Driving Supply to Marketplaces: Optimal Platform Pricing when Suppliers Share Inventory

Victor Martínez-de-Albéniz[1] • Catarina Pinto[2] • Pedro Amorim[3]

Abstract

Problem definition: Marketplace platforms such as Amazon or Farfetch provide a convenient meeting point between customers and suppliers and have become an important element of e-commerce. This sales channel is particularly interesting for suppliers that sell seasonal goods under a tight time frame, because they provide expanded reach to potential customers, even though it entails lower margins. In this dyadic relationship, a supplier needs to optimize when to share inventory with the platform, and the platform needs to set the right commission structure during the season.

Academic / Practical Relevance: We characterize supplier participation into the platform in a dynamic setting, and to link it to inventory levels, demand rates, time left in the season, and commission structure. This directly drives the pricing decision made by the platform. We thus provide a framework to evaluate platform pricing policies taking into account supplier responses.

Methodology: We use an optimal control framework with limited inventory supply and a stochastic demand process. We study the conditions under which the supplier will accept to participate and use the platform as a sales channel. We also study the optimal static and dynamic commission structure that the platform should employ.

Results: We find that it will only participate if inventory is high relative to the time left to sell the items. As a result, the platform can only offer limited supply at the beginning of the season. Given this behavior, we find that the platform and the system are always better off with dynamic pricing, while the supplier prefers static pricing. Interestingly, when the inventory decision is contingent on the platform pricing policy, the platform will often find it beneficial to commit to a static price to incentivize the supplier to stock up, highlighting that inability to commit to fixed commissions may destroy value through double marginalization effects.

Managerial Implications: Our work suggests that platforms should commit to static commission structures, so as to provide the incentive for suppliers to stock up, and refrain from dynamically adjusting the commission even though these are better in the short term.

Submitted: June 30, 2020

Keywords: Marketplaces, e-commerce platforms, revenue management, inventory sharing, pricing, optimal control.

[1]valbeniz@iese.edu. IESE Business School, University of Navarra, Av. Pearson 21, 08034 Barcelona, Spain.

[2]cpires@iese.edu. IESE Business School, University of Navarra, Av. Pearson 21, 08034 Barcelona, Spain. Catarina Pinto's research was supported by FCT scholarship.

[3]pamorim@fe.up.pt. INESC TEC, Faculdade de Engenharia da Universidade do Porto, R. Dr. Roberto Frias, Porto 4200-465, Portugal.

1. Introduction

Since the development of the internet in the 1990s, online selling channels have taken various forms, from the vertically integrated direct sales model (Dell as an early innovator for this business model, or Amazon), to the extremely decentralized peer-to-peer sites (eBay). Each configuration has merits and shortcomings. For example, more integrated models allow the platform to generate higher efficiencies in the form of reduced costs (efficient logistics from scale and optimization of product flows) or faster lead-times (visibility and control over inventory). At the same time, more decentralized models provide more flexibility (lighter assets) and faster scalability (no need to invest in inventory to grow the product assortment). The marketplace model has emerged as a reasonable compromise between these extremes: a marketplace platform is typically in charge of customer acquisition through marketing activities; enlisting vendors to incorporate their products to the platform assortment; and collecting orders and payments, that can then be fulfilled by either the platform, e.g., Fulfillment by Amazon, or the vendor. With this type of organization, the platform is able to quickly offer enormous assortments, without the corresponding inventory investments, thereby making the platform very attractive to customers. For example, the luxury fashion marketplace Farfetch offered 355,000 different stock-keeping units (SKUs) from boutiques and brands in 2018 (Cheng 2018). The marketplace approach allowed Farfetch to grow much faster than competitor Yoox Net-A-Porter, with high profitability in the form of 25-30% commission on sales (Ferraro et al. 2019). The success of Amazon has also been fueled by the marketplace model: in 2019, 53% of Amazon sales came from third-party vendors, for which it takes a commission of 6-25% (Johnson 2020). Other examples of marketplaces are online travel agents, that simply include hospitality vendors in their platform (in this case supply takes the form of a perishable capacity), and collect a commission of 15% at Booking (Booking.com 2020) vs. 14-20% at Airbnb (Airbnb Help Center 2020); food delivery services, that include restaurants as vendors, for a commission of 30-35% at Deliveroo or UberEats (Shead 2020); or transportation platforms such as Uber or Didi Chuxing, who charge commissions of 12-28% (Derwin 2020).

One of the reasons behind the success of the marketplace model is that, for suppliers, joining the marketplace platform can be as easy as installing a smartphone application or registering on a website. These involve very few commitments: from a financial perspective, almost all the fees to the platform are variable costs that are only paid whenever revenue is generated for the supplier. As a result, the main consideration for a supplier is to decide whether to reserve the possibly limited inventory for its own channels, or to open it up to the platform (and if there are multiple choices, which one), with the risk of generating a lower margin for the sale. When the supplier decision can evolve over time, e.g., when inventory is sold over months, changing sharing decisions creates

high supply uncertainty for the platform, and as a result strongly impacts operations and platform profits. As a manager of one of the online fashion marketplaces that we interacted with puts it: "We have witnessed that more [inventory] depth in supply translates, almost linearly, into more sales". Clearly, platforms are worried about their own supply, which comes from supplier's sharing decisions. These can be influenced by the platform through the commission structure they propose to their vendors. But designing those can be challenging: for instance, rewards for drivers were mentioned as one of the main drivers of the Uber's lack of profitability, and have accounted for a substantial portion of the losses reported after its initial public offering (Rapier and Wolverton 2019). For these reasons, we need to better understand how platform-supplier interactions shape outcomes, so as to prescribe good decisions to both the supplier and the platform. Specifically, when should the supplier share its inventory, and how does this depend on the incentives given by the platform? Which commission structure should we expect? How does it impact profitability for platform and supplier? These are the research questions that we tackle.

In this paper, we focus on situations where the inventory is limited and perishable: short life-cycle goods such as fashion or electronic equipment in an e-commerce site, or hotel capacity in an online travel agency. Our work is less applicable to marketplaces where there is no forward selling of inventory, such as food delivery or transportation, because in those settings the supplier sharing decision typically does not change over time. In our setting, suppliers can commit their inventory to the platform, and, in the blink of an eye, can overcome the geographical limitations imposed by a physical store, at the expense of a percentage of the revenue. When products depreciate over time, there is pressure to sell items early, while customers are willing to purchase the items at full price. Depending on factors such as the number of units available or timing, the supplier will experience different payoffs whether it decides to keep the inventory to itself or to make it available to a platform. Hence, suppliers will adopt different inventory sharing strategies throughout the selling season. We first formulate the supplier's inventory sharing problem as a continuous-time optimal control problem, and characterize the optimal solution: the supplier should divide the season into two distinct periods. Upon the start of the season, the supplier should restrain sales only to its own direct channel (e.g., its physical stores). Sharing shall occur only after a specific time point is reached, which decreases on the available inventory level and increases in the margin taken from the transaction by the platform.

This optimal sharing policy is, however, not convenient for the platform, because it implies that few items will be available at the beginning of the season (only those in an over-stock situation). We thus study the pricing problem for the platform, i.e., its optimal commission structure, taking into account the supplier's best response. The platform can either set a static or dynamic policy, in which commission is constant or can vary over time respectively. We compare both options and

show that dynamic pricing maximizes the expected revenue captured by the platform during the season, as well as the total profits of the system (sum of platform and supplier profits). However, the supplier is better off under static pricing. Interestingly, under the optimal dynamic pricing policy, all the upside from an expanded market reach is captured by the platform, so the incentive to stock up is reduced for the supplier. As a result, when inventory decisions are endogenous, our finding can be reversed: it may be optimal for the platform to commit to a fixed margin over the season, so as to incentivize the supplier to buy higher inventory quantities and increase profits for both supplier and platform. This suggests that "predictable" platforms that commit to fixed commission fees over the entire season may be more profitable because, even though they suffer from limited supplies at the beginning of the season, they are able to convince their suppliers to take more inventory risks thereby increasing overall supply to the platform.

The paper makes several contributions to the literature. To the best of our knowledge, our work is the first to characterize supplier participation into a platform in a dynamic setting, and to link it to inventory levels, demand rates, time left in the season, and commission structure. Besides being valuable to help suppliers engage with marketplaces, this characterization is also helpful for further studies of marketplaces. Specifically, it can be used to control for the endogenous participation decision of the supplier, which is a typical limitation of empirical models studying online platforms. Note that alternative methods require complex structural models, which assume a fixed choice structure (see for instance Ming et al. 2019). Second, our pricing results show that dynamic pricing is superior for the platform, but can have unintended negative consequences on supplier incentives to carry inventory. This suggests that platforms must be careful in using dynamic commissions with their vendors, and highlights the dangers of dynamic decisions when strategic behaviors are present (Calvo and Martínez-de Albéniz 2016, Bandi et al. 2018).

The rest of the paper is organized as follows. In §2, we review the literature related to our research. In §3, we present the supplier decision problem and characterize the optimal inventory sharing strategy. Section 4 analyzes the optimal pricing policy for the platform, while §5 discusses the impact of such policy on procurement. We finalize our paper in §6 with a discussion of our findings.

2. Literature Review

Our work is related to several lines of research. First, we adapt existing inventory and revenue management solution techniques, which we use to analyze the supplier's inventory sharing problem. Second, we analyze incentives in two-sided marketplaces. Third, our work is an instance of supply chain games.

In our analysis, we first take the supplier's perspective, in which, for each inventory level and time left in the season, a decision must be made about sharing or not. This can be interpreted as a revenue management problem with a binary decision space, for which there is a rich literature: the seminal work of Gallego and Van Ryzin (1994) finds optimal prices within a given time window and different leftover capacity. As time passes, sellers should decrease prices to increase demand, except after a sale at which point the optimal price jumps up. Markdown optimization provides similar insights (Bitran et al. 1998, Smith and Achabal 1998, Caro and Gallien 2012). In this work, optimal inventory sharing has a similar structure, as explicitly discussed in §3. In comparison with these works, we use a decomposition technique of Axsäter (1990) that allows us decompose the complex inventory problem into a set of optimal stopping time sub-problems. This approach has been used in the inventory literature to integrate variations in the state of the system, such as multiple echelons (Muharremoglu and Tsitsiklis 2008, Berling and Martínez-de Albéniz 2016a), cost changes (Berling and Martínez-de Albéniz 2011) or expediting (Berling and Martínez-de Albéniz 2016b). Specifically, we are able to recursively determine the structure of the optimal policy for the binary variable (to share or not to share), as these works also do.

After understanding the supplier perspective, we look at the platform-buyer relationship, within a marketplace. Two-sided marketplaces have been an active topic in recent years (Benjaafar and Hu 2020, Chen et al. 2020). Many models have been developed to study the supply-demand matching practices, especially in ride-hailing settings. Different angles have been provided on the question, including the characterization of prices and wages (Cachon et al. 2017, Taylor 2018, Bai et al. 2019, Hu and Zhou 2019), spatial and inter-temporal equilibria (Bimpikis et al. 2019, Banerjee et al. 2015, Hu et al. 2018, Afeche et al. 2018, Besbes et al. 2018, 2019), endogenous choices (Ming et al. 2019) or driver behavior (Benjaafar and Hu 2020, Chen et al. 2019). These works typically involve competition between agents, while this paper considers the behavior of one agent in the supply side, a necessary ingredient of market equilibrium models. In contrast with on-the-spot matching situations typical in ride-hailing, in our case we need to develop the temporal dimension in detail, while models in the literature either consider steady-state behaviors or two-period situations (Hu et al. 2018). Furthermore, while in this literature prices are endogenous, in our case prices tend to be fixed and constant, and product availability is endogenously determined. Models that focus specifically on retail marketplaces have also been developed. Ryan et al. (2012) discuss coordination issues and potential conflict between marketplaces and direct channels, in a stylized single period model. Ferreira et al. (2016) or Martínez-de-Albéniz et al. (2017) provide detailed studies about demand forecasting and pricing for the marketplace, although in these works supply is taken as an input of the model, while in our case it depends on supplier actions.

Finally, our work is also related to the broad research on supply chain games (Cachon 2003).

Specifically, we optimize platform prices given that the supplier responds optimally, as in Lariviere and Porteus (2001) and Cachon and Lariviere (2005) but in a multi-period setting. Supplier-buyer multi-period negotiations have been studied too, see Martínez-de Albéniz and Simchi-Levi (2013) and references therein: in those works, the emphasis is on the inter-temporal strategic interaction between agents. In contrast, in this paper we focus on open-loop policies, resulting in a simple best-response function of the supplier, even though it involves solving a complex optimal control problem.

3. Supplier Inventory Sharing Strategy

We start by modeling the inventory sharing decision of the supplier. Using optimal control theory, we find a closed-form solution by solving a Hamilton-Jacobian-Bellman (HJB) equation (Bertsekas 1995). We demonstrate our findings in a real-world scenario from the luxury fashion industry and discuss the impact of the supplier's optimal strategy for the platform.

3.1 Model Formulation

We consider a dyadic supply chain with a supplier and a platform. The supplier owns a direct channel where inventory is located and is directly sold to consumers. In addition, the platform operates a marketplace model, such as Farfetch (Ferraro et al. 2019), where the supplier can make inventory available at any time, and if the item is sold the platform retains a commission.

There is a finite sales window of duration T, that we call the *season*, e.g., the Spring-Summer fashion season between early January and end of June. During the season, the item is sold at a price π, in both channels, which is a standard practice to avoid channel conflicts, e.g., the manufacturing recommended price in the fashion industry. After the season ends, the supplier gets rid of the left-over items at a salvage value v, which is assumed to be lower than the sourcing cost c, and to simplify the exposition assumed to be $v = 0$. We assume that the starting inventory available to the supplier is fixed and equal to N units, although this assumption will be relaxed in §5.

Demand is generated by a Poisson process in every channel: we denote λ_D the arrival rate of customers to the direct supplier channel, and λ_P that of customers to the platform. For simplicity we assume that the sales intensity is constant across the season, but the analysis can be extended to situations where λ_D and λ_P vary over time, provided that the demand process remains memoryless.

When a sale occurs in the direct channel, the supplier captures the full margin π. In contrast, when the sale occurs in the platform, the supplier retains $(1 - \theta)\pi$ and the platform captures $\theta\pi$, with $\theta \in [0, 1]$ denoting the commission fee or platform margin. Because selling through the platform expands the demand rate, the supplier faces the following trade off: it can decide not to

share inventory and reserve it for its own direct channel, capturing the full margin; or it can make it available to the platform as well as its own channel, hence selling in both channels but sharing some of the margin with the platform in case of an online sale. This setting is depicted in Figure 1.

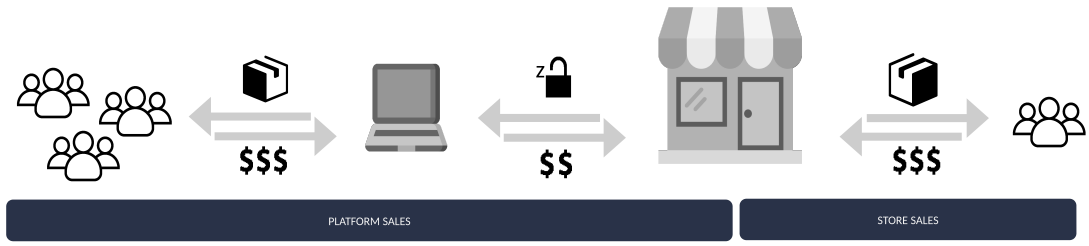

Figure 1: The supplier can decide to sell only at its own direct channel ($z = 0$) or also share its inventory with the platform, adding an extra sales channel ($z = 1$). When sharing inventory, the supplier increases the total sales volume at the expense of a lower margin in each online sale.

We can now formulate the *Inventory Sharing* problem as follows. Let t be the time-to-go to the end of the season, i.e., time is counted backwards: in a season of length T, $t = T$ at the start of the season and $t = 0$ at the end. Let z_t be a binary variable corresponding to the supplier inventory sharing decision at t: the supplier can decide to *not share* ($z_t = 0$) or to *share* ($z_t = 1$) its inventory with the platform.

As the system is Markovian, it is sufficient to consider inventory sharing policies of the form $z_t = \mathbf{z}(t, N_t)$ where N_t is the state of the system at time t, i.e., how many inventory units are left, with N inventory units at the season start. Given an *Inventory Sharing* policy z and a sales horizon $t > 0$, the expected revenue captured by the supplier can be formulated as follows:

$$\sup_{\mathbf{z}} \mathbb{E}\left[\int_0^T e^{-r\tau}\Big(\pi d\Lambda_{D,\tau} + (1-\theta)\pi \mathbf{z}(\tau, N_{T-\tau}) d\Lambda_{P,\tau}\Big) d\tau\right] \tag{1}$$

where r corresponds to a discount rate and $\Lambda_{D,\tau}$ and $\Lambda_{P,\tau}$ are the (stochastic) counting process of customers arriving at the store (direct channel) or the platform during the interval $[0, \tau]$, respec-

tively; and $d\cdot$ represents the differential operator. Starting from $N_T = N$, inventory is updated when $N_\tau > 0$ by $dN_{T-\tau} = -d\Lambda_{D,\tau} - \mathbf{z}(\tau, N_{T-\tau})d\Lambda_{P,\tau}$.

The formulation of the *Inventory Sharing* problem is a variation of *Dynamic Pricing* optimal control. Indeed, the binary decision z_t in the *Inventory Sharing* problem can be related to the effective dynamic price $\hat{p}_t = \frac{\pi\lambda_D + (1-\theta)\pi z_t \lambda_P}{\lambda_D + z_t\lambda_P}$ in a revenue management context, e.g., in Gallego and Van Ryzin (1994). In our case, the average margin is equal to the ratio between the total revenue earned by the supplier and the demand rate (with or without platform), while the demand rate can take two values; in *Dynamic Pricing*, the average margin can be controlled and the demand rate is a function of it. This analogy is summarized in Table 1.

	Average Margin	Demand rate
Inventory Sharing	$\frac{\pi\lambda_D + (1-\theta)\pi z_t \lambda_P}{\lambda_D + z_t\lambda_P}$	$\lambda_D + z_t\lambda_P$
Dynamic Pricing	p_t	$\lambda(p_t)$

Table 1: Comparison between *Inventory Sharing* and *Dynamic Pricing*.

3.2 General Solution Procedure

The general problem described in Equation (1) can be solved using optimal control, by formulating the problem recursively through the value function $V_n(t)$, which is defined as the maximum expected revenue made by the supplier when n units are available to sell from period t to the end of the season, i.e.,

$$V_n(t) \doteq \sup_{\mathbf{z}} \mathbb{E}\left[\int_0^t e^{-r\tau}\Big(\pi d\Lambda_{D,\tau} + (1-\theta)\pi\mathbf{z}(\tau, N_{t-\tau})d\Lambda_{P,\tau}\Big)d\tau\right] \tag{2}$$

where we start from $N_t = n$, and inventory is updated when $N_\tau > 0$ by $dN_{t-\tau} = -d\Lambda_{D,\tau} - \mathbf{z}(\tau, N_{t-\tau})d\Lambda_{P,\tau}$. The closed-form expression for the revenue-to-go function can be found by solving the HJB equation (cf. Appendix):

$$\frac{dV_n}{dt} = \sup_{z\in\{0,1\}}\left[\Big(\pi\lambda_D + z(1-\theta)\pi\lambda_P\Big) - \Big(\lambda_D + z\lambda_P\Big)\Delta V_n(t) - rV_n(t)\right] \tag{3}$$

where $\Delta V_n(t) \doteq V_n(t) - V_{n-1}(t)$. From Equation (3) we can immediately observe that it is only optimal to share inventory if the change in expected revenue when sharing an extra inventory unit ($z = 1$) is higher than the change in expected revenue when not sharing ($z = 0$), i.e., when $(1-\theta)\pi \geq \Delta V_n(t)$. In order to solve Equation (3) in closed form, we first introduce some general observations that will simplify our exposition.

Lemma 1. *At time t, the revenue-to-go is increasing in the number of inventory units, i.e., $V_{n+1}(t) \geq V_n(t), n \in \{1, \ldots, N\}, t \geq 0$.*

The intuition (and proof) behind the lemma is quite simple: as the supplier has one more unit to sell, it can apply the same sharing policy and still have an extra unit that can be sold, thereby increasing the revenue-to-go.

Lemma 2. *At time t, the marginal revenue-to-go is decreasing in the number of inventory units, i.e., $\Delta V_n(t) \geq \Delta V_{n+1}(t), n \in \{1, \ldots, N\}, t \geq 0$.*

This lemma shows that the revenue-to-go is concave in the number of units available. Indeed, because the potential demand is limited, the sales probability decreases with the number of units, thereby making marginal revenue decreasing.

Lemma 3. *Given a number of units n, the revenue-to-go is increasing with the time left until the end of the season, i.e., $V_n(t_2) \geq V_n(t_1), n \in \{1, \ldots, N\}, t_2 \geq t_1 > 0$.*

This lemma is also intuitive: as more time is available for sale, the supplier can use the same sharing policy and still have some additional time for selling the remaining inventory, thereby increasing the revenue-to-go.

Lemma 4. *Given a number of units n, the marginal revenue-to-go is increasing with the time left until the end of the season, i.e., $\Delta V_n(t_2) \leq \Delta V_n(t_1), n \in \{1, \ldots, N\}, t_2 \geq t_1 > 0$.*

The intuition behind this lemma resembles that of Lemma 2: as more time is available, any additional unit of inventory is more valuable, given that the probability of selling it is higher.

Taken together, the lemmas allow us to solve recursively the HJB equation, leading to the following closed-form structure of the revenue-to-go.

Theorem 1. *For any n:*

- *Since $\Delta V_n(t)$ increases in t, there is an unique t_n where $\Delta V_n(t) = (1-\theta)\pi$. It is optimal to set $z^*(t,n) = 1$ if and only if $t \leq t_n$, and $z^*(t,n) = 0$ otherwise.*
- *For each interval $I_{nj} \doteq [t_{j-1}, t_j]$, with $j \in \{1, \ldots, n\}$, and for $I_{nn+1} \doteq [t_n, \infty)$, we have that for $t \in I_{nj}, j \in \{1, \ldots, n+1\}$, $V_n(t) \doteq V_{nj}(t)$ where $V_{nj}(t)$ can be written as*

$$V_{nj}(t) = C_{nj} + \sum_{i=0}^{n-j} A_{nji} t^i e^{-(\lambda_D + \lambda_P)t} + \sum_{i=0}^{j-2} B_{nji} t^i e^{-\lambda_D t}. \tag{4}$$

Theorem 1 thus completely characterizes the optimal policy of the inventory sharing problem. It identifies for each n, a critical time epoch t_n, such that unit n should only be shared with the platform if the time until the end of the season is lower than t_n. We thus have a simple threshold

policy to guide the sharing decision. Moreover, t_n can be characterized by numerically searching for the solution to $\Delta V_n(t) = (1-\theta)\pi$ (which has a unique solution, as a consequence of Lemma 4; when there is no finite solution then we can set $t_n = \infty$), where V_n and ΔV_n can also be described in closed form using Equation (4). Note that the recursive formulas for A_{nji}, B_{nji} and C_{nj} are provided in the proof of the theorem.

The resulting revenue-to-go structure is depicted in Figure 2. From the graphs, we observe that $V_n(t)$ increases with t (Lemma 3), as well as $\Delta V_n(t)$ (Lemma 4). The time at which $\Delta V_n(t)$ reaches value $(1-\theta)\pi$ is precisely the threshold t_n. Thus, the selling season is divided into two periods: an earlier period ($t \geq t_n$), where the supplier should sell only at the store, and a later period ($t \leq t_n$) where the supplier should share its inventory with the platform. Notice also from the figures that ΔV_n is decreasing in n (Lemma 2), resulting in the following corollary.

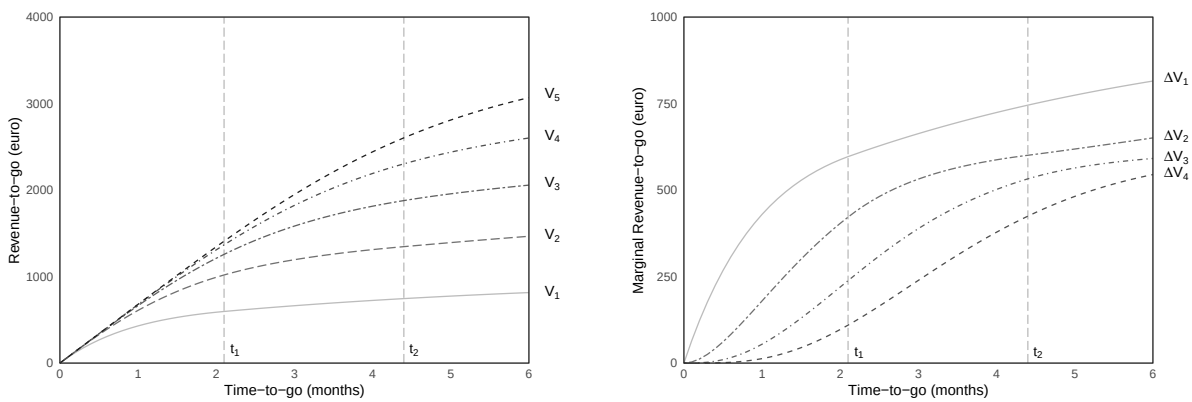

Figure 2: $V_n(t)$ (left) and $\Delta V_n(t)$ (right).

Corollary 1. *For any n, $t_{n+1} \geq t_n$.*

The structure of Theorem 1 is intuitive: early in the season, there is plenty of time to sell an item. Therefore, it is beneficial for the supplier to reserve the limited inventory to the store and receive full item value. As time passes, the opportunities to sell at the store decrease. With less time to sell, the supplier has an incentive to increase demand and match it with the existing capacity, thereby making it optimal to share the inventory with the platform. This behavior is similar to the one found in Gallego and Van Ryzin (1994) and Martínez-de Albéniz and Talluri (2011), where the optimal selling price decreases with time, for any given inventory level.

3.3 Application

We demonstrate our results with an application in luxury fashion retailing. Due to a high seasonality, short selling seasons, and high depreciation of clothes due to the introduction of trends for the new season (Fall/Winter or Spring/Summer), this context fits well our model of a finite selling horizon and Poisson demand (Caro and Martínez-de Albéniz 2015, Caro et al. 2020).

Our base case is based on information available from a real scenario. A store (supplier) carries a dress from a renowned luxury brand, styled according to the Spring-Summer season. An online marketplace platform like Farfetch is a potential second channel where the item can be sold. This selling season spans from the beginning of January to the end of June. Time is measured in months, so $T = 6$. The ticket price of the dress was 2,000€, with a salvage value of about 1,000€ i.e., $\pi = 1,000$€. The customer demand rate at the supplier, λ_D, was approximately 0.2 units per month (direct sale), while the online demand rate, λ_P, was approximately 0.8 units per month (platform sale). Also, during this period, the platform margin was about 400€, i.e., $\theta = 0.40$.

Optimal inventory sharing strategy. We start by applying Theorem 1 to compute t_n. The optimal *Inventory Sharing* policy for our base case is depicted in Figure 3, as well as one simulated scenario of customer arrivals for $N = 2$ and $N = 4$. When $N = 2$, both sales occur in the store, as the optimal policy for $n = 2$ until late January and $n = 1$ until March is to not share. For $N = 4$, the two additional sales occur online.

Table 2: Sensitivity Analysis of base case parameters.

	Base Case	λ_D tests		λ_P tests		θ tests	
π	**1000**	1000	1000	1000	1000	1000	1000
θ	**0.40**	0.40	0.40	0.40	0.40	0.20	0.60
λ_D	**0.2**	0.1	0.3	0.2	0.2	0.2	0.2
λ_P	**0.8**	0.8	0.8	0.5	1.5	0.8	0.8
t_1	**2.14** (Apr 25th)	2.97 (Mar 31st)	1.70 (May 9th)	2.62 (Apr 11th)	1.54 (May 13th)	3.04 (Mar 29th)	1.47 (May 15th)
t_2	**4.38** (Feb 18th)	5.99 (Jan 1st)	3.52 (Mar 15th)	5.42 (Jan 18th)	3.11 (Mar 27th)	5.51 (Jan 15th)	3.47 (Mar 16th)
t_3	>**6** (always)	> 6 (always)	5.27 (Jan 22th)	> 6 (always)	4.63 (Feb 11th)	> 6 (always)	5.49 (Jan 16th)
t_4	>**6** (always)	> 6 (always)	> 6 (always)	> 6 (always)	> 6 (always)	> 6 (always)	> 6 (always)

Table 2 contains the results of a sensitivity analysis on the model parameters λ_D, λ_P and θ. We observe that when any of these parameters increase, it is better to postpone sharing (later date to start sharing, i.e., lower t_n). Indeed, for higher λ_D, the supplier is able to sell more in its direct channel (at higher margins), so sharing inventory with the platform becomes less attractive, hence sharing later. For higher λ_P, the probability of selling online is higher, thus, it is still possible to sell inventory online in shorter time at the end of the season, making later sharing preferable for

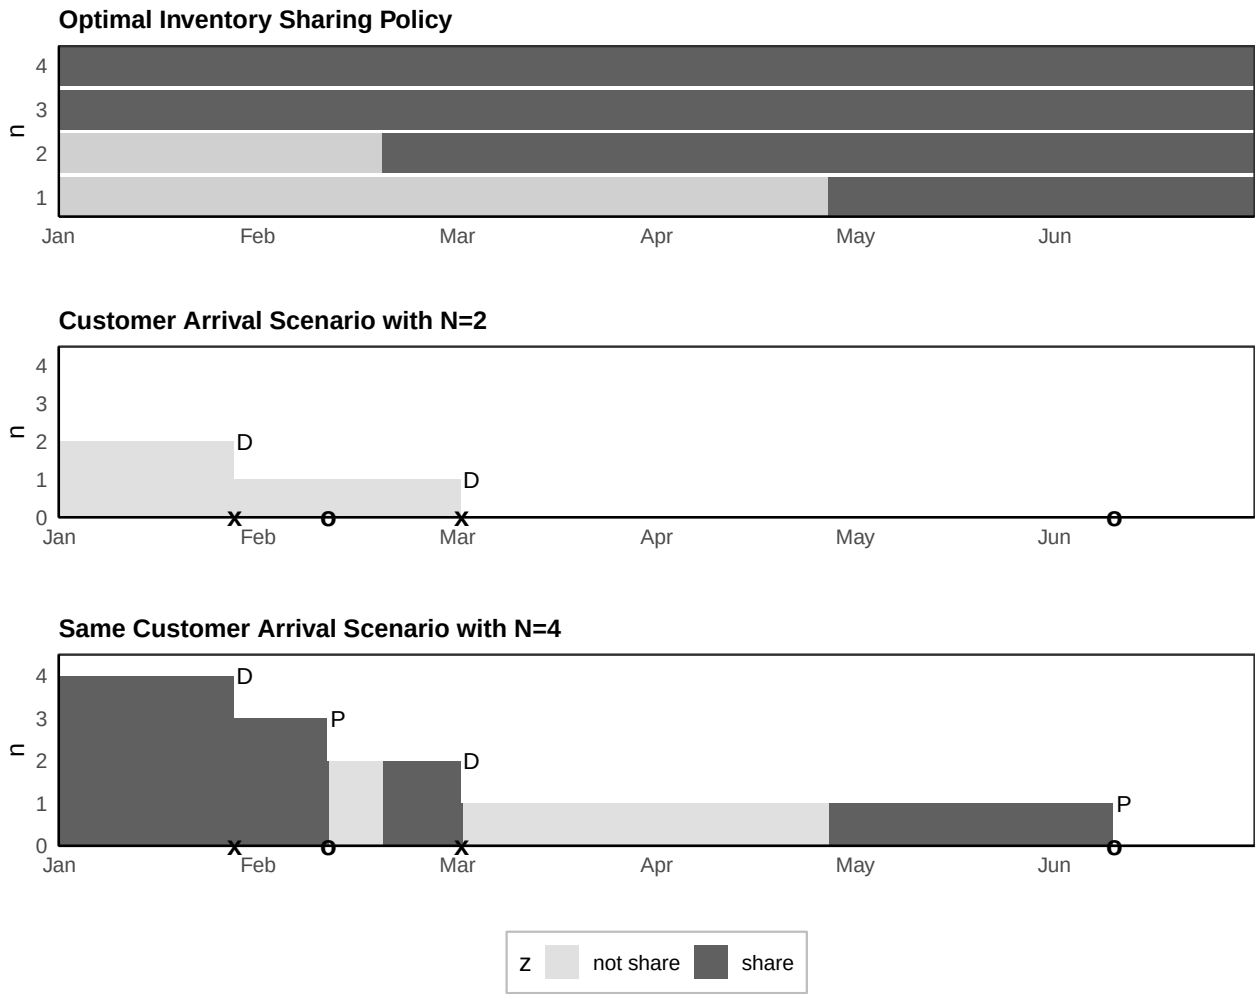

Figure 3: Optimal inventory sharing policy and example of a simulated scenario for $N = 2$ and $N = 4$. The supplier uses the policy depicted in the top graph to decide if it should share or not its inventory with the platform depending on t and n. The letters D and P in the graphs represent sales in the direct and platform channels, respectively. Customer arrivals at the store are marked by an x and at the platform with an o in the x-axis.

the supplier. For higher θ, the supplier keeps a lower margin in each online sale, thus sharing is not as interesting financially, again postponing the date for sharing with the platform.

Incentivizing suppliers to share. Receiving items only towards the end of the season is not in the platform's best interest. Upon such scenario, what can platforms do to encourage suppliers to share their inventory as soon as they receive the products? As explained before, in every online

sale the supplier and the platform share revenue such that the platform retains a margin θ. This margin can be used to reverse engineer the process and encourage inventory sharing throughout the season, i.e., to make the supplier set $z_t = 1$ for all $t \leq T$.

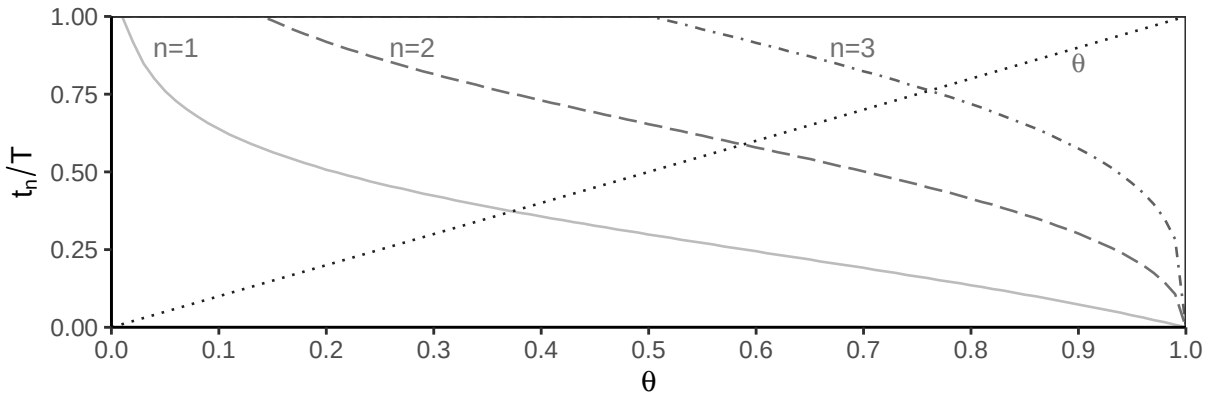

Figure 4: Percentage of sharing time during the season with respect to θ. As θ increases, the supplier shares during a shorter period in the season.

However, the platform faces a trade off between the revenue retained in each sale, $\theta\pi$, and the time during which the supplier shares inventory, $t \in [0, t_n]$. As depicted in Figure 4, as θ increases, the period of online sales is reduced (if $z_t = 0$, no revenue can be earned by the platform as there is no inventory shared and, therefore, no inventory available for sale online). As a result, the platform needs to decide θ so as to maximize the revenue captured, which is a function of the margin retained in each online sale and the time during which inventory is shared by the supplier. In the next section, we analyze this trade off, describing the optimal revenue management strategy for the platform.

4. Platform Optimization: Incentivizing Early Sharing

In order to encourage early-on inventory sharing, the platform can vary the value of the retained margin θ. To this end, the platform can either set a value for θ that is kept throughout the season (static pricing), or make θ time dependent (dynamic pricing).

4.1 Static Pricing (SP)

When establishing a single value of θ for the whole season, the platform needs to trade off the margin retained in each sale with the time during which inventory is shared by the supplier. While

the general structure of the problem is difficult to analyze, we are able to show that, when there is a single unit ($N = 1$), the platform expected revenue is well-behaved and there is a single value of θ^* that maximizes it. Note that this case is practically relevant in many contexts, e.g., luxury fashion marketplaces where suppliers typically have only one unit of a particular SKU in stock.

Theorem 2. *Let Π be the expected revenue captured by the platform over the season. For $N = 1$, Π is unimodal and there is a single value θ^* that maximizes it.*

Theorem 2 thus suggests that the platform can simply increase θ as long as $d\Pi/d\theta > 0$ to find the optimal margin. The intuition behind the result is the following. If $\theta = 0$, $z_t = 1$, and $\Pi = 0$. As θ increases, the retained revenue per unit sold increases, thus Π increases. However, for higher θ, the time span where $z_t = 1$ becomes shorter, as shown in Figure 4, and it turns out that the number of units sold on average (heavily dependent on t_1/T), even if it is neither convex nor concave, decreases sufficiently fast so as to guarantee that Π is unimodal. At some point (θ^*), it reaches a maximum, and afterwards it decreases until Π reaches zero again when $\theta = 1$, when $z_t = 0$ for all $t \in [0, T]$.

Unfortunately, it is difficult to analyze the variation of Π with respect to θ for $N \geq 2$, because the expected number of inventory units sold over the season is hard to compute (given that it depends on the sample paths of customer arrivals in both channels). Therefore, we resort to numerically computing θ^* for higher values of N. Specifically, using the base case parameters from §3.3, we conduct 1,000 simulations of customer arrivals at the direct channel and the platform under different values of N and $\theta \in [0, 1]$, with increments of 0.01. The supplier shares or not its inventory at time-to-go t according to the optimal sharing policy characterized in Theorem 1. We then select the θ^* that maximizes the expected revenue captured by the platform during the season. Figure 5 shows the average number of units sold in the season for $N = 1$ and $N = 4$, and the correspondent revenue captured. We find that the unimodal structure of Π remains: for $N \geq 2$, Π first increases and then decreases, so there is a single value of θ^* that maximizes the platform's expected revenue. In this illustration, $\theta^* = 0.60$ and $\theta^* = 0.82$ for $N = 1$ and $N = 4$, respectively.

4.2 Dynamic Pricing (DP)

Under dynamic pricing, the platform encourages inventory sharing by setting a θ_t for each time epoch t that always makes the supplier better off when sharing. In other words, for each t, it sets the highest possible θ_t so that $z_t = 1$. The following theorem identifies the value of θ_t that maximizes the expected revenue for the platform under such conditions.

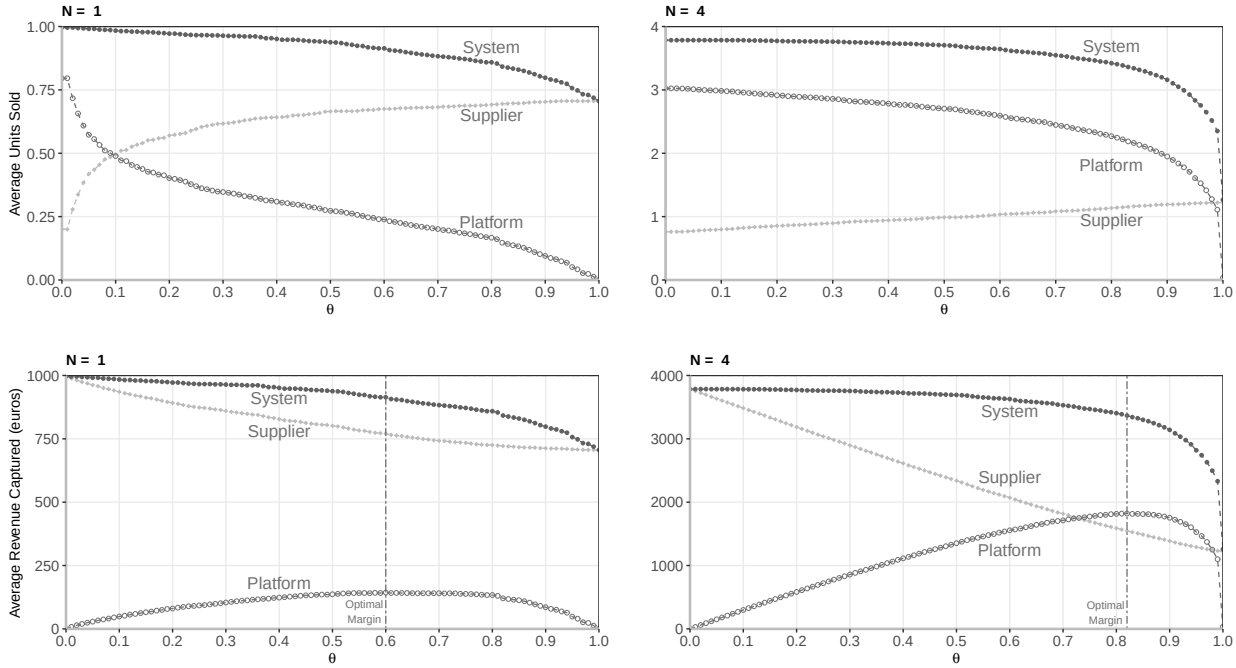

Figure 5: Average result of the 1,000 simulations for $N = 1$ (left) and $N = 4$ (right), to determine the optimal value of θ under static pricing.

Theorem 3. *For any n, let $V_n^0(t)$ be the supplier expected revenue, solution of the Inventory Sharing problem when $\theta = 1$, so that $z_\tau^* = 0$ for all $\tau \leq t$; let $V_n^1(t)$ the solution when $\theta = 0$, so that $z_\tau^* = 1$ for all $\tau \leq t$. As before, define $\Delta V_n^0(t) \doteq V_n^0(t) - V_{n-1}^0(t)$.*

Under a dynamic pricing policy, it is optimal for the platform to set $\theta_t^* = \left(1 - \Delta V_n^0(t)/\pi\right)$. In such case, it is optimal for the supplier to set $z_t^* = 1$ for all t, and its expected revenue over the season is $V_N^0(T)$. The expected revenue captured by the platform is $V_N^1(T) - V_N^0(T)$.

The theorem hence demonstrates that, under dynamic pricing, the optimal value of θ_t^* is achieved when the supplier is always indifferent between sharing and not sharing its inventory. This is intuitive: if the platform sets $\theta_t > \theta_t^*$, the amount earned by the supplier with each online sale is too low so that it is better not to share inventory; Otherwise, if the platform sets $\theta_t < \theta_t^*$, the supplier shares the inventory but platform leaves margin on the table. Therefore, θ_t^* ensures both revenue maximization for the platform and indifference in actions for the supplier. This way, the platform leaves the supplier with the same profits as when there is no online channel, and captures all the additional expected revenue from the additional online sales. Interestingly, while in the static pricing case only limited results for $N = 1$ can be obtained, under dynamic pricing the proof technique is different (based on finding the minimum acceptable price for the supplier) and we are

able to fully characterize the optimal policy.

We depict this result in our base case of §3.3. In Figure 6, we observe that θ_t^* increases with n and decreases with t (recall that t represents the time left in the season, with $t = T$ corresponding to January 1st and $t = 0$ to June 30th), as shown in Lemmas 2 and 4. We see in Figure 7 that as the selling horizon becomes longer, the revenue captured by the platform grows much faster than the one captured by the supplier, for $N = 1$ and 4, except when the probability of selling all the units in the store becomes very high. Finally, Figure 8 illustrates the same insight for different values of starting inventory N, confirming that under dynamic pricing, the platform keeps all the upside from the online channel while the supplier's margin only grows slowly with more inventory.

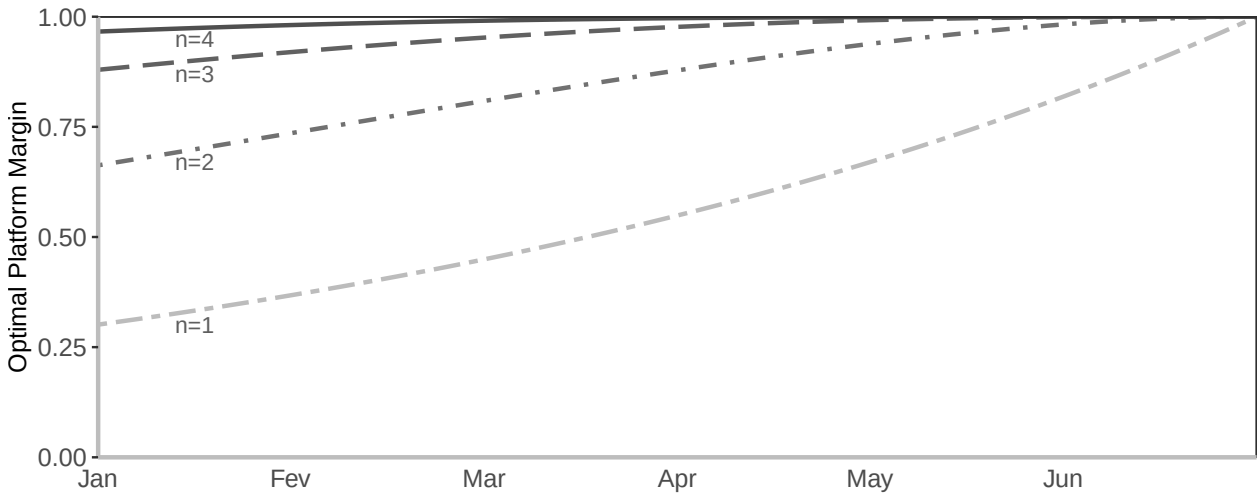

Figure 6: Optimal θ_t^* over time, for $n \in \{1, 2, 3, 4\}$, under dynamic pricing.

4.3 Comparing Pricing Schemes

Table 3 compares the expected revenue captured over the season under dynamic and static pricing policies, for multiple values of N. As expected from the structure of the optimal dynamic pricing policy, it is always better to use dynamic pricing, from the perspective of the platform and the system (i.e., the total revenue for both supplier and platform). Indeed, optimal dynamic prices lead to complete inventory sharing ($z_t^* = 1$ for all t), which maximizes sales for the system, and transfer all the upside from online sales to the platform. In other words, under dynamic pricing, inventory is shared during the entire sales window: inefficiencies, which arise from the "reservation" of inventory for the direct channel, are eliminated. This can be seen as a manifestation of double marginalization, where inventory supply is restricted below the efficient level.

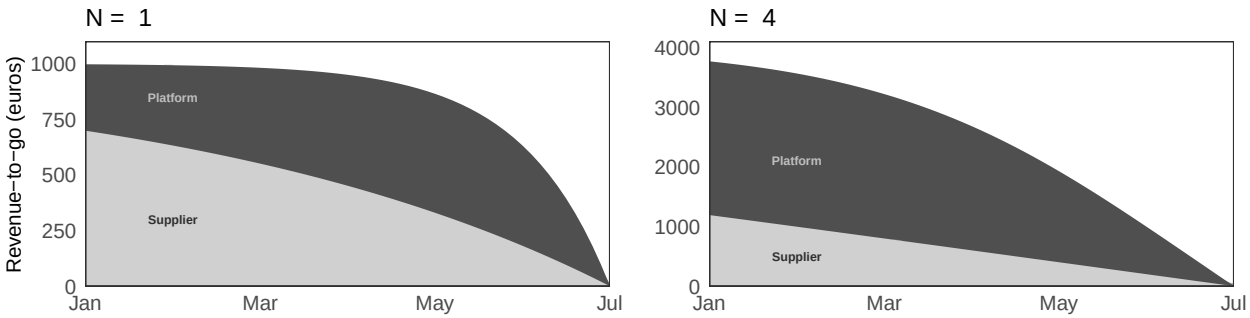

Figure 7: Expected revenue-to-go under DP for $N = 1$ (left) and $N = 4$ (right).

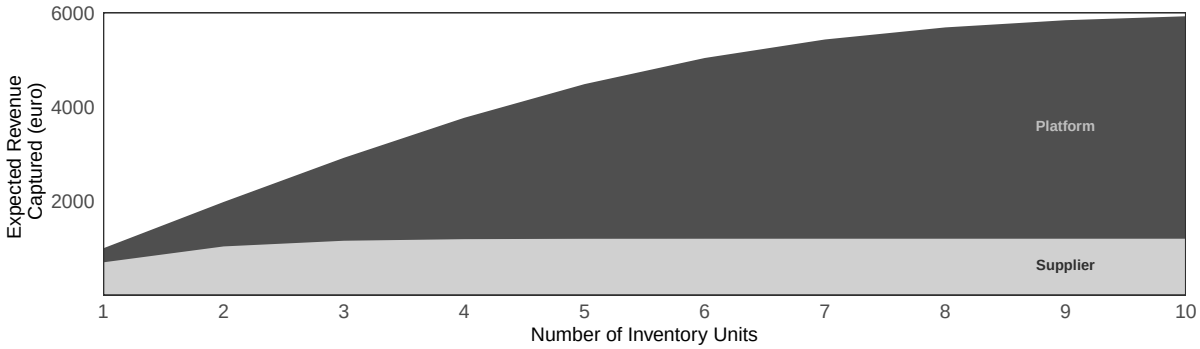

Figure 8: Revenue captured over the season under a DP for different values of N.

From the supplier's perspective, the picture is radically different: under dynamic pricing, the supplier is indifferent between sharing or not, leading to a low expected revenue V_N^0. As a result, the supplier is in fact better off with static pricing. The difference can be significant, for instance when $N = 4$, the supplier captures 1,547€ under static pricing, against 1,190€ under dynamic pricing (23% less).

This suggests that dynamic pricing in the marketplace is superior from an efficiency perspective, even though it transfers value from supplier to platform, which is intuitive because the pricing is indeed set to maximize platform profits.

At the same time, one might be concerned about the negative indirect effects of such value transfer. Interestingly, we see that the supplier profits might not be increasing in the number of inventory units N. Hence, not only the supplier carries all the inventory risk, but it also becomes more vulnerable in front of the platform when inventory is high. For instance, under static pricing, the supplier is better off when committing to only 4 units of inventory (almost 6% more revenue captured than with $N = 10$), showing that, for the supplier, it is not always better to have more

Revenue							
	Dynamic Pricing			Static Pricing			
N	Supplier	Platform	System	θ^*	Supplier	Platform	System
1	699€	299€	998€	0.60	771€	143€	914€
2	1,036€	944€	1,980€	0.68	1,207€	516€	1,723€
3	1,157€	1,762€	2,918€	0.77	1,430€	1,099€	2,529€
4	1,190€	2,577€	3,767€	0.82	**1,547€**	1,820€	3,367€
5	1,198€	3,284€	4,482€	0.88	1,539€	2,551€	4,090€
6	1,200€	3,837€	5,036€	0.91	1,516€	3,181€	4,697€
7	1,200€	4,230€	5,430€	0.93	1,494€	3,692€	5,186€
8	1,200€	4,486€	5,686€	0.94	1,481€	4,054€	5,535€
9	1,200€	4,639€	5,839€	0.95	1,454€	4,294€	5,748€
10	**1,200€**	**4,723€**	**5,923€**	0.95	1,464€	**4,431€**	**5,895€**

Table 3: Expected revenue captured by the platform and the supplier under optimal dynamic pricing and a static pricing policies.

stock. Under such scenario, the platform captures only 1,820€, 41% of what it would capture if $N = 10$. Similarly, the increase of profits with N under dynamic pricing are all positive, but disproportionally small for the supplier, in comparison with the platform, as shown in Figure 8 and Table 3.

In other words, incentives for early-on sharing might have undesirable implications in procurement, as suppliers might be indirectly drawn to keep N small, compared to what would be beneficial for the supply chain. The impact of incentives on procurement will be further explored in the following section.

5. Implications On Procurement

5.1 Endogenous Stocking Quantity

Typically, suppliers will know the pricing conditions of the platform before the start of a season, especially if the relationship is recurrent. In such cases, the supplier will use the information available to determine its own inventory purchasing strategy. For the supplier, committing to the appropriate number of inventory units in the season start is key, especially because under the marketplace model, it bears all inventory risks alone. In this section, we endogenize the optimal inventory decision of the supplier, and discuss how endogenous inventories affect the platform's pricing strategy.

When the platform adopts dynamic pricing, supplier profits is given by $V_N^0(T)$ according to Theorem 3. As a result, the optimal inventory decision under dynamic pricing is given by

$$\max_{N \geq 0} \{V_N^0(T) - cN\} \tag{5}$$

where c is the sourcing cost of the item for the supplier. It turns out that, because the supplier does not capture any of the upside from platform sales, this is a standard newsvendor problem, which results in a critical fractile solution where the largest N such that $Pr[\text{Poisson}_{\lambda_D T} \geq N] \geq c/\pi$.

In contrast, under static pricing, the supplier's decision will directly depend on the commission level θ that the platform decides. The optimal inventory decision in this case can be expressed as

$$\max_{N \geq 0} \{V_N(T|\theta) - cN\} \tag{6}$$

and hence the optimality condition is given by the highest N such that $\Delta V_N(T|\theta) \geq c$. As a result, the platform's static price optimization problem, i.e., setting θ, becomes more difficult because now the platform needs to take into consideration the impact of θ on N^*. As before, the problem is not tractable to provide analytical solutions, so we proceed numerically.

Under static pricing, N^* depends on the value of θ chosen by the platform. In Figure 9 we demonstrate the optimal supplier profit (first graph) and correspondent N^* (second graph) for different values of $\theta \in [0, 1]$. The third and last graphs show the correspondent expected profit for the platform and the system when the supplier commits to $N^*(\theta)$, respectively. When comparing the platform profit under both pricing strategies, we can see that the platform is better off most of the time by choosing static pricing ($\theta \in [0.08, 0.82]$), with maximum profit captured for $\theta^* \approx 0.49$. In comparison, under dynamic pricing the supplier will set $N^* = 1$, leading to an expected profit of 299€ for the platform.

Our findings lead to the following conclusion: while dynamic pricing is efficient when N is given, it hurts the incentives of the supplier to stock high quantities and as a result the platform and the system suffer. In contrast, by committing to a not-too-high commission (in Figure 9, $\theta < 0.82$), the platform is able to guarantee sufficient margin for the supplier that leads to both a higher profit for the supplier, due to higher sales brought by the platform, a higher stocking quantity N^*, and ultimately a higher profit for the platform as well. In other words, while dynamic pricing solves *double marginalization during the season*, by pushing the supplier to share the inventory during the entire sales period, it creates *double marginalization before the season* by restricting stocking quantities, as in Lariviere and Porteus (2001). Under such conditions, the platform is no longer better off using a dynamic pricing policy, but with static pricing.

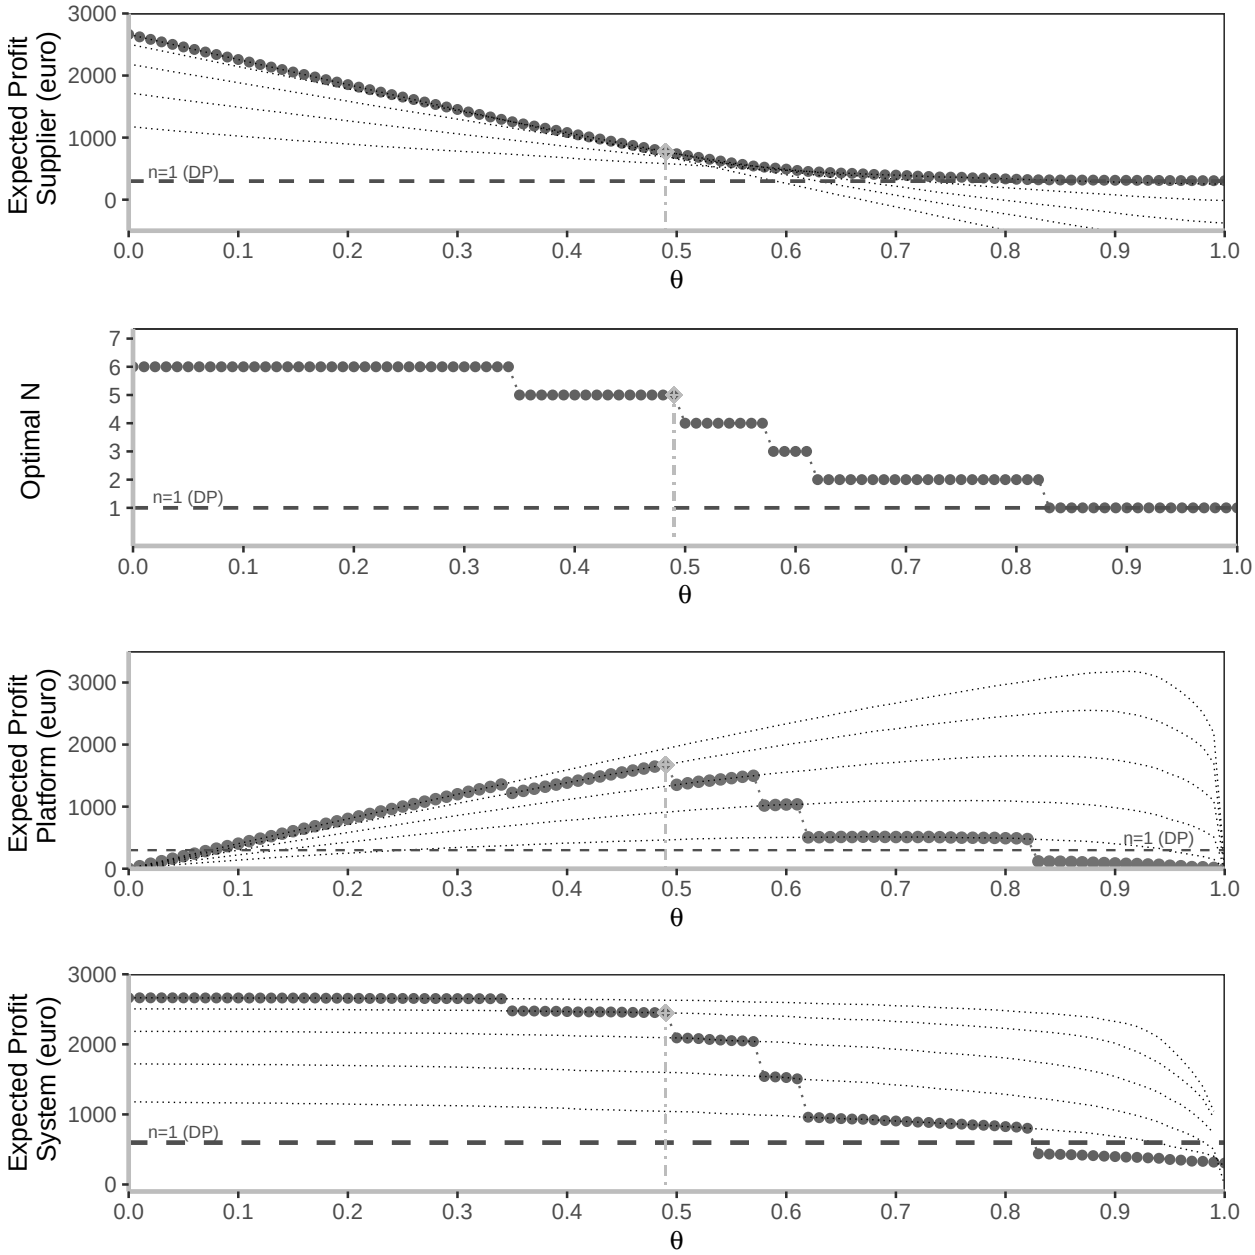

Figure 9: Impact of incentives on procurement. From the top: optimal profit for the supplier, correspondent $N^*(\theta)$, platform profit, and system profit.

5.2 Static vs. Dynamic Pricing

For the platform, choosing a static or dynamic pricing policy depends on the purchasing cost of the supplier ($c = 400$€ per inventory unit in the example of Figure 9) and customer arrival rates

to both channels. We conduct here a series of numerical simulations for 1000 different random scenarios of customer arrivals. For each scenario, we vary the following parameters: $\theta \in [0, 1]$, purchasing cost $c \in [0, 800]$, and $\lambda_P \in [0.2, 1.7]$, while we keep $\lambda_D = 0.5$ and $\pi = 1000$. For each parameter combination, we consider different values of $N \in \{1, \ldots, 10\}$, and compute the expected profits for supplier and platform. Under a static pricing approach, the supplier sets $N^*(\theta)$ so to maximize its profit, followed by the platform, who sets θ^*. Under a dynamic pricing approach, the supplier sets N^* (as the platform terms vary throughout the season as shown in Theorem 3).

Our results are presented in Figure 10, where *Arrival Rates Ratio* corresponds to λ_P/λ_D. We can observe that under static pricing, as the arrival rate ratio increases from 1 (top graph) to 2 (middle graph), the supplier is able to commit to a higher number of inventory units for the same purchase cost, as more units can be sold through the platform. The zig-zag pattern is also easily explainable: as purchase costs increase, supplier's profitability decreases, being better off with lower values of $N^*(\theta)$. In order to prevent the supplier from reducing N^*, the platform reduces θ as purchase costs increase. However, there is a moment when the platform finds it more profitable to let the supplier pick a lower N, so it becomes optimal to set a higher θ. At such moment, the number of inventory units decreases one unit and there is a jump in θ, which decreases again until the next threshold. Under dynamic pricing, the graph presented (bottom graph) can be generalized to all arrival rate ratios. This occurs since, as λ_P increases, the platform retains most of the additional revenue earned. Thus, the supplier's profitability is only affected by the increase in purchase costs.

The optimal policy for different values of purchase costs and arrival rate ratios is depicted in Figure 11. We can observe that, contrary to the results in §4, the platform is most often better off with static pricing, benefiting from dynamic pricing only under low purchase costs and when the customer arrival rate to the platform is lower or similar to the customer arrival rate to the store.

6. Conclusion

In this paper, we build a model to understand the behavior of suppliers in marketplace platforms. Specifically, we analyze the participation decision of the supplier along the sales season, to find that it is optimal to share the supplier's inventory with the platform when inventory level is high compared with the time left in the season. The optimal policy is thus defined by a remaining time threshold below which inventory sharing is optimal (Theorem 1). This implies that the platform only has access to supply late in the season.

To manage this possibly limited supply situation, we study how pricing terms given by the platform affect the sharing decision. Under static pricing, i.e., when the percentage retained by the platform upon a sale is fixed, the platform's optimization problem is well-behaved, at least when

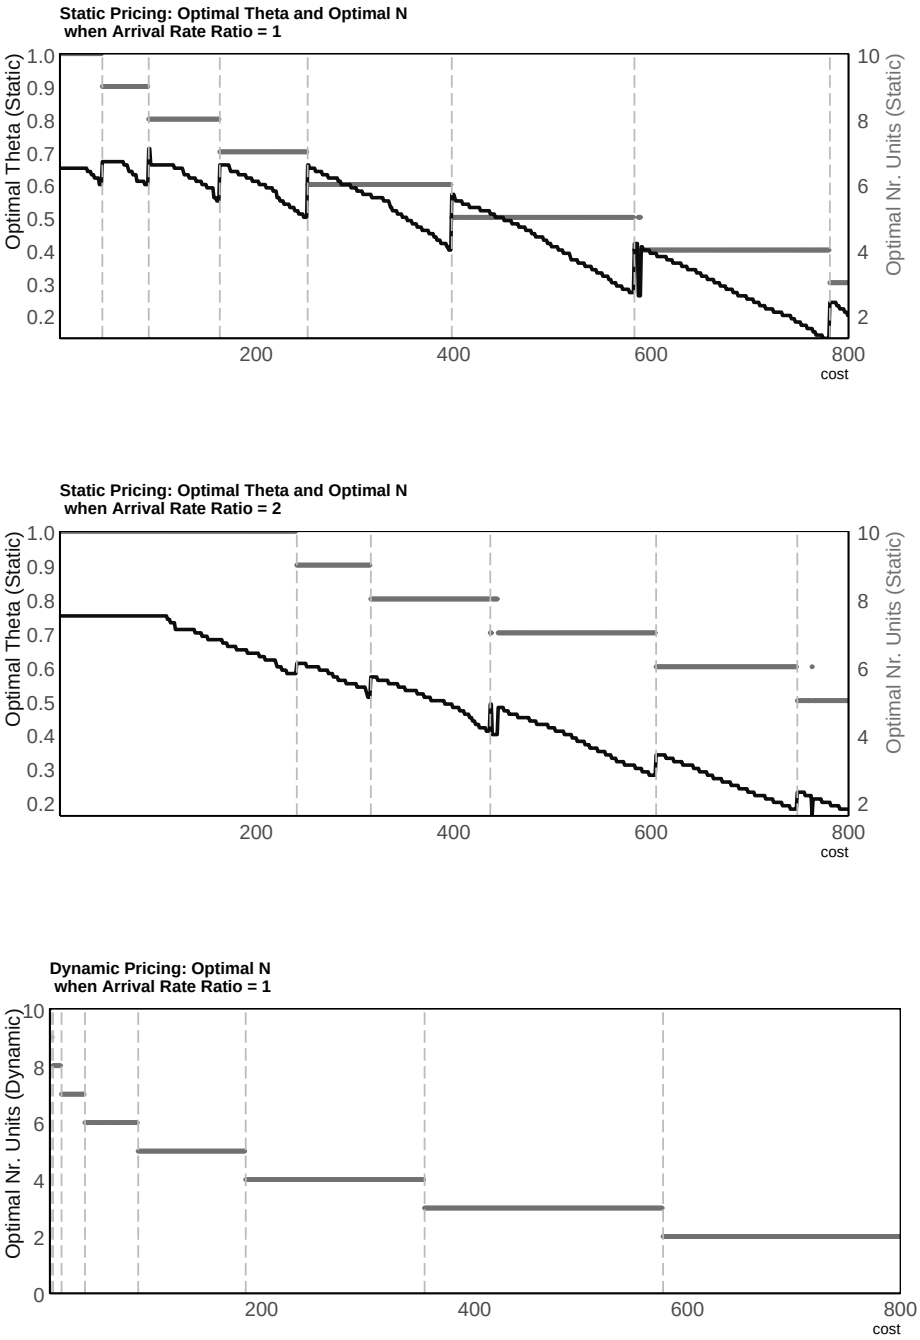

Figure 10: Optimal θ^* (and $N^*(\theta)$) for varying purchase costs, under Static (top and middle) and Dynamic (bottom) Pricing policies.

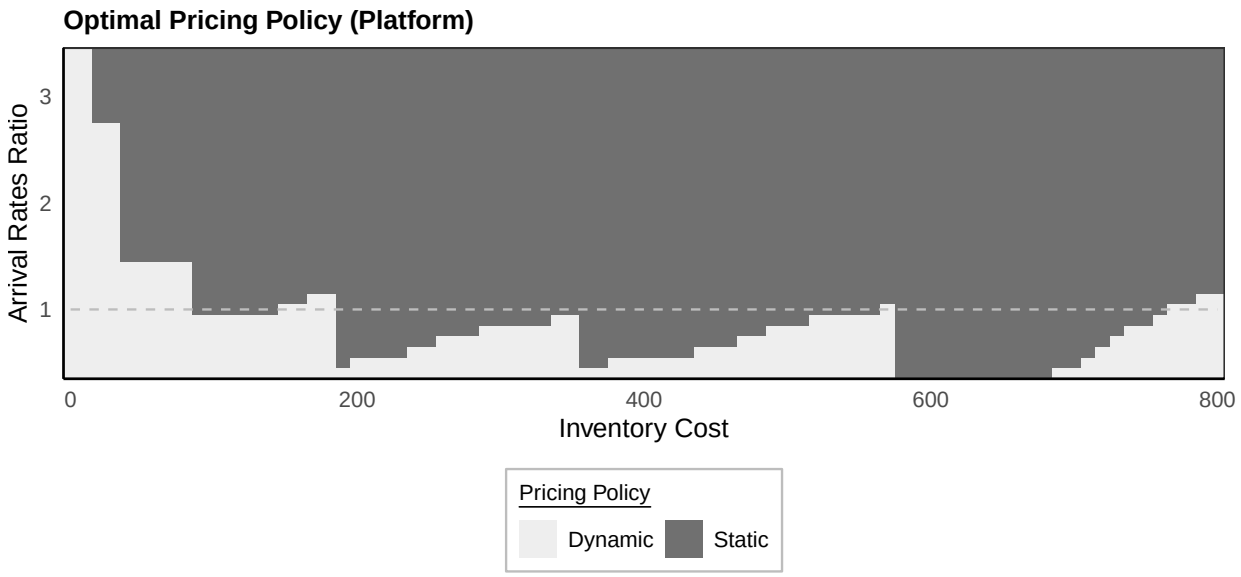

Figure 11: Pricing policy that maximizes supplier's profit.

there is one unit of inventory for sale (Theorem 2). We find this regime less appealing for the platform than dynamic pricing: under a simple dynamic pricing rule that we characterize (Theorem 3), it is possible to incentivize the supplier to share inventory during the entire season, thereby eliminating supply limitation inefficiencies similar to the double marginalization phenomenon, maximizing platform profits and at the same time system profits. We thus show that the platform is always better off using dynamic pricing, while the supplier prefers static pricing.

We finally study the dynamics of the supplier-platform relationship when inventory is endogenous, i.e., when the supplier has the power to set N based on the pricing policy imposed by the platform. The results from this analysis reverse the previous findings. Namely, we show that dynamic pricing creates double marginalization before the season: lack of revenue guarantees to the supplier lead to lower inventory quantities, which protect the supplier from abusive platform commissions, and by the same token hurt the platform because supply and sales are lower than what they could be in a centralized system. We find that, under endogenous inventory, the platform is most of the time better off with static pricing, with two exceptions: when purchase cost or the arrival rate to the platform are low. We thus provide a new example of double marginalization in digital marketplaces, revealing that incentives for early-on inventory sharing can have a negative impact in supplier's procurement levels.

There are a number of research directions worth exploring in the future. Our current work is a useful first step to analyze the temporal aspect of participation in marketplaces of perishable

goods, and can be readily integrated in econometric models studying supply-demand equilibria. Furthermore, our model has focused on a single-product context, and it would be interesting to explore the effect of competition in optimal commission structure, as in Martínez-de Albéniz and Talluri (2011). Last but not least, our model takes the demand side as given, in the form of a Poisson arrival process. It is well known that two-sided markets attract more demand when more supply is available. When the platform demand depends on the inventory sharing decision, the platform may want to reduce its fees, so as to ensure higher supply, leading to higher demand and thus higher profits.

References

Afeche, P., Z. Liu, and C. Maglaras. 2018. Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance. Columbia Business School Research Paper.

Airbnb Help Center 2020, Accessed May 29. What is the Airbnb service fee? https://www.airbnb.ie/help/article/1857/what-is-the-airbnb-service-fee.

Axsäter, S. 1990. Simple solution procedures for a class of two-echelon inventory problems. *Operations research* 38 (1): 64–69.

Bai, J., K. C. So, C. S. Tang, X. Chen, and H. Wang. 2019. Coordinating supply and demand on an on-demand service platform with impatient customers. *Manufacturing & Service Operations Management* 21 (3): 556–570.

Bandi, C., A. Moreno, D. Ngwe, and Z. Xu. 2018. Opportunistic Returns and Dynamic Pricing: Empirical Evidence from Online Retailing in Emerging Markets. Working paper, HBS.

Banerjee, S., C. Riquelme, and R. Johari. 2015. Pricing in ride-share platforms: A queueing-theoretic approach. Available at SSRN 2568258.

Benjaafar, S., and M. Hu. 2020. Operations management in the age of the sharing economy: what is old and what is new? *Manufacturing & Service Operations Management* 22 (1): 93–101.

Berling, P., and V. Martínez-de Albéniz. 2011. Optimal inventory policies when purchase price and demand are stochastic. *Operations research* 59 (1): 109–124.

Berling, P., and V. Martínez-de Albéniz. 2016a. A characterization of optimal base-stock levels for a multi-stage serial supply chain. *Naval Research Logistics (NRL)* 63 (1): 32–46.

Berling, P., and V. Martínez-de Albéniz. 2016b. Dynamic speed optimization in supply chains with stochastic demand. *Transportation Science* 50 (3): 1114–1127.

Bertsekas, D. P. 1995. *Dynamic programming and optimal control*, Volume 1. Athena scientific Belmont, MA.

Besbes, O., F. Castro, and I. Lobel. 2018. Spatial capacity planning. Available at SSRN 3292651.

Besbes, O., F. Castro, and I. Lobel. 2019. Surge pricing and its spatial supply response. Columbia Business School Research Paper No. 18-25. Available at SSRN: https://ssrn.com/abstract=3124571.

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

Bitran, G., R. Caldentey, and S. Mondschein. 1998. Coordinating clearance markdown sales of seasonal products in retail chains. *Operations Research* 46 (5): 609–624.

Booking.com 2020, Accessed May 29. How much commission do I pay? https://partner.booking.com/en-gb/help/commission-invoices-tax/how-much-commission-do-i-pay.

Cachon, G. P. 2003. Supply chain coordination with contracts. *Handbooks in operations research and management science* 11:227–339.

Cachon, G. P., K. M. Daniels, and R. Lobel. 2017. The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management* 19 (3): 368–384.

Cachon, G. P., and M. A. Lariviere. 2005. Supply chain coordination with revenue-sharing contracts: strengths and limitations. *Management science* 51 (1): 30–44.

Calvo, E., and V. Martínez-de Albéniz. 2016. Sourcing strategies and supplier incentives for short-life-cycle goods. *Management Science* 62 (2): 436–455.

Caro, F., and J. Gallien. 2012. Clearance pricing optimization for a fast-fashion retailer. *Operations Research* 60 (6): 1404–1422.

Caro, F., A. G. Kök, and V. Martínez-de Albéniz. 2020. The future of retail operations. *Manufacturing & Service Operations Management* 22 (1): 47–58.

Caro, F., and V. Martínez-de Albéniz. 2015. Fast fashion: Business model overview and research opportunities. In *Retail supply chain management*, 237–264. Springer.

Chen, X., Z. Li, L. Ming, and W. Zhu. 2019. The Incentive Game under Target Effects in Ridesharing: A Structural Econometric Analysis. Available at SSRN 3479675.

Chen, Y.-J., T. Dai, C. G. Korpeoglu, E. Körpeoğlu, O. Sahin, C. S. Tang, and S. Xiao. 2020. Innovative online platforms: Research opportunities. *Manufacturing & Service Operations Management* 22 (3): 430–445.

Cheng, A. 2018. What Farfetch's IPO Filing Says About The $300 Billion Luxury Fashion Industry. *Forbes* August 21:online.

Derwin, J. 2020. Uber competitor DiDi's new reward system has slashed pay rates for Australian drivers, as they work punishing hours and battle it out for limited jobs. *Business Insider* January 10:online.

Ferraro, F., A. Baldo, and I. Sastre Boquet. 2019. Digital Online Competition in Fashion and Luxury Goods: YNAP versus Farfetch. Technical report, IESE Case SM-1656-E.

Ferreira, K. J., B. H. A. Lee, and D. Simchi-Levi. 2016. Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management* 18 (1): 69–88.

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

Hu, B., M. Hu, and H. Zhu. 2018. Surge pricing and two-sided temporal responses in ride-hailing. Available at SSRN 3278023.

Hu, M., and Y. Zhou. 2019. Price, wage and fixed commission in on-demand matching. Available at SSRN 2949513.

Johnson, T. 2020. How To Sell on Amazon in 2020: Setup, Fees, & Strategy. *Tinuiti* March 2:online.

Lariviere, M. A., and E. L. Porteus. 2001. Selling to the newsvendor: An analysis of price-only contracts. *Manufacturing & Service Operations Management* 3 (4): 293–305.

Martínez-de-Albéniz, V., S. Nasini, and A. Planas. 2017. Using Clickstream Data to Improve Campaign Effectiveness in Flash Sales. *Production and Operations Management* forthcoming:NA.

Martínez-de Albéniz, V., and D. Simchi-Levi. 2013. Supplier–Buyer Negotiation Games: Equilibrium Conditions and Supply Chain Efficiency. *Production and Operations Management* 22 (2): 397–409.

Martínez-de Albéniz, V., and K. Talluri. 2011. Dynamic price competition with fixed capacities. *Management Science* 57 (6): 1078–1093.

Ming, L., T. I. Tunca, Y. Xu, and W. Zhu. 2019. An empirical analysis of market formation, pricing, and revenue sharing in ride-hailing services. Working paper, IESE Business School.

Muharremoglu, A., and J. N. Tsitsiklis. 2008. A Single-Unit Decomposition Approach to Multiechelon Inventory Systems. *Operations Research* 56 (5): 1089–1103.

Rapier, G., and T. Wolverton. 2019. Uber lost $5.2 billion in 3 months. Here's where all that money went. *Business Insider* August 9:online.

Ryan, J. K., D. Sun, and X. Zhao. 2012. Competition and coordination in online marketplaces. *Production and Operations Management* 21 (6): 997–1014.

Shead, S. 2020. Restaurant owners in Britain call on Deliveroo to drop commission fees. *CNBC* May 1:online.

Smith, S., and D. Achabal. 1998. Clearance Pricing and Inventory Policies for Retail Chains. *Management Science* 44 (3): 285–300.

Taylor, T. A. 2018. On-demand service platforms. *Manufacturing & Service Operations Management* 20 (4): 704–720.

Proofs

Derivation steps in the definition of the dynamical control problem

One can informally derive the HJB equation by considering what happens over a small time period δ. The supplier sells one item directly to the consumer with probability $\lambda_D\delta$, sells one item through the platform with probability $z\lambda_P\delta$, and no items with probability $1-\lambda_D\delta-z\lambda_P\delta$. Thus, the expected revenue function can be formulated as:

$$\begin{aligned}V_n(t) &= \sup_{\mathbf{z}}\left[\delta\lambda_D\Big(\pi+V_{n-1}(t)\Big)+\delta\mathbf{z}\lambda_P\Big(\pi+V_{n-1}(t)\Big)+\Big(1-\delta\lambda_D-\mathbf{z}\delta\lambda_P\Big)V_n(t)-\delta\frac{dV_n}{dt}-r\delta V_n(t)+o(\delta)\right]\\ &= \sup_{\mathbf{z}}\left[\delta\Big(\pi\lambda_D+\mathbf{z}(1-\theta)\pi\lambda_P\Big)-\delta\Big(\lambda_D+\mathbf{z}\lambda_P\Big)\Big(V_n(t)-V_{n-1}(t)\Big)+V_n(t)-\delta\frac{dV_n}{dt}-r\delta V_n(t)+o(\delta)\right]\end{aligned}$$

Rearranging and taking the limit as $\delta\to 0$ leads to the HJB equation below. By solving this equation, one can find the closed-form expression for the revenue-to-go function.

$$0=\sup_{\mathbf{z}}\left[\Big(\pi\lambda_D+\mathbf{z}(1-\theta)\pi\lambda_P\Big)-\Big(\lambda_D+\mathbf{z}\lambda_P\Big)\Delta V_n(t)-\frac{dV_n}{dt}-rV_n(t)\right]$$

$$\frac{dV_n}{dt}=\sup_{\mathbf{z}}\left[\Big(\pi\lambda_D+\mathbf{z}(1-\theta)\pi\lambda_P\Big)-\Big(\lambda_D+\mathbf{z}\lambda_P\Big)\Delta V_n(t)-rV_n(t)\right]$$

Proof of Lemma 1

With an inventory of $n+1$ items, there is an extra item that can be sold when compared with an inventory of n items. As a result, we can implement the same policy as for n, and when there is one unit left, we can still make additional revenue. Therefore, the expected revenue with $N=n+1$ items is higher than the expected revenue with $N=n$ items, i.e., $V_{n+1}(t)\geq V_n(t), \forall n\in\{1,\dots,N\}, \forall t>0$.

Proof of Lemma 2

We prove this result by contradiction. First, note that for all n, $V_n(0)=0$ so that $\Delta V_n(0)=0$ for all n. In addition, since V_n is the solution to the HJB equation, it is continuously differentiable in t, and so is ΔV_n.

Suppose that there exists $t>0$ and $n<\infty$ such that $\Delta V_n(t)-\Delta V_{n+1}(t)<0$. Let t_0 be the infimum of such t, and n the smallest value satisfying the inequality, i.e., for all $t\leq t_0$, for all n, $\Delta V_n(t)-\Delta V_{n+1}(t)\geq 0$ (since $\Delta V_n-\Delta V_{n+1}$ is continuous). By construction, $\Delta V_n(t_0)-\Delta V_{n+1}(t_0)=0$, and because $\Delta V_n-\Delta V_{n+1}$ is continuously differentiable, it must be the case that $\frac{d}{dt}\Big(\Delta V_n-\Delta V_{n+1}\Big)(t_0)<0$. But at the same time,

$$
\begin{aligned}
\frac{dV_{n+1}}{dt}(t_0) &= \sup_{\mathbf{z}} \left[\Big(\pi\lambda_D + z(1-\theta)\pi\lambda_P\Big) - \Big(\lambda_D + z\lambda_P\Big)\Delta V_{n+1}(t_0) \right] - rV_{n+1}(t_0) \\
&= \sup_{\mathbf{z}} \left[\Big(\pi\lambda_D + z(1-\theta)\pi\lambda_P\Big) - \Big(\lambda_D + z\lambda_P\Big)\Delta V_{n}(t_0) \right] - rV_{n+1}(t_0) \\
&= \frac{dV_n}{dt}(t_0) + rV_n(t_0) - rV_{n+1}(t_0);
\end{aligned}
$$

and

$$
\begin{aligned}
\frac{dV_n}{dt}(t_0) &= \sup_{\mathbf{z}} \left[\Big(\pi\lambda_D + z(1-\theta)\pi\lambda_P\Big) - \Big(\lambda_D + z\lambda_P\Big)\Delta V_{n}(t_0) \right] - rV_{n}(t_0) \\
&\geq \sup_{\mathbf{z}} \left[\Big(\pi\lambda_D + z(1-\theta)\pi\lambda_P\Big) - \Big(\lambda_D + z\lambda_P\Big)\Delta V_{n-1}(t_0) \right] - rV_{n}(t_0) \\
&\quad \text{because } \Delta V_{n-1}(t_0) - \Delta V_n(t_0) \geq 0, \\
&= \frac{dV_{n-1}}{dt}(t_0) + rV_{n-1}(t_0) - rV_n(t_0).
\end{aligned}
$$

This implies that

$$
\frac{d\Delta V_{n+1}}{dt}(t_0) = -r\Delta V_{n+1}(t_0) = -r\Delta V_n(t_0) \leq \frac{d\Delta V_n}{dt}(t_0),
$$

i.e., a contradiction. This means that for all t, for all n, $\Delta V_n(t) - \Delta V_{n+1}(t) \geq 0$. ■

Proof of Lemma 3

The value of an item is maximal at the season start and minimum at its end ($V_n(t=0) = 0$). Thus, for a given number of inventory units n, $V_n(t+1) \geq V_n(t)$, as the expected sum of the revenue captured during the season cannot increase as time-to-go decreases.

Proof of Lemma 4

We prove this result by contradiction. Suppose that there exists $t > 0$ and $n < \infty$ such that $\Delta V_n(t+1) - \Delta V_n(t) < 0$. Let t_0 be the infimum of such t, and n the smallest value satisfying the inequality, i.e., for all $t \leq t_0$, for all n, $\Delta V_n(t+1) - \Delta V_n(t) \geq 0$ (since ΔV_n is continuous). By construction, $\Delta V_n(t_0+1) - \Delta V_n(t_0) = 0$, and thus $\frac{d\Delta V_n}{dt}(t_0+1) - \frac{d\Delta V_n}{dt}(t_0) < 0$. But at the same time,

$$
\begin{aligned}
\frac{dV_n}{dt}(t_0+1) &\geq \frac{dV_{n-1}}{dt}(t_0+1) + rV_{n-1}(t_0+1) - rV_n(t_0+1). \\
&\geq \frac{dV_{n-1}}{dt}(t_0+1) - r\Delta V_n(t_0+1) \\
&\geq \frac{dV_{n-1}}{dt}(t_0+1) - r\Delta V_n(t_0) \\
\frac{d\Delta V_n}{dt}(t_0+1) &\geq \frac{d\Delta V_n}{dt}(t_0), \text{ because } -r\Delta V_n(t_0) \leq \tfrac{dV_n}{dt}(t_0)
\end{aligned}
$$

Which is a contradiction. This means that for all t, for all n, $\Delta V_n(t+1) - \Delta V_n(t) \geq 0$. ■

Proof of Theorem 1

To prove Theorem 1, we show by induction that there is an unique t_n for each value of n and characterize the structure of V_{nj} according to Equation (4).

Let $\lambda = \lambda_D + z\lambda_P$ and $\gamma = \pi\lambda_D + z(1-\theta)\pi\lambda_P$. To initiate the induction, we show the propositions are true for $n = 1$. For t close to zero, $z = 1$ and the solution to $V_1^{'}(t) = \gamma - (\lambda + r)V_1(t)$ is $V_1(t) = \frac{\gamma}{(\lambda+r)} + A_{110}e^{-(\lambda+r)t}$ with $A_{110} = -\frac{\gamma}{(\lambda+r)}$, because $V_1(0) = 0$. As t grows, $\Delta V_1(t) = V_1(t) - V_0(t) = V_1(t)$ increases, so that there exists a unique t_1 satisfying the above equation, defined as $t_1 = \frac{1}{(\lambda+r)}\big(ln\big((\lambda + r) - \theta\lambda_P\big) - ln\big(\theta\lambda_D\big)\big)$ [4].

Hence, for $t \leq t_1$, $V_1(t) = V_{11}(t) = C_{11} + A_{110}e^{-\lambda t}$ with $C_{11} = \gamma/(\lambda + r)$ and $A_{110} = -C_{11}$. For $t \geq t_1$ (i.e, $j > 1$), $A_{1j0} = 0$, $B_{1j0} = [V_{11}(t_1) - C_{1j}]e^{(\lambda_D+r)t_1}$ and $C_{1j} = b/(\lambda_D + r)$, in order for $V_1(t)$ to be differentiable in all its spectrum. Hence the hypothesis is true for $n = 1$.

Assume the induction property is true for $n-1$. For $n \geq 2$, the proposed formulation, in the interval $[t_{j-1}, t_j]$, with $j \geq 1$, satisfies the HJB equation. For t close to zero, $V_{nj}^{'} = \gamma - \lambda\Delta V_{nj}(t) - rV_{nj}(t)$. According to Lemma 4 as t grows, $\Delta V_{nj}(t)$ increases, with $V_{nj}(t) > V_{n-1,j}(t)$, $\forall n$, with $\Delta V_n(t) > 0$. Assume $\Delta V_n(t)$ increases up to t_{n-1}. We aim to show there is a unique point that turns $(1-\theta)\pi \geq \Delta V_n(t)$ into an equality. For $t > t_{n-1}$,

$$\begin{aligned}
\Delta V_{n,n-1}(t)' &= V'_{n,n-1}(t) - V'_{n-1,n-1}(t) \\
&= (\gamma - \lambda\Delta V_{n,n-1}(t) - rV_{n,n-1}(t)) - (\pi\lambda_D - \lambda_D\Delta V_{n-1,n-1}(t) - rV_{n-1,n-1}(t)) \\
&= (\gamma - \pi\lambda_D) - \lambda\Delta V_{n,n-1}(t) + \lambda_D\Delta V_{n-1,n-1}(t) - r(V_{n,n-1}(t) - V_{n-1,n-1}(t)) \qquad (7) \\
&= (\pi\lambda_D + (1-\theta)\pi\lambda_P - \pi\lambda_D) - (\lambda + r)\Delta V_{n,n-1}(t) + \lambda_D\Delta V_{n-1,n-1}(t) \\
&= (1-\theta)\pi\lambda_P - (\lambda + r)\Delta V_{n,n-1}(t) + \lambda_D\Delta V_{n-1,n-1}(t)
\end{aligned}$$

Assume there exists a value of t such that $\Delta V_{n,n-1}(t)' = 0$, therefore $\Delta V_{n,n-1}(t)'' = -(\lambda + r)\Delta V_{n,n-1}(t)' + \lambda_D\Delta V_{n-1,n-1}(t)' = \lambda_D\Delta V_{n-1,n-1}(t)' > 0$. Therefore, $\Delta V_{n,n-1}(t)' > 0$, since, if

[4]

$$\begin{gathered}
\Delta V_1(t) = V_1(t) = \frac{(\lambda_D+\lambda_P(1-\theta))\pi}{(\lambda+r)}\big(1 - e^{-(\lambda+r)t_1}\big) = (1-\theta)\pi \\
1 - e^{-(\lambda+r)t_1} = (1-\theta)\frac{(\lambda+r)}{\lambda_D+\lambda_P(1-\theta)} \\
e^{-(\lambda+r)t_1} = 1 - \frac{(\lambda+r)-\theta(\lambda+r)}{\lambda_D+\lambda_P(1-\theta)} \\
e^{-(\lambda+r)t_1} = \frac{\lambda_D}{\lambda-\theta\lambda_P} \\
t_1 = \frac{1}{(\lambda+r)}\Big(ln\big((\lambda + r) - \theta\lambda_P\big) - ln\big(\theta\lambda_D\big)\Big)
\end{gathered}$$

otherwise, $\Delta V_{n,n-1}(t)''(0) < 0$ when $\Delta V_{n,n-1}(t)' = 0$, which is not true. $\Delta V_{n,n-1}(t)$ is, thus, strictly increasing. Given that $\Delta V_n(t)$ is strictly increasing, there is an unique point t_n at which $(1-\theta)\pi \geq \Delta V_n(t)$ becomes an equality.

We follow with the characterization of $V_n(t)$. For $t \leq t_n$, differentiating Equation (4) leads to:

$$\begin{aligned} V_{nj}^{'} = & \sum_{i=0}^{j-3} \Big(-\lambda_D B_{n,j,i} + (i+1) B_{n,j,i+1} \Big) e^{-\lambda_D t} t^i \; - \lambda_D B_{n,j,j-2} e^{-\lambda_D t} t^{j-2} + \\ & + \sum_{i=0}^{n-j-1} \Big(-\lambda A_{n,j,i} + (i+1) A_{n,j,i+1} \Big) e^{-\lambda t} t^i \; - \lambda A_{n,j,n-j} e^{-\lambda t} t^{n-j} \end{aligned} \tag{8}$$

Simple identification of terms yields that

$$\begin{aligned} & C_{nj} = \frac{\gamma}{(\lambda+r)} C_{n-1,j} + \frac{\gamma}{(\lambda+r)} \\ & A_{n-1,j,n-j} = 0 && \text{, for } i = n-j \\ & A_{n,j,i} = \tfrac{\lambda}{(\lambda+r)i} A_{n-1,j,i-1} && \text{, for } i \in \{1,\ldots,n-j\} \\ & B_{n,j,j-2} = \tfrac{\lambda}{(\lambda+r)-\lambda_D} B_{n-1,j,j-2} && \text{, for } i = j-2 \\ & B_{n,j,i} = \tfrac{\lambda}{(\lambda+r)-\lambda_D} B_{n-1,j,i} + \textstyle\sum_{k=i+1}^{j-2} (-1)^{k-i} k \tfrac{\lambda}{((\lambda+r)-\lambda_D)^{(k+1)}} B_{n-1,j,k} && \text{, for } i \in \{0,\ldots,j-3\}, \text{ and} \\ & A_{nj0} = \Big[V_{n-1,j}(t_{N-1}) - (C_{nj} + \textstyle\sum_{i=0}^{j-2} B_{nji} t_{j-1}^i e^{-\lambda_D t_{j-1}} + \sum_{i=1}^{n-j} A_{nji} t_{j-1}^i e^{-(\lambda+r)t_{j-1}}) \Big] e^{(\lambda+r)t_{j-1}}. \end{aligned}$$

For $t > t_n$, $V_{nj}^{'} = b - \lambda_D \Delta V_n(t) - r V_n(t)$. By differentiation of Equation (4), we get:

$$V_{nj}^{'} = \sum_{i=0}^{j-3} \Big(-\lambda_D B_{n,j,i} + (i+1) B_{n,j,i+1} \Big) e^{-(\lambda_D+r)t} t^i \; - \lambda_D B_{n,j,j-2} e^{-\lambda_D t} t^{j-2} \tag{9}$$

Simple identification of terms yields that

$$\begin{aligned} & C_{nj} = \frac{\pi\lambda_D}{(\lambda_D+r)} C_{n-1,j} + \frac{\pi\lambda_D}{(\lambda_D+r)} = \frac{n\pi\lambda_D}{(\lambda_D+r)} \\ & A_{nji} = 0 \\ & B_{n,j,i} = \tfrac{\lambda_D}{(\lambda_D+r)i} B_{n-1,j,i-1} && \text{, for } i \in \{1,\ldots,j-2\} \\ & B_{nj0} = \Big[V_{n-1,j}(t_{N-1}) - (C_{nj} + \textstyle\sum_{i=0}^{j-2} B_{nji} t_{j-1}^i e^{-(\lambda_D+r)t_{j-1}} \Big] e^{(\lambda_D+r)t_{j-1}} \; . \end{aligned}$$

With $r = 0$ we get that for $t \leq t_n$,

$C_{nj} = C_{n-1,j} + \gamma/\lambda$,

$A_{n-1,j,n-j} = 0$ for $i = n-j$,

$A_{n,j,i} = \lambda/i A_{n-1,j,i-1}$ for $i \in \{1,\ldots,n-j\}$,

$B_{n,j,j-2} = \lambda/\lambda_P B_{n-1,j,j-2}$ for $i = j-2$,

$B_{n,j,i} = \lambda/\lambda_P B_{n-1,j,i} + \sum_{k=i+1}^{j-2} (-1)^{k-i} k\lambda/(\lambda_P^{(k+1)}) B_{n-1,j,k}$ for $i \in \{0,\ldots,j-3\}$, and

$A_{nj0} = \left[V_{n-1,j}(t_{N-1}) - (C_{nj} + \sum_{i=0}^{j-2} B_{nji} t_{j-1}^i e^{-\lambda_D t_{j-1}} + \sum_{i=1}^{n-j} A_{nji} t_{j-1}^i e^{-\lambda t_{j-1}})\right] e^{\lambda t_{j-1}}$.

And for $t > t_n$:

$C_{nj} = \pi\lambda_D/\lambda_D + C_{n-1,j} = nb/\lambda_D$, $A_{nji} = 0$,

$B_{n,j,i} = \lambda_D/i B_{n-1,j,i-1}$ for $i \in \{1, \ldots, j-2\}$, and

$B_{nj0} = \left[V_{n-1,j}(t_{N-1}) - (C_{nj} + \sum_{i=1}^{j-2} B_{nji} t_{j-1}^i e^{-\lambda_D t_{j-1}})\right] e^{\lambda_D t_{j-1}}$.

This completes the induction. ■

Proof of Corollary 1

As $\Delta V_n(t) \geq \Delta V_{n+1}(t)$ (cf.Lemma 4), the time point at which $(1-\theta)\pi \geq \Delta V_n(t)$ becomes an equality for n, t_n, occurs before the same point for $n+1$, t_{n+1}. Given that t_n is unique (cf. Theorem 1), $t_{n+1} > t_n, \forall n$.

Proof of Theorem 2

Let $r = 0$. Assume there is one single unit to be sold during the interval $[0, T]$. Recall that time is counted backwards, i.e., $t = T$ is the beginning of the season and $t = 0$ is the end of the season. We have shown that for the seller, it is optimal to sell it directly at the store between T and t_1, and to sell it both directly and through the platform between t_1 and 0. According to our previous results, if we denote by θ the percentage commission of the platform, t_1 is uniquely defined by

$$1 - \theta = (1 - \zeta\theta)\left(1 - e^{-\lambda t_1}\right)$$

where $\lambda = \lambda_D + \lambda_P$ and $\zeta = \lambda_P/\lambda$.

Letting $y = e^{-\lambda t_1}$, we have that

$$y = 1 - \frac{1-\theta}{1-\zeta\theta} = \frac{(1-\zeta)\theta}{1-\zeta\theta}$$

so

$$\theta = \frac{y}{1-\zeta(1-y)}.$$

Now, consider the profit Π made by the platform: during $[t_1, T]$, the seller sells it with infinitesimal intensity λ_D, so the probability that it has not been sold at t_1 is equal to $e^{-\lambda_D(T-t_1)} = e^{-\lambda_D T} y^{-1+\zeta}$; similarly during $[0, t_1]$, the probability that it is sold is equal to $1 - e^{-(\lambda_D+\lambda_P)t_1} = 1 - y$, and conditional on a sale, the probability that the sale comes from the platform is $\lambda_P/\lambda = \zeta$. Finally, the revenue earned by the platform is $\theta\pi$, where π is the selling price. Hence,

$$\Pi(y) = \pi \frac{y}{1-\zeta(1-y)} e^{-\lambda_D T} y^{-1+\zeta} \zeta(1-y) = \pi\zeta e^{-\lambda_D T} \frac{y^\zeta(1-y)}{1-\zeta(1-y)}$$

As a result

$$\frac{\Pi'}{\Pi} = \frac{\zeta}{y} - \frac{1}{1-y} - \frac{\zeta}{1-\zeta(1-y)}$$

Consider a critical point y_0 of Π, so that $\Pi'(y_0) = 0$. In other words,

$$\frac{\zeta}{1-\zeta(1-y_0)} = \frac{\zeta}{y_0} - \frac{1}{1-y_0}$$

Then

$$\begin{aligned}
\frac{d}{dy}\left(\frac{\Pi'}{\Pi}\right) &= \frac{\Pi''(y_0)}{\Pi(y_0)^2} \\
&= -\frac{\zeta}{y_0^2} - \frac{1}{(1-y_0)^2} + \frac{\zeta^2}{(1-\zeta(1-y_0))^2} \\
&= -\frac{\zeta}{y_0^2} - \frac{1}{(1-y_0)^2} + \left(\frac{\zeta}{y_0} - \frac{1}{1-y_0}\right)^2 \\
&= \frac{\zeta^2}{y_0^2} - \frac{\zeta}{y_0^2} - \frac{2\zeta}{(1-y_0)y_0} - \frac{1}{(y_0^2-2y_0+1)} + \frac{1}{(1-y_0)^2} \\
&= \frac{\zeta^2}{y_0^2} - \frac{\zeta}{y_0^2} - \frac{2\zeta}{(1-y_0)y_0} \\
&= -\frac{\zeta(1-\zeta)}{y_0^2} - \frac{2\zeta}{y_0(1-y_0)} \\
&\leq 0.
\end{aligned}$$

Hence Π is quasi-concave, because it can have no local minimizer. ■

Proof of Theorem 3

Let $r = 0$. Under the proposed policy, we can show that the supplier is always indifferent between $z_t = 0$ and $z_t = 1$. When this is the case, for all t, $V_n(t) = V_n^0(t)$, because the HJB equation can be written as

$$\begin{aligned}
\frac{dV_n}{dt} &= \sup_{\mathbf{z}}\left[\Big(\pi\lambda_D + \mathbf{z}\pi(1-\theta_t)\lambda_P\Big) - \Big(\lambda_D + \mathbf{z}\lambda_P\Big)\Delta V_n(t) - rV_n(t)\right] \\
&\quad \text{using } \Delta V_n = (1-\theta)\pi \text{ , as } V'_{n,z=0} = V'_{n,z=1} \\
&= \sup_{\mathbf{z}}\left[\Big(\pi\lambda_D + \mathbf{z}\Delta V_n^0(t)\lambda_P\Big) - \Big(\lambda_D + \mathbf{z}\lambda_P\Big)\Delta V_n^0(t) - rV_n^0(t)\right] \\
&= \pi\lambda_D + \mathbf{z}\lambda_P\Delta V_n^0(t) - \lambda_D\Delta V_n^0(t) - \mathbf{z}\lambda_P\Delta V_n^0(t) - rV_n^0(t) \\
&= \pi\lambda_D - \lambda_D\Delta V_n^0(t) - rV_n^0(t) \\
&= \frac{dV_n^0}{dt}
\end{aligned}$$

and hence $V_n(0) = V_n^0(0) = 0$ automatically implies the property for all $t \geq 0$. This proves that the proposed pricing policy gives the supplier an expected profit of $V_N^0(T)$ and results in $z_t^* = 1$ being optimal (formally, it is indifferent, but the platform could give it an infinitesimal margin ϵ to break

the tie, leading to the same expected revenue).

With the current policy, the total revenue captured by the system (i.e., jointly by the supplier and the platform), denoted $\Pi_n(t)$, satisfies $\Pi_n(t) = V_n^1(t)$, because $\Pi_n(0) = V_n^1(0) = 0$ and

$$\begin{aligned}
\frac{d\Pi_n}{dt} &= \sup_{\mathbf{z}} \left[\pi\Big(\lambda_D + \lambda_P\Big) - \Big(\lambda_D + \lambda_P\Big)\Delta W_n(t) - rW_n(t) \right] \\
&= \sup_{\mathbf{z}} \left[\Big(\pi\lambda_D + \mathbf{z}\pi\lambda_P\Big) - \Big(\lambda_D + \mathbf{z}\lambda_P\Big)\Delta V_n^1(t) - rV_n^1(t) \right] \\
&= \frac{dV_n^1}{dt}
\end{aligned}$$

Hence, with this policy, the value captured by the system is maximal, while the revenue captured by the supplier is kept at minimum: it thus maximizes the platform's expected revenue captured, which is equal to $V_N^1(T) - V_N^0(T)$. ■