. 2017) leads to the low utilization of the cars and possibly results in increasing traffic in cities. Hu (2020) shows that both temporal and spatial pooling (by delay matching and designating pickup and drop-off areas, respectively), common notions in the operations management, can resolve the issue to a large extent (see also Yan et al. 2020 for an approach called "dynamic waiting" which is similar to temporal pooling). Besides, with a unique gender perspective, Guo et al. (2020) examine how female users' safety concerns affect the system configuration of ride-hailing platforms. Lastly, Asadpour et al. (2020) study the implications of utilization-based minimum earning regulations on ride-hailing service providers. As ride-hailing has penetrated our daily lives, studies on various public policies could have a considerable impact.
3. Food ordering and delivery. The food delivery services, such as Uber Eats, while removing the hassle of cooking, are adding to our landfills. The packaging that comes with the food orders is often excessive or not recyclable. It is critical to consider the environmental downside in the strategic and operational decisions in this emerging industry.
4. Short-term rental. Airbnb is thought to be crowding out long-term renters. Empirical research is desirable to quantify such an impact and guide policy making to curb any potential downside.
5. Clothes-sharing. Clothes-sharing firms such as Rent the Runway and Style Lend may be encouraging customers—who know that they can resell or share their purchases—to buy more clothes than they need. This not only increases the volume of clothes in circulation but incurs further environmental costs, such as cleaning, shipping, and packaging, whenever these items change hands. A holistic view needs to be taken to evaluate the overall impact of the notion of sharing clothes.

11. Concluding Remarks

The field of sharing economy and innovative marketplaces is an exciting area. The problems of matching supply with demand in those contexts can be well rooted in practice and connected with the classic fields in operations management such as inventory, supply chain, revenue management, and queueing. This article aims at building bridges between those classic problems and the new applications of sharing economy, which allow us to transfer questions, techniques, intuitions, and insights from one side to the other.

Moreover, the new applications also motivate us to learn, invent and develop new questions, techniques, intuitions and insights, which help us to better understand the essences of the problem of matching supply with demand. In particular, the new applications in sharing economy and innovative marketplaces typically may have the following features, in contrast to the classic settings, which need to be considered:

1. Many independent sellers. The suppliers are independent decision makers and there are a large number of them. Analyzing such suppliers' behavior may require approximation tools such as mean-field approximation (see, e.g., Balseiro et al. 2015) and non-atomic game theory.
2. Supply uncertainty. The supply side has uncertainty, due to the independent, self-scheduling behavior by suppliers.
3. Platform intermediation. The platform can influence the decision making of sellers and buyers who may have private information. The theory of mechanism design is a desirable toolset for designing moderation schemes.
4. Stakeholder welfare. As there are multiple parties involved, it is critical to take into account the welfare of all parties such as sellers, buyers, and the platforms.

Finally, the applications of sharing economy are likely to evolve (e.g., crowdsourced drivers and human couriers may be replaced by autonomous cars and delivery robots), but the theme of matching supply with demand will not fade away.

Acknowledgments

The authors thank Mingliu Chen and Yun Zhou for careful reading of an earlier version of the paper. The authors also thank an anonymous reviewer for her/his constructive and insightful suggestions, which improved the paper. This research is in part supported by the Natural Sciences and Engineering Research Council of Canada [Grant RGPIN-2015-06757].

References

Adelman, D. 2007. Dynamic bid prices in revenue management. *Oper. Res.* **55**(4): 647–661.

Afèche, P., Z. Liu, C. Maglaras. 2018a. Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance. Available at http://ssrn.com/abstract=3120544 (accessed date November 21, 2020).

Afèche, P., Z. Liu, C. Maglaras. 2018b. A tale of timescales: Surge pricing and dynamic matching for hotspot demand shock in ridehailing networks. Presented at 2018 INFORMS Revenue Management and Pricing Section Conference.

Agrawal, V. V., A. Atasu, L. N. Van Wassenhove. 2019. OM forum–New opportunities for operations management research in sustainability. *Manuf. Serv. Oper. Manag.* **21**(1): 1–12.

Allon, G., A. Bassamboo, E. B. Çil. 2012. Large-scale service marketplaces: The role of the moderating firm. *Management Sci.* **58** (10): 1854–1872.

Allon, G., M. Cohen, P. Sinchaisri. 2018. The impact of behavioral and economic drivers on gig economy workers. Available at http://ssrn.com/abstract=3274628 (accessed date November 21, 2020).

Arnosti, N., R. Johari, Y. Kanoria. 2020. Managing congestion in matching markets. *Manuf. Serv. Oper. Manag.* Forthcoming.

Asadpour, A., I. Lobel, G. van Ryzin. 2020. Minimum earnings regulation and the stability of marketplaces. Available at http://ssrn.com/abstract=3502607 (accessed date November 21, 2020).

Ata, B., N. Barjesteh, S. Kumar. 2019. Spatial pricing: An empirical analysis of taxi rides in New York City. Working paper.

Baccara, M., S. Lee, L. Yariv. 2020. Optimal dynamic matching. *Theor. Econ.* **15**(3): 1221–1278.

Bai, J., C. S. Tang. 2020. Can two competing on-demand service platforms be both profitable? Available at http://ssrn.com/abstract=3282395 (accessed date November 21, 2020).

Bai, J., K. C. So, C. S. Tang, X. Chen, H. Wang. 2019. Coordinating supply and demand on an on-demand service platform with impatient customers. *Manuf. Serv. Oper. Manag.* **21**(3): 556–570.

Balseiro, S. R., O. Besbes, G. Y. Weintraub. 2015. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Sci.* **61**(4): 864–884.

Balseiro, S. R., O. Candogan, H. Gurkan. 2019. Intermediation in online advertising. M. Hu, ed. *Sharing Economy: Making Supply Meet Demand*, chap. 21. Springer, Cham, 505–528.

Balseiro, S. R., D. B. Brown, C. Chen. 2020. Dynamic pricing of relocating resources in large networks. *Management Sci.* Forthcoming.

Banerjee, S., C. Riquelme, R. Johari. 2015. Pricing in ride-share platforms: A queueing-theoretic approach. Proceedings of the 2015 ACM Conference on Economics and Computation. EC '15, pp. 639–639.

Benjaafar, S., M. Hu. 2020. Operations management in the age of the sharing economy: What is old and what is new? *Manuf. Serv. Oper. Manag.* **22**(1): 93–101.

Benjaafar, S., G. Kong, X. Li, C. Courcoubetis. 2019. Peer-to-peer product sharing: Implications for ownership, usage and social welfare in the sharing economy. *Management Sci.* **65**(2): 477–493.

Benjaafar, S., J.-Y. Ding, G. Kong, T. Taylor. 2020. Labor welfare in on-demand service platforms. *Manuf. Serv. Oper. Manag.* Forthcoming.

Bernstein, F., G. DeCroix, N. B. Keskin. 2019. Competition between two-sided platforms under demand and supply congestion effects. *Manuf. Serv. Oper. Manag.* Forthcoming.

Bertsimas, D. J., G. Van Ryzin. 1991. A stochastic and dynamic vehicle routing problem in the Euclidean plane. *Oper. Res.* **39** (4): 601–615.

Bertsimas, D. J., G. Van Ryzin. 1993. Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles. *Oper. Res.* **41**(1): 60–76.

Besbes, O., F. Castro, I. Lobel. 2018. Spatial capacity planning. Available at http://ssrn.com/abstract=3292651 (accessed date November 21, 2020).

Besbes, O., F. Castro, I. Lobel. 2019. Surge pricing and its spatial supply response *Management Sci.* Forthcoming.

Bimpikis, K., Y. Papanastasiou. 2019. Inducing exploration in service platforms. M. Hu, ed. *Sharing Economy: Making Supply Meet Demand*, chap. 10. Springer, Cham, 193–216.

Bimpikis, K., O. Candogan, D. Saban. 2019. Spatial pricing in ride-sharing networks. *Oper. Res.* **67**(3): 744–769.

Bimpikis, K., W. J. Elmaghraby, K. Moon, W. Zhang. 2020. Managing market thickness in online B2B markets. *Management Sci.* Forthcoming.

Please Cite this article in press as: Hu, M. From the Classics to New Tunes: A Neoclassical View on Sharing Economy and Innovative Marketplaces. *Production and Operations Management* (2020), https://doi.org/10.1111/poms.13330

Braverman, A., J. Dai, X. Liu, L. Ying. 2019. Empty-car routing in ridesharing systems. *Oper. Res.* **67**(5): 1437–1452.

Cachon, G. P. 2003. Supply chain coordination with contracts. S. Graves, T. de Kok, eds. *Handbooks in Operations Research and Management Science: Supply Chain Management*, vol. 11. Elsevier, Amsterdam, 227–339.

Cachon, G. P., P. Feldman. 2011. Pricing services subject to congestion: Charge per-use fees or sell subscriptions? *Manuf. Serv. Oper. Manag.* **13**(2): 244–260.

Cachon, G. P., K. Daniels, R. Lobel. 2017. The role of surge pricing on a service platform with self-scheduling capacity. *Manuf. Serv. Oper. Manag.* **19**(3): 337–507.

Castillo, J. C., D. Knoepfle, G. Weyl. 2017. Surge pricing solves the wild goose chase. Proceedings of the 2017 ACM Conference on Economics and Computation. EC '17, pp. 241–242.

Chen, M., M. Hu. 2020a. Courier dispatch in on-demand delivery. Available at http://ssrn.com/abstract=3675063 (accessed date November 21, 2020).

Chen, Y., M. Hu. 2020b. Pricing and matching with forward-looking buyers and sellers. *Manuf. Serv. Oper. Manag.* **22**(4): 717–734.

Chen, X., D. Simchi-Levi. 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, X., Z. Pang, L. Pan. 2014. Coordinating inventory control and pricing strategies for perishable products. *Oper. Res.* **62** (2): 284–300.

Chen, M., M. Hu, J. Wang. 2019a. Food delivery service and restaurant: Friend or foe? Available at http://ssrn.com/abstract=3469971 (accessed date November 21, 2020).

Chen, Y.-J., C. Maglaras, G. Vulcano. 2019b. Design of an aggregated marketplace under congestion effects: Asymptotic analysis and equilibrium characterization. M. Hu, ed. *Sharing Economy: Making Supply Meet Demand*, chap. 11. Springer, Cham, 217–248.

Chen, Z., M. Hu, Y. Zhou. 2019c. Optimal dynamic matching with impatient demand and patient supply. Available at http://ssrn.com/abstract=3329000 (accessed date November 21, 2020).

Chen, L., Y. Cui, J. Liu, X. Liu. 2020a. Wage schemes and platform competition. Available at http://ssrn.com/abstract=3392700 (accessed date November 21, 2020).

Chen, N., G. Gallego, P. Gao, S. Kou. 2020b. Dealership or marketplace: A dynamic comparison. Available at https://ssrn.com/abstract=3472161 (accessed date November 21, 2020).

Chen, Y.-J., T. Dai, C. G. Körpeoğlu, E. Körpeoğlu, O. Sahin, C. S. Tang, S. Xiao. 2020c. OM forum–Innovative online platforms: Research opportunities. *Manuf. Serv. Oper. Manag.* **22**(3): 430–445.

Chu, L. Y., Z. Wan, D. Zhan. 2018. Harnessing the double-edged sword via routing: Information provision on ride-hailing platforms. Available at http://ssrn.com/abstract=3266250 (accessed date November 21, 2020).

Cohen, M. C., R. Zhang. 2019. Competition and coopetition for two-sided platforms. Available at http://ssrn.com/abstract=3028138 (accessed date November 21, 2020).

Cohen, M., M.-D. Fiszer, B. J. Kim. 2020. Frustration-based promotions: Field experiments in ride-sharing. *Management Sci.* Minor revision.

Cui, R., J. Li, D. J. Zhang. 2020a. Reducing discrimination with reviews in the sharing economy: Evidence from field experiments on Airbnb. *Management Sci.* **66**(3): 1071–1094.

Cui, Y., A. Y. Orhun, M. Hu. 2020b. Under the same roof: Value of shared living in airbnb. Available at http://ssrn.com/abstract=3136138 (accessed date November 21, 2020).

Dong, L., L. Ren, D. J. Zhang. 2018. Financing small and medium-size enterprises via retail platforms. Available at http://ssrn.com/abstract=3257899 (accessed date November 21, 2020).

Federgruen, A., A. Heching. 1999. Combined pricing and inventory control under uncertainty. *Oper. Res.* **47**(3): 454–475.

Federgruen, A., M. Hu. 2015. Multi-product price and assortment competition. *Oper. Res.* **63**(3): 572–584.

Federgruen, A., M. Hu. 2017. Competition in multi-echelon systems. R. Batta, J. Peng, eds. *INFORMS Tutorials in Operations Research.* INFORMS, Catonsville, MD, 178–211.

Feng, G., G. Kong, Z. Wang. 2020. We are on the way: Analysis of on-demand ride-hailing systems. *Manuf. Serv. Oper. Manag.* Forthcoming.

Freund, D., S. G. Henderson, D. B. Shmoys. 2019. Bike sharing. M. Hu, ed. *Sharing Economy: Making Supply Meet Demand*, chap. 18. Springer, Cham, 435–459.

Gallego, G., G. Van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* **40**(8): 999–1020.

Garg, N., H. Nazerzadeh. 2020. Driver surge pricing. Available at https://ssrn.com/abstract=3390346 (accessed date November 21, 2020).

Gorbushin, A., M. Hu, Y. Zhou. 2020. Courier sharing in food delivery. Available at http://ssrn.com/abstract=3548978 (accessed date November 21, 2020).

Guda, H., U. Subramanian. 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., C. S. Tang, Y. Tang, Y. Wang. 2020. Gender-based operational issues arising from on-demand ridehailing platforms: Safety concerns and system configuration. Available at http://ssrn.com/abstract=3260427 (accessed date November 21, 2020).

Gurvich, I., A. Ward. 2014. On the dynamic control of matching queues. *Stoch. Syst.* **4**(2): 479–523.

Gurvich, I., M. Lariviere, A. Moreno. 2019. Operations in the on-demand economy: Staffing services with selfscheduling capacity. M. Hu, ed. *Sharing Economy: Making Supply Meet Demand*, chap. 12. Springer, Cham, 249–278.

Hassin, R., M. Haviv. 2003. *To Queue or Not to Queue: Equilibrium Behavior in Queueing Systems*, vol. 59. International Series in Operations Research & Management Science, Springer US.

He, L., H.-Y. Mak, Y. Rong, Z.-J. M. Shen. 2017. Service region design for urban electric vehicle sharing systems. *Manuf. Serv. Oper. Manag.* **19**(2): 309–327.

Hu, M., ed. 2019. *Sharing Economy: Making Supply Meet Demand*. Springer Series in Supply Chain Management, vol. 6. Springer, Cham.

Hu, M. 2020. Spatial or temporal pooling solves wild goose chase. Available at http://ssrn.com/abstract=3676065 (accessed date November 21, 2020).

Hu, M., Y. Liu. 2019. Precommitments in two-sided market competition. Available at http://ssrn.com/abstract=3734828 (accessed date November 21, 2020).

Hu, M., L. Wang. 2019. Created unequal: Bundling with crowdsourced products. Available at https://ssrn.com/abstract=3338314 (accessed date November 21, 2020).

Hu, M., Y. Zhou. 2018. Price, wage and fixed commission in ondemanmd matching. Available at http://ssrn.com/abstract=2949513 (accessed date November 21, 2020).

Hu, M., Y. Zhou. 2020. Dynamic type matching. *Manuf. Serv. Oper. Manag.* Forthcoming.

Hu, B., M. Hu, H. Zhu. 2020a. Surge pricing and two-sided temporal responses in ride-hailing. *Manuf. Serv. Oper. Manag.* Forthcoming.

Hu, M., J. Wang, H. Wen. 2020b. Share or solo? Individual and social choices in ride-hailing. Available at http://ssrn.com/abstract=3675050 (accessed date November 21, 2020).

Ibrahim, R. 2018. Managing queueing systems where capacity is random and customers are impatient. *Prod. Oper. Manag.* **27** (2): 234–250.

Jacob, J., R. Roet-Green. 2019. Ride solo or pool: The impact of sharing on optimal pricing of ride-sharing services. Available at http://ssrn.com/abstract=3008136 (accessed date November 21, 2020).

Jiang, Z.-Z., G. Kong, Y. Zhang. 2020. Making the most of your regret: Workers' relocation in on-demand platforms. *Manuf. Serv. Oper. Manag.* Forthcoming.

Kabra, A., E. Belavina, K. Girotra. 2020. Bike-share systems: Accessibility and availability. *Management Sci.* **66**(9): 3803–3824.

Karp, R. M., U. V. Vazirani, V. V. Vazirani. 1990. An optimal algorithm for on-line bipartite matching. Proceedings of the Twenty-Second Annual ACM Symposium on Theory of Computing. ACM, pp. 352–358.

Kendall, D. G. 1951. Some problems in the theory of queues. *J. R. Statist. Soc. B Methodol.* **13**(2): 151–185.

Kreps, D. M., J. A. Scheinkman. 1983. Quantity precommitment and bertrand competition yield cournot outcomes. *Bell J. Econo.* **14**(2): 326–337.

Li, J., A. Moreno, D. J. Zhang. 2019. Agent pricing in the sharing economy: Evidence from Airbnb. M. Hu, ed. *Sharing Economy: Making Supply Meet Demand*, chap. 20. Springer, Cham, 485–503.

Liu, Y., X. Wang, S. Gilbert, G. Lai. 2018. Pricing, quality and competition at on-demand healthcare service platforms. Available at http://ssrn.com/abstract=3253855 (accessed date November 21, 2020).

Liu, Z., D. J. Zhang, F. Zhang. 2020. Information sharing on retail platforms. *Manuf. Serv. Oper. Manag.* Forthcoming.

Ma, W., D. Simchi-Levi. 2020. Algorithms for online matching, assortment, and pricing with tight weight-dependent competitive ratios. *Oper. Res.* Forthcoming.

Ma, H., F. Fang, D. C. Parkes. 2019. Spatio-temporal pricing for ridesharing platforms. Proceedings of the 2019 ACM Conference on Economics and Computation. EC '19, pp. 583–583.

Ming, L., T. I. Tunca, Y. Xu, W. Zhu. 2020. An empirical analysis of price formation, utilization, and value generation in ride sharing services. Available at https://ssrn.com/abstract=3338088 (accessed date November 21, 2020).

Özkan, E., A. R. Ward. 2020. Dynamic matching for real-time ride sharing. *Stoch. Syst.* **10**(1): 29–70.

Petruzzi, N. C., M. Dada. 1999. Pricing and the newsvendor problem: A review with extensions. *Oper. Res.* **47**(2): 183–194.

Rochet, J. C., J. Tirole. 2003. Platform competition in two-sided markets. *J. Eur. Econ. Assoc.* **1**(4): 990–1029.

Shen, Z.-J. M, X. Su. 2007. Customer behavior modeling in revenue management and auctions: A review and new research opportunities. *Prod. Oper. Manag.* **16**(6): 713–728.

Song, J.-S., P. Zipkin. 2013. Supply streams. *Manuf. Serv. Oper. Manag.* **15**(3): 444–457.

Taylor, T. A. 2018. On-demand service platforms. *Manuf. Serv. Oper. Manag.* **20**(4): 704–720.

Truong, V.-A., X. Wang. 2019. Prophet inequality with correlated arrival probabilities, with application to two sided matchings. Available at https://ssrn.com/abstract=3338088 (accessed date November 21, 2020).

Villani, C. 2009. *Optimal Transport: Old and New*, vol. 338. Springer-Verlag, Grundlehren der mathematischen Wissenschaften. Berlin, Heidelberg.

Yan, C., H. Zhu, N. Korolko, D. Woodard. 2020. Dynamic pricing and matching in ride-hailing platforms. *Nav. Res. Log.* **67**(8): 705–724.

Zhu, H., M. Ray, S. Ray. 2019. Short-term housing rentals and corporatization of platform pricing. Available at http://ssrn.com/abstract=3344260 (accessed date November 21, 2020).

Zipkin, P. 2008. On the structure of lost-sales inventory models. *Oper. Res.* **56**(4): 937–944.

Please Cite this article in press as: Hu, M. From the Classics to New Tunes: A Neoclassical View on Sharing Economy and Innovative Marketplaces. *Production and Operations Management* (2020), https://doi.org/10.1111/poms.13330