

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Jost, Peter-J.; Ressi, Anna

Article — Published Version What can I do for you? Optimal market segmentation in service markets

Production and Operations Management

Provided in Cooperation with: John Wiley & Sons

Suggested Citation: Jost, Peter-J.; Ressi, Anna (2022) : What can I do for you? Optimal market segmentation in service markets, Production and Operations Management, ISSN 1937-5956, Wiley, Hoboken, NJ, Vol. 31, Iss. 7, pp. 2838-2852, https://doi.org/10.1111/poms.13720

This Version is available at: https://hdl.handle.net/10419/265036

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



NC ND http://creativecommons.org/licenses/by-nc-nd/4.0/





DOI: 10.1111/poms.13720

ORIGINAL ARTICLE



What can I do for you? Optimal market segmentation in service markets

Peter-J. Jost¹ D | Anna Ressi² D

¹Chair of Organizational Theory, WHU – Otto Beisheim School of Management, Burgplatz, Vallendar, Germany

²WHU – Otto Beisheim School of Management, Burgplatz, Vallendar, Germany

Correspondence

Anna Ressi, WHU – Otto Beisheim School of Management, Burgplatz 2, D-56179 Vallendar, Germany. Email: anna.ressi@whu.edu

Handling Editor: Haresh Gurnani

Abstract

This paper considers price competition in a market where two firms sell a homogeneous service to a continuum of customers differing with respect to some exogenous characteristic. Our paper's novelty consists of explicitly acknowledging a distinctive property of many services in that firms incur customer-specific service costs after the contract is signed. Hence, not only the customers' willingness-to-pay and as such demand but also the firms' supply are related to customer characteristics. In this paper, we shed light on the implications thereof for optimal pricing and market segmentation strategies in a monopoly as well as a duopoly market. Importantly, we stress the profitability of services by demonstrating that firms in highly competitive industries still earn positive expected profits in equilibrium.

KEYWORDS

Bertrand paradox, market segmentation, personalized pricing, services

1 | INTRODUCTION

Recent advances in information technologies and the growing use of Internet transactions enable firms and other institutions to gather, store, and process unprecedented volumes of consumer data. So far, a particular focus has been on the online sale of consumer goods, where online searches and shopping behavior can be traced back to so-called cookies, web beacons, or Etags and stored in remarkable detail. These large data sets can then be utilized to more closely tailor the firm's marketing mix, such as online ads, product features, and pricing strategies, to customer characteristics (Choe et al., 2018; Obermiller et al., 2012; Valletti & Wu, 2020). Whereas the new reality of big data has refueled the interest of marketing and management scholars in price discrimination and inspired a vast literature thereof, the implications of the recent technological developments for service industries are not well understood. This is surprising as there is ample empirical and anecdotal evidence documenting that the use of the increased availability of personal data is particularly prominent in service industries such as business consultancy or catering (Arora et al., 2008; Choudhary et al., 2005).

A distinctive feature of various services is that firms incur customer-specific costs after the contract is signed. Hence, not only the customers' willingness-to-pay and as such demand but also the firms' supply are related to customer characteristics. An illustrative example is the market for support services, for example, in the realm of enterprise software applications. In recent years, it has become increasingly more common for firms to offer customer service through the use of implementation support, one-to-one training, repair schedules, and uptime guarantees (Ghose & Huang, 2009). Depending on the customers' characteristics (e.g., the scope of the software's application), the customer's willingness-to-pay, as well as the frequency and complexity of support and training services, and therefore also the firm's service costs vary. These features are also present, for example, in the banking context, where a customer's willingness-to-pay for their account correlates with their utilization level for the bank's service offerings.

Our work centers on trying to understand the implications of customer-specific service costs for monopoly as well as

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2022 The Authors. Production and Operations Management published by Wiley Periodicals LLC on behalf of Production and Operations Management Society.

Accepted by Haresh Gurnani, after three revisions.

competitive firms that offer a single homogeneous service at different contractual terms. We address the following main research questions:

- 1. What are the optimal pricing and the resulting market segmentation strategies of service providers in a monopoly and duopoly market?
- 2. Can service providers earn positive profits despite selling a homogeneous service in a competitive market?
- 3. Can service providers benefit from being able to employ personalized pricing?

Motivated by these questions, we develop an analytical framework in which we explicitly acknowledge that not only demand but also a firm's supply are related to customer characteristics. We consider both a monopoly market and Bertrand competition between two symmetric firms. Firms sell a homogeneous service to a continuum of customers by offering a menu of contracts. Customers are marked by an exogenous idiosyncratic characteristic θ , uniformly distributed in [0, 1], which constitutes a reasonably accurate and observable proxy for the utilization level. Firms, first, simultaneously announce their portfolio of service contracts, and customers choose the offer that leads to the greatest utility. For each customer, firms then decide whether to accept or reject the demand. A crucial feature of our model is that not only the customers' willingness-to-pay but also the service costs increase in θ . As the characteristics of customers are observable, firms will have an incentive to not serve those with a high level of θ that, given a specific price, cause unprofitably high service costs. This means that firms are to some extent able to restrict customers' self-selection opportunities, and the market will be segmented by interconsumer heterogeneity. In this market environment, we are interested in the firms' equilibrium pricing and the resulting market segmentation strategies.

In our analysis, we differentiate between two cases that pertain to the contractability of our observable proxy for the utilization level, θ . We will see that this assumption crucially determines the degree to which firms are able to restrict customers' self-selection. In the first case, θ is contractible such that firms can offer contracts that stipulate a price for the service as well as a maximum permissible θ for this specific price. In the example of the IT support market mentioned before, there is a specific number of people within a firm who must have access to a specific software. The number of licenses held by the customer then correlates with the service provider's expected scope of their customer support activities and can serve as a reasonable proxy that is contractible. The service provider can, hence, publicly advertise a menu of contracts that stipulate a price as well as a maximum number of licenses. In the second case, θ is observable but not contractible. Contractually specifying a characteristic might, for example, be prohibitively costly. For example, IT implementation support crucially depends on the complexity of the customers' IT infrastructure. Whereas

complexity might be observable to the service provider, it is hardly objectively measurable and can, hence, be regarded as noncontractible. The service provider might, nevertheless, use it when advertising its menu of contracts by, for example, offering a *Basic*, *Intermediate*, and *Premium* contract for different prices, and informally linking the eligibility to a particular contract to the complexity of the customer's IT infrastructure. The customer then demands one of these contracts and the service provider will agree if this business relationship promises to be profitable.

We derive several interesting results. Among others, basic intuition might predict that price competition between two symmetric firms offering a service that is homogeneous for each given customer sparks a continuous undercutting mechanism that ultimately leads to zero economic profits. Notably, whereas this intuition might hold for the pricing of typical products, we show that this is indeed not the case in service markets. The main driver in our paper is that, with every price reduction, the market segment that is being served shifts to the left as the firm will be able to sell to customers marked by a lower level of θ but will abandon some customers marked by unprofitably high levels of θ . We show that the undercutting process will reach a point where profits from further undercutting the competitor are lower than increasing the price to again target customers at the upper part of the spectrum. However, then firms find themselves in the same situation where they started, and the undercutting process starts anew. A crucial difference to previous studies on market segmentation, hence, is that there does not exist an equilibrium in pure strategies but only in mixed strategies, which results in positive expected profits.

Whereas in most of the paper we assume that the number of contracts firms can offer is finite, relaxing this assumption presents an even more intriguing case. Note that in this scenario the firms' portfolios of service contracts might encompass an individual price for each potential customer. Hence, they can essentially engage in personalized pricing and compete in infinitely many Bertrand markets for a single customer. Interestingly, our analysis shows that competitive firms do not use their ability to offer personalized pricing, which is why expected equilibrium profits are still positive, and the Bertrand paradox does not hold. Moreover, we show that firms are better off in terms of expected profits when having the possibility to employ personalized prices than when they are limited to charge only a finite number of different prices.

The paper is organized as follows. The next section discusses related literature and highlights our contribution. In Section 3, we present our basic model. Section 4 contains the derivation of our results for the monopoly market. Section 5 extends our analysis to a competitive market. The scenario where firms are able to offer infinitely many prices and can, in principle, employ personalized pricing is presented in Section 6. Finally, in Section 7, we conclude with some final remarks and provide avenues for further research. Proofs are relegated to the Supporting Information.

Production and Operations Management -

2 | RELATED LITERATURE

Our paper contributes to three areas of research: the literature on monopoly and competitive market segmentation, personalized pricing, and the Bertrand paradox. Generally, a firm's potential customer base can be segmented by employing price discrimination, where a given product is sold to customers for different prices.¹ According to Stole (2007), firms can either engage in direct price discrimination where there exists an observable component of customer characteristics that prices may be conditioned upon. They then exploit intercustomer heterogeneities in customers' individual demand schedules. Using the terminology of Pigou (1920), direct price discrimination encompassed third-degree price discrimination as well as first-degree or perfect price discrimination, respectively, as an extreme case. When a component of customer characteristics is not observable, firms can also engage in indirect price discrimination (or, following Pigou, 1920, second-degree price discrimination). They then usually exploit intra-customer heterogeneities that refer to a given customer's evaluation of successive units by structuring appropriate contracts prompting customers to choose the one that is most profitable for the firm given the self-selection constraint. The market can also be segmented by offering different versions of a product that are then sold for different prices, that is, a product line. Thereby, the firm again utilizes intracustomer differences in customers' valuation, for example for quality, by using appropriate price-quality schedules as a sorting mechanism (Stole 2007).²

The fact that customers obtain different service levels when purchasing a contract might be reminiscent of the literature on competitive product line design or versioning (see, e.g., Champsaur & Rochet, 1989; De Fraja, 1996; Johnson & Myatt, 2006; Stole, 1995). However, in our paper, these different service levels should not be understood as service qualities. Instead, firms sell a single homogeneous service at different contractual terms to a continuum of customers differing in the expected level of ex post utilization. Accordingly, in our paper, offering multiple service contracts is theoretically not equivalent to offering a product line with different versions. Producing a specific quality of a product entails the same production costs irrespective of the customer that buys it and their valuation of the quality. In contrast, if a firm in our model offers its service at different contractual terms, not every customer buying a specific contract utilizes the service to the same extent and, hence, does not engender the same costs for the firm.³ Hence, the core competitive dynamics of our service model crucially differ from other models of product line design with differentiated products and warrant a thorough theoretical analysis.

According to the definitions described above, our model with contractible characteristics rather relates to third-degree price discrimination as firms can perfectly control customers' self-selection by specifying a maximum level of the characteristic θ and group them in different market segments. The results of our paper generally confirm the findings of Holmes

(1989) and Armstrong and Vickers (2001), indicating that third-degree price discrimination typically increases profits under best response symmetry and sufficient competition.⁴ However, whereas Holmes (1989) analyzes a model with two market segments where consumers are exogenously grouped in a weak and strong market, Armstrong and Vickers (2001) consider a model in which consumers in each market segment differ by their transportation costs, that is, each segment is of a Hotelling type. Thus, the crucial difference of both models to our paper is that market segmentation is exogenously instead of endogenously given and that there are only finitely many segments.

Our setting with noncontractible but observable characteristics essentially employs a mixture of the two approaches of direct and indirect price discrimination mentioned above.⁵ By being able to reject unprofitable customers, a firm can utilize the observability of intercustomer heterogeneities to create market segments, which is again reminiscent of thirddegree price discrimination. However, the crucial difference is that, because characteristics are not contractible, firms have to some extent resign to customer self-selection in case two contracts create an internal rivalry. At the same time, because our model does not involve intracustomer heterogeneities, it does not directly compare to other studies that involve pure self-selection. The reason for the lack of comparable studies investigating price discrimination in the context of noncontractible characteristics is that results are economically immediate. If a firm sells a single typical product, it lacks a feasible segmenting mechanism and will be forced to charge a single price, which, in the face of competition will equal marginal cost. We show that this is in stark contrast to a situation where firms face customer-specific service costs. These costs actually equip firms with a feasible segmenting mechanism as they now have an incentive to reject a customer's demand that will lead to unprofitably high service costs. Hence, the scantness of a differentiated analysis concerning services is an important omission in the literature that we seek to address.

By considering an infinite number of market segments, our paper also relates to the literature on personalized pricing, where firms can offer a customized price designed for one specific buyer; see Zhang (2009) for an overview. Most papers on competitive personalized pricing assume ex ante heterogeneous firms with horizontally and/or vertically differentiated products (Bester and Petrakis, 1996; Choe et al., 2018; Choudhary et al., 2005; Ghose & Huang, 2009; Shaffer & Zhang, 2002). Given that firms are ex ante symmetric, existing papers argue that the opportunity to employ personalized pricing reduces profits. Choudhary et al. (2005, p. 1122) even state the following:

It is immediate to show in this model that, if the qualities of the firms are the same, PP adds no value – the result is Bertrand competition, with both firms pricing at marginal cost.

In our context of services with infinitely many market segments, the conclusion of Choudhary et al. (2005) would

apply to each segment and would, hence, predict zero overall profits for both firms. We challenge this intuition by showing that symmetric firms who are in principle able to offer personalized prices still earn positive expected profits in equilibrium.

So far, there only exists a few papers showing that firms' ability to employ personalized pricing can be beneficial for firms. In Ghose and Huang (2009), two multiproduct firms can also earn positive profits and be better off under personalized pricing and personalized quality. Different from us, they assume the existence of intraconsumer heterogeneities, which represent both quality and brand preferences. These brand preferences allow each firm to appropriate a subset of the customer space by setting the price to match the consumer surplus of the rival and thereby earn positive profits. Related to this, Chen et al. (2020) show that, in a Hotelling model, firms can be better off with personalized pricing if customers engage in identity management. In Shaffer and Zhang (2002), vertically and horizontally differentiated firms can also profit from personalized pricing if they are ex ante asymmetric regarding their market size. The crucial difference between this paper's result and ours is that, although competitive firms are in principle able to offer a different price to each customer, they will instead create market segments where multiple customers are charged the same price. What is more, our results suggest that expected profits when firms are able to employ personalized pricing are higher than when they are limited to charge only a finite number of different prices. We hence complement this stream of literature by demonstrating that, even in highly competitive industries and symmetric firms, service providers can earn positive profits in equilibrium and be better off when being able to employ personalized pricing than when they are not.

Finally, when the number of market segments tends to infinity, firms essentially compete in infinitely many Bertrand markets. As such, our paper also speaks to the theoretical literature on the Bertrand paradox (Bertrand, 1883). Research in previous years has shown that the zero-profit equilibrium outcome crucially depends on properties of the demand and supply specification as well as on other modeling features. Harrington (1989), for example, shows that zero profits is the only equilibrium outcome when firms produce at constant marginal cost and market demand is bounded, continuous, downward sloping, and has a finite choke price. In the absence of these assumptions, the Bertrand paradox does not hold (see, e.g., Baye and Morgan, 1999, 2002; Dastidar, 1995; Kaplan & Wettstein, 2000; Kreps & Scheinkman, 1983; or Hoernig, 2002, 2007; Routledge, 2010, or Szech & Weinschenk, 2013). In most of these models, the equilibrium price-setting behavior by firms is characterized by mixed strategies as in the present paper.⁶ We add to this literature on the Bertrand paradox by showing that firms earn positive expected profits in Bertrand competition when supply and demand depend on customer characteristics.

3 | THE BASIC MODEL

We consider two different market structures: first, a monopoly market and, second, a duopoly market where firms compete in prices to sell a single homogeneous service. In this context, we are interested in the firms' optimal market segmentation strategies. In the following, we describe the market in more detail.

On the demand side, there exists a continuum of customers, each demanding at most one unit of the service. They differ with respect to an exogenous and observable characteristic θ , which constitutes a reasonably accurate proxy for their expected ex post utilization level and is positively related to their willingness-to-pay. We assume θ to be uniformly distributed on the line segment [0, 1]. Formally, the utility of a particular customer with characteristic θ of buying the homogeneous service for a price of *p* is given by

$$u = \theta - p. \tag{1}$$

On the supply side, we either consider a monopoly firm M or a duopoly with firm 1 and firm 2. The cost of supplying the service to a particular customer after having signed the contract, that is, the customer-specific service cost, is positively related to the customer's θ . Formally, the profit of firm $i, i \in \{M, 1, 2\}$, when selling the service for a price p_i to a customer with characteristic θ is given by

$$\Pi_i = p_i - c\theta,\tag{2}$$

where $c \in (0, 1)$.

Assume that firm $i \in \{M, 1, 2\}$ offers a range of $n \in \mathbb{N} \setminus \{1\}$ different contracts for the single homogeneous service and hence segments the market into *n* different customer groups.⁷ In the course of the paper, we differentiate between two cases where, first, θ is contractible (and as such of course also observable), and firms can employ third-degree price discrimination. We then denote a firm's contract schedule by $\pi_i = ((p_{i1}, \theta_{i1}^{max}), (p_{i2}, \theta_{i2}^{max}), \dots, (p_{in}, \theta_{in}^{max}))$, where w.l.o.g. prices are ordered in a descending manner, $p_{i1} \ge p_{i2} \ge \dots \ge p_{in}$, and θ_{ij}^{max} , $j = 1, \dots, n$, denotes the maximum level of θ that a customer is allowed to hold in order to be eligible for buying the contract for a price of p_{ij} . In the second case, θ is still observable but not contractible. A firm's contract schedule then only formally specifies *n* different prices for the service, that is, $\pi_i = (p_{i1}, p_{i2}, \dots, p_{in})$.

We model firms' decisions and the interaction with customers as a three-stage game: In the first stage, firms simultaneously and publicly set their service contracts π_i . In the second stage, each customer articulates a demand for the contract that gives them the highest utility and, in the case of contractible θ , that they are eligible to. In the third and final stage, firms decide whether or not to accept, and, finally, profits and utilities realize. We use the subgame-perfect equilibrium concept and solve the game by backward induction.⁸ Let us begin by analyzing the third stage of the game. Firm *i* will sell its service for a price of p_{ij} only when this trade is profitable. This will only be the case if the customer's characteristic is sufficiently low, that is, if $\theta \leq p_{ij}/c$. Let us denote this threshold by $\overline{\theta}_{ij}$. If θ is not contractible but observable, the firm will only accept a demand if $\theta \leq \overline{\theta}_{ij}$. In case θ is contractible, the firm will always set $\theta_{ij}^{max} \leq \overline{\theta}_{ij}$ and accepts a customer demanding price p_{ij} whenever $\theta \leq \theta_{ij}^{max}$. On the demand side, which becomes relevant at the second stage of the game, a customer marked by θ will demand the service only if his willingness-to-pay for the service is sufficiently high, that is, if $\theta \geq p_{ij}$. We denote this threshold by $\underline{\theta}_{ij}$.

We henceforth call $S_{ij} = [\underline{\theta}_{ij}, \overline{\theta}_{ij}]$ a firm's market segment that is generated with price p_{ij} .⁹ Given $\overline{\theta}_{ij} = \theta_{ij}^{max}$ and unless there is a better alternative, all customers with a characteristic of $\theta \in S_{ij}$ demand the price p_{ij} , firms accept, and the market will be segmented by intercustomer heterogeneity θ . Finally, $D_{ij}(\pi_i; \pi_{-i}) \subseteq [0, 1], -i \neq i, -i, i \in \{1, 2\}$, defines the actual demand that comprises all customers who actually buy the service of firm *i* at price p_{ij} . The remainder of the paper will now be concerned with solving the first stage of the game, which eventually leads to a characterization of the *n* service contracts and resulting market segments.

4 | THE MONOPOLY MARKET

To highlight the peculiarities of our market environment and develop a clear intuition for our results, we first present the most simple case of a monopoly firm. Because the scenario with contractible characteristics where contracts stipulate a price as well as a maximum permissible level θ_{ij}^{max} is more common in reality, it constitutes a natural starting point that we discuss in the following section. We then analyze the market with noncontractible characteristics in Section 4.2. The monopoly case will then serve as a benchmark and springboard that renders it easier to grasp the intuition of our results in the duopoly market.

4.1 | Contractible customer characteristics

If θ is contractible, a given service contract stipulates a price p_{Mj} as well as a corresponding maximum permissible level θ_{Mj}^{max} . For example, when implementing a new software system, larger firms require more comprehensive customer support activities. Software providers may then offer different contracts for their after-sale customer service based on the number of employees that need the software and hold a license.¹⁰

To begin with, let us develop an intuition for how the monopolist can, in isolation, maximize the profit stemming from the most expensive contract sold for a price of p_{M1} . Because all customers with $\theta > \overline{\theta}_{M1} = p_{M1}/c$ are not profitable for the monopolist, the optimal maximum customer characteristic for this contract will be directly determined by $\overline{\theta}_{M1}$ and as such by the price. We can, therefore, treat p_{M1} as the monopolist's single choice variable regarding its first contract. The respective profit can then be formulated as follows:

$$\Pi_M(p_{M1}) = \int_{\underline{\theta}_{M1}}^{\underline{\theta}_{M1}^{max}} (p_{M1} - c\theta) \, d\theta, \qquad (3)$$

with $\underline{\theta}_{M1} = p_{M1}$ and $\underline{\theta}_{M1}^{max} = \min{\{\overline{\theta}_{M1}, 1\}}$. Note that as long as $0 \le \overline{\theta}_{M1} < 1$ the profit is increasing in p_{M1} because both the price and the size of the market segment increase. Hence, customers on the upper side of the market are always the most profitable to target, and it will never be optimal to set p_{M1} so that $\overline{\theta}_{M1} < 1$. When increasing the price so that $\overline{\theta}_{M1}$ is above 1 and, hence, $\underline{\theta}_{M1}^{max} = 1$, the monopolist faces a trade-off between decreasing the number of customers it would serve and charging the customers that are still served a higher price.

Next, we take into account that the monopolist seeks to maximize the combined profits from all n > 1 contracts sold at respective prices p_{Mj} , j = 1, ..., n.¹¹ As profits are increasing in p_{Mj} as long as $0 \le \overline{\theta}_{Mj} < 1$, the monopolist will never set its prices so that it leaves customers between the first and the second targeted market segment unserved. From this it follows that $\overline{\theta}_{Mj+1} \ge \underline{\theta}_{Mj}$. In case $\overline{\theta}_{Mj+1} > \underline{\theta}_{Mj}$, it will of course always be optimal to charge customers in the intersection the higher price p_{Mj} . Hence, the optimal θ^{max} of a given contract j > 1 will be determined by the preceding contract, that is, $\theta_{Mj}^{max} = \underline{\theta}_{Mj-1} = p_{Mj-1}$, for j > 1. This implies that we can focus on the monopolist's optimal prices, and the associated levels of θ_{Mj}^{max} included in the contracts will immediately follow. We can therefore treat p_{Mj} as the monopolist's choice variables.

Maximizing the monopolist's profit function

$$\Pi_{M}(p_{Mj}) = \int_{p_{M1}}^{1} (p_{M1} - c\theta) d\theta$$
$$+ \sum_{j=2}^{n} \int_{p_{Mj}}^{p_{Mj-1}} (p_{Mj} - c\theta) d\theta \qquad (4)$$

with respect to prices p_{Mj} , j = 1, ..., n leads to the monopolist's optimal pricing schedules (see the Supporting Information). The following proposition summarizes our results.

Proposition 1. In case of a monopolist market for a service with contractible customer characteristics, the firm's optimal market segments are determined by the pricing schedule

$$p_{Mj}^* = \frac{n - (j - 1) - (n - j)c}{1 + n(1 - c)} \text{ for } j = 1, \dots, n, \ n \in \mathbb{N} \setminus \{1\}.$$
(5)

Each contract selling for a price of p_{Mj}^* also stipulates a maximum level of $\theta_{Mj}^{max} = 1$ for j = 1 and $\theta_{Mj}^{max} = p_{Mj-1}^*$ for j > 1.

The monopolist's profits are then given by

$$\Pi_M^* = \frac{(1-c)^2 n}{2+2(1-c)n}.$$
(6)

All in all, our results show that, akin to standard models of third-degree price discrimination, the optimal pricing decision leads the monopolist firm to first target customers with the highest characteristic θ on the upper side of the market. For every additional market segment, it will then successively reduce the price to target additional customer groups that comprise customers with lower levels of θ . Also, each contract will be sold to the same number of customers, that is, $p_{Mi} - p_{Mi-1} = p_{Mi+1} - p_{Mi}$, for all j > 1.

4.2 | Noncontractible customer characteristics

We now derive the optimal contracts $\pi_M^* = (p_{M1}^*, p_{M2}^*, \dots, p_{Mn}^*)$ when customer characteristics are not contractible and, therefore, the monopolist is not able to formally specify a maximum level θ_{Mj}^{max} for each price *j*. As mentioned in the introduction in greater detail, firms then offer a portfolio of contracts that formally differ in price while justifying these price differences by only informally referring to θ .

We start our analysis with the simple case where the monopolist can only offer two contracts, that is, n = 2, at prices p_{M1} and p_{M2} . We argue that, unlike Section 4.1, it will not be optimal to have intersecting market segments and hence an internal rivalry between contracts. Instead, market segments will be adjacent to each other. We prove this claim by contradiction. Suppose that the two contracts optimally stipulate prices such that $\overline{\theta}_{M2} = p_{M2}^*/c < p_{M1}^* = \underline{\theta}_{M1}$ and, hence, leave customers in between the two contracts unserved. Similar to before, however, profits from a single contract are increasing in price. Therefore, the monopolist can boost its profits by increasing p_{M2} at least to $p_{M2}/c =$ p_{M1}^* , which implies that the initial prices could have never been the optimal ones. Consequently, the optimal prices have to fulfil the condition $p_{M2}^*/c \ge p_{M1}^*$. Next, suppose that the two contracts optimally involve an intersection, that is, $\overline{\theta}_{M2} =$ $p_{M2}^*/c > p_{M1}^* = \underline{\theta}_{M1}$. The crucial difference to Section 4.1 is that the monopolist cannot use a maximum permissible level of θ to force customers in the intersection to pay the higher price p_{M1}^* instead of p_{M2}^* . Hence, the firm has to succumb to customers' self-selection. The immediate consequence is that all customers at the intersection will, of course, always choose the lower price. But then, the monopolist can, for example, boost its total profits by increasing p_{M1}^* to $p'_{M1} =$ θ_{M2} , which would resolve the intersection of the two market segments. Thereby, it will leave profits from the contract stipulating p_{M2}^* unaffected. At the same time, because we know that $\overline{\theta}_1 \geq 1$, the market segment generated by p'_{M1} will stay the same, but all customers in this segment will be charged a higher price. Consequently, the initial prices could have never been optimal, which contradicts our assumption from the beginning. Overall, in the optimum, the monopolist will have to set its prices such that the two market segments are adjacent to each other. Hence, it holds that $\overline{\theta}_{M2} = \underline{\theta}_{M1}$, which implies that $p_{M2}^* = c p_{M1}^*$.¹² The same line of argumentation equally holds for any pair

The same line of argumentation equally holds for any pair of contracts within a larger set of n > 2 contracts. It then follows that for any given market segment targeted by price p_{Mj} the price p_{Mj+1} of the next market segment that lies to the left is directly determined by the former price p_{Mj} . Then, any range of service contracts determined by the monopolist's pricing schedule $\pi_M = (p_{M1}, p_{M2}, ..., p_{Mn})$ will immediately follow from the price of the first segment on the upper side of the market where customers with the highest characteristic θ buy. For any given price of the first market segment, p_{M1} , every other price of additional market segments is therefore given by $p_{Mj} = c^{j-1} p_{M1}$.

Finally, we are now able to characterize the monopolist's optimal market segmentation strategy given that it can offer $n \in \mathbb{N} \setminus \{1\}$ service contracts. Building on the observations above, we can treat the monopolist's price p_{M1} as its choice variable to maximize profits:

$$\Pi_{M}(p_{M1}) = \sum_{j=1}^{n} \int_{c^{j-1}p_{M1}}^{\min\left\{1, c^{j-2}p_{M1}\right\}} \left(c^{j-1}p_{M1} - c\theta\right) d\theta$$
$$= p_{M1}^{2} \frac{\left(1 - c\right)^{2} \left(1 - c^{2n}\right)}{2c\left(1 - c^{2}\right)} - \frac{1}{2c} (\max\left\{p_{M1}, c\right\} - c\right)^{2}.$$
(7)

The following Proposition 2 presents our next result.

Proposition 2. In case of a monopolist market for a service with observable but noncontractible customer characteristics, the firms' optimal market segments are determined by the pricing schedule

$$p_{Mj}^* = c^{j-1} p_{M1}^* \text{ for } j = 1, \dots, n, \ n \in \mathbb{N} \setminus \{1\},$$
(8)

where the price targeted for the most upper market segment equals $p_{M1}^* = \frac{1+c}{2+(1-c)c^{2n-1}}$. The monopolist's profits are then given by

$$\Pi_M^* = \frac{c(1-c)(1-c^{2n})}{4c+2(1-c)c^{2n}}.$$
(9)

In contrast to the previous Section 4.1, the monopolist's optimal pricing schedule will result in adjacent market segments that are shrinking in size for j > 1.¹³ This means that the number of customers that any additional market segment

comprises is higher than that of the next market segment on the left side of the market, where customers with lower characteristics lie. Formally, $\frac{\partial}{\partial j}(p_{Mj}/c - p_{Mj}) < 0$. Furthermore, contracts do not involve internal rivalries. Finally, a comparison of prices and profits shows that both are lower when customer characteristics are noncontractible than when they are. Of course, this is due to the fact that in the former case the monopolist is restricted by the customers' self-selection.

It is immediate to show that our market outcome crucially differs from that of a firm selling a regular consumer good instead of a service. In the former case, the monopolist does, of course, never have an incentive to reject a customer's demand for the lowest price. Therefore, it lacks a feasible segmentation mechanism and will offer its product for a single price p. In contrast, we show that a service provider benefits from the observability of customer characteristics. The reason is that the firm incurs customer-specific costs after the contract is signed, which creates the incentive to reject certain customers that are not profitable. This enables the firm to credibly restrict customers' self-selection and thereby charge different prices that segment the market into multiple customer groups.

5 | THE DUOPOLY MARKET

We now extend our model to a competitive market with two firms, $i \in \{1, 2\}$, and search for a symmetric equilibrium.¹⁴ We again begin with the case where customer characteristics are contractible. Finally, we will see that our results also translate to situations where customer characteristics are not contractible.

5.1 | Contractible customer characteristics

Assume that two competitive firms $i, i \in \{1, 2\}$, offer a range of *n* different contracts and hence segment the market into *n* different customer groups. Whether in competition firm *i* is able to realize the potential demand of market segment $S_{ij} = [\underline{\theta}_{ij}, \overline{\theta}_{ij}]$ for the service associated with price p_{ij} of course depends on whether the competitor -i offers those customers a price p_{-il} that is lower than p_{ij} , where $-i \neq i, -i, i \in \{1, 2\}$ and l = 1, ..., n denotes a particular contract of the competitor in its menu of contracts. Two ranges of contracts that are determined by the pricing schedules $\pi_i = ((p_{i1}, \theta_{i1}^{max}), (p_{i2}, \theta_{i2}^{max}), ..., (p_{in}, \theta_{in}^{max})), i \in \{1, 2\}$, are then equilibrium offers if π_i maximizes the overall profits of firm *i* given its competition's choice π_{-i} , that is,

$$\max_{p_{ij}} \quad \Pi_i(\pi_i; \pi_{-i}) = \sum_{D_{ij}(\pi_i; \pi_{-i})} \left(\int_{D_{ij}(\pi_i; \pi_{-i})} \left(p_{ij} - \theta c \right) d\theta \right).$$
(10)

Akin to our previous section, we can again focus on the firms' pricing decisions and the optimal θ_{ij}^{max} will follow

from that. To grasp the basic intuition of our subsequent line of argumentation, we proceed in two steps: We first argue that firm *i* will never choose a pure strategy, that is, a given range of contracts π_i with certainty. This implies that there does not exist an equilibrium in pure strategies but only in mixed strategies. To derive the firms' pricing schedules, leading to a characterization of the firms' portfolio of contracts and the resulting market segments, we then concentrate on a symmetric equilibrium in mixed strategies with support over all pure strategies that are best responses to the competitor's pure strategies.

5.1.1 | The basic intuition for firms' equilibrium strategies

To understand the basic intuition for the competitive forces, let us first restrict our attention to the competition for the most profitable segment of customers targeted with the firms' highest price p_{i1} while leaving aside the remaining n - 1 contracts. Suppose that firm 1 thinks that firm 2 will set the price for its first contract p_{21} such that $\overline{\theta}_{21} \ge 1$. Firm 1's best response would now be to marginally undercut this price by setting $p_{11} = p_{21} - \varepsilon$, with $\varepsilon > 0$ but infinitely small. Since all customers who first intended to buy the service for price p_{21} from firm 2 now switch to firm 1, sticking to price p_{21} leads to zero profits. Now it is firm 2 that has an incentive to marginally undercut firm 1's price to recapture its customers. Obviously, there will be a back and forth between firm 1 undercutting firm 2's price p_{21} and firm 2 undercutting firm 1's price p_{11} .

This reasoning is reminiscent of firms' undercutting behavior in Bertrand product competition, which would end when both firms price their product at marginal costs and make zero economic profits. The crucial difference between the standard Bertrand product market and our market with services is that the market segments $[\underline{\theta}_{ii}, \overline{\theta}_{ij}] = [p_{ij}, p_{ij}/c]$ shift to the left for every reduction in price p_{ii} . If the service becomes cheaper, an additional group of customers with a lower characteristic θ and thus a lower willingness-to-pay will now find it beneficial to buy. At the same time, firms will find it unprofitable to sell the service to those customers with a high θ as they entail overly high service costs after the contract is signed. Consequently, for every reduction in price p_{ii} , an even larger share of customers on the upper side of the market remains unserved. At some point in time, we will show that this undercutting process for the first contract will reach a specific price \tilde{p} such that overall profits with $p_{i1} = \tilde{p}$ and its other contracts are the same as when the firm would stop the undercutting that targets customers on the lower end of the market and instead shift the focus back to the upper side of the market. The first contract will then target the customers with the highest levels of θ with a price of $p'_{i1} = \tilde{p}/c.^{15}$

Consider the case where the undercutting process is at the point where firm 2 charges $p_{21} = \tilde{p}$ and firm 1 thinks about whether to further undercut firm 2 by setting $p_{11} = \tilde{p} - \varepsilon$, or

to target the upper market segment with a price of $p'_{11} = \tilde{p}/c$ instead. For deriving firm 1's best response, note that \tilde{p} is always lower and \tilde{p}/c is always higher than the monopolist's price of $p_{M1} = \frac{n - (n-1)c}{1 + n(1-c)}$ derived in Section 4.1.¹⁶ It is well known that a firm's profits are increasing in the price as long as the price is below the monopoly price and vice versa. Because profits are the same with \tilde{p} and \tilde{p}/c , profits with $p_{11} = \tilde{p}/c$ are strictly higher than profits with $p_{11} = \tilde{p} - \varepsilon$. Hence, firm 1's best response to firm 2 setting $p_{21} = \tilde{p}$ is $p_{11} = \tilde{p}/c$. Because firm 2 will now once again have an incentive to undercut the price of firm 1, we are back to the situation we already had and the undercutting process begins anew. Hence, the best response to any price $p_{-i1} \in [\tilde{p}, \tilde{p}/c]$ of the competitor is always between \tilde{p} and \tilde{p}/c .¹⁷ However, as none of these prices between \tilde{p} and \tilde{p}/c are mutually best responses, there only exists an equilibrium in mixed strategies. In the next section, we derive a mixed strategy equilibrium over the support of prices that are a best response to the prices of the competitor.

5.1.2 | Firms' equilibrium market segmentation strategies

To see how firms will set their remaining prices, let us first start with the case where n = 2 before generalizing our line of argumentation to $n \in \mathbb{N} \setminus \{1\}$. We first consider, for illustrative purposes, the situation where firm 1 expects firm 2 to target the upper side of the market and set its prices such that market segments do not overlap and, hence, $p_{22} = c p_{21}$. What is firm 1's best response? We already know from the previous section that, with its first contract, firm 1 will marginally undercut p_{21} to steal the most profitable customers on the upper side of the market. How will it set its second price? As profits increase in price, firm 1 will not leave a gap between its two market segments. Also, creating intersecting market segments, that is, $p_{12}/c > p_{11}$, and charging customers at the intersection the higher price can never be a best response of firm 1. This is because firm 1 will then never make a profit with its second contract as $p_{22} < p_{12}$. Hence, firm 1's best response is to steal firm 2's demand by setting its prices p_{11} and p_{12} to marginally undercut both of firm 2's prices, thereby again creating adjacent market segments itself. Similar to the previous section, this will again result in a downward spiral until firms reach the specific price \tilde{p} such that firm *i*'s overall profits with $p_{i1} = \tilde{p}$ and $p_{i2} = c\tilde{p}$ are the same as when the firm would stop the undercutting and set $p_{i1} = \tilde{p}/c$ and $p_{i2} = \tilde{p}$, and the undercutting begins anew. All in all, a firm's best response to its rival setting its market segments adjacent to one another is also to create adjacent market segments, but none of these strategies is a mutual best response. The same argumentation analogously holds for n > 2.

Now suppose that firm 1 expects firm 2 to target the most profitable customers with the highest levels of θ but create an overlap that is greater than the optimal intersection of the



FIGURE 1 (a) Illustration of firm 2's market demands in the absence of firm 1 for prices $p_{21} = p_{M1}$ and $p_{22} = p_{M2}$, and both firms' market demands when firm 1 sets p_{11} and p_{12} as a best response. (b) Illustration of firm 1's best response when the undercutting spiral has reached a point where it becomes attractive to again target customers with the highest level of θ

monopolist derived in Section 4.1. The best response of firm 1 is to marginally undercut firm 2's first price and then set its second market segment such as to create the monopolist's profit-maximizing overlap. Hence, strategies with a sufficiently large intersection can never be a best response to any of its rival's strategies.

We can now go on to consider the final case where firm 1 expects firm 2 to create contracts with an overlap that is equal or lower than the optimal intersection of the monopolist in Section 4.1. In the latter case, increasing the intersection to the monopolist's optimal overlap cannot be profit-maximizing as then $p_{12} > p_{22}$, and firm 1 will not make any profits with its second contract. Hence, it is optimal for firm 1 to marginally undercut firm 2's prices and thereby also have overlapping market segments itself (see Figure 1a). This will lead to an undercutting spiral with overlapping market segments until we are at \tilde{p} . Suppose that firm 1 will now be the first to find it profitable to target the customers on the upper side of the market with a price of $p_{11} = \tilde{p}/c$. How will firm 1 now set its second price p_{12} ? As before, creating an intersection between one's own first and second contract cannot be optimal as firm 1's second price will not generate

- Production and Operations Management

any profits. Hence, firm 1 will set its second price p_{12} marginally lower than firm 2's first price p_{21} (see Figure 1b), and we are back at the downward spiral with adjacent market segments we discussed in the previous paragraph.

The key insight here is that, similar to the previous Section 5.1.1, prices that create overlapping market segments will never be chosen in equilibrium by iteratively eliminating strategies that cannot be a best response to given prices of the rival.¹⁸ Given our set of strategies for the first price, $p_{i1} \in [\tilde{p}, \tilde{p}/c]$, we know from the previous paragraph that a contract that specifies $p_{i1} = \tilde{p}/c$ can only be a best response if it does not involve an intersection. Hence, all contracts with $p_{i1} = \tilde{p}/c$ that involve overlapping market segments will be eliminated. We then marginally move to the left, $p_{i1} = \tilde{p}/c$ – ε , and see that all pairs of contracts with an intersection can only be a best response to a contract with $p_{i1} = \tilde{p}/c$ involving an intersection itself. Yet, these contracts are exactly the ones that we eliminated in the previous step. Proceeding in this way eventually results in the final set of pricing strategies that do not involve overlapping market segments, but none of these strategies are mutual best responses.

The same line of argumentation transfers to situations where n > 2 but, depending on the initial situation, it might require a significantly larger number of undercutting cycles until both firms' contracts create adjacent market segments. Suppose, for example, that n = 3 and firm 2 starts with intersections between its first and second as well as between its second and third contract that are lower than the monopolist's optimal intersections. Analogously to before, this will lead to an undercutting spiral where each firm sets its prices marginally lower than its rival's prices. Eventually, one firm, say firm 1, finds it optimal to again target the customers on the upper side of the market. We already know that it will do so by setting $p_{11} = \tilde{p}/c$ and that it will optimally set its second price to marginally undercut its rival's first price. Thereby, firm 1 resolves the intersection between its first and second contract while stealing firm 2's demand for its first contract sold at p_{21} . Firm 1 can now use its third contract to also steal firm 2's second market segment targeted with p_{22} . However, akin to before, firm 1 can only achieve that by setting the same intersection as firm 2 does. Overall, firm 1's first and second market segments are now adjacent, whereas the second and third market segments intersect. Firm 2 will now react by marginally undercutting each of firm 1's prices, prompting a new undercutting spiral with two adjacent and two intersecting market segments. At the point where one firm finds it optimal to target the upper side of the market, it will now face a situation where the rival firm's first and second market segments are adjacent. Marginally undercutting the respective prices finally also resolves the last intersection, and we again land in our familiar undercutting spiral without internal rivalries.

We now focus on this reduced set of contracts that do not involve overlapping market segments. As in Section 4.2, any range of service contracts determined by the pricing schedule $\pi_i = ((p_{i1}, \theta_{i1}^{max}), (p_{i2}, \theta_{i2}^{max}), \dots, (p_{in}, \theta_{in}^{max}))$ will then immediately follow from the price of the first segment. We there-

for again treat firm *i*'s price p_{i1} as its strategic variable from now on.

We can now calculate the critical price \tilde{p} such that firm *i*'s total profits when charging \tilde{p}/c as the price for the most upper segment and its total profits when charging \tilde{p} are the same:

$$\int_{\frac{\tilde{p}}{c}}^{1} \left(\frac{\tilde{p}}{c} - c\theta\right) d\theta + \sum_{j=2}^{n} \int_{c^{j-2}\tilde{p}}^{c^{j-3}\tilde{p}} \left(c^{j-2}\tilde{p} - c\theta\right) d\theta$$
$$= \sum_{j=1}^{n} \int_{c^{j-1}\tilde{p}}^{c^{j-2}\tilde{p}} \left(c^{j-1}\tilde{p} - c\theta\right) d\theta.$$
(11)

Note that the sum of profits from the market segments j = 2, ..., n created with a price of \tilde{p}/c are identical to the sum of profits from the market segments 1, ..., n - 1 created with a price of \tilde{p} . Hence, these terms cancel out, and the above condition can be simplified to

$$\int_{\frac{\tilde{p}}{c}}^{1} \left(\frac{\tilde{p}}{c} - c\theta\right) d\theta = \int_{c^{n-1}\tilde{p}}^{c^{n-2}\tilde{p}} \left(c^{n-1}\tilde{p} - c\theta\right) d\theta.$$
(12)

Solving the above expression with respect to \tilde{p} yields the critical price that defines the range from which the price targeted at the upper market segment will be drawn:

$$\tilde{p} = \frac{c^2}{1 - (1 - c)\sqrt{1 - c^{2n}}}.$$
(13)

Finally, we can derive the equilibrium distribution function that specifies the choice of p_{i1} . For this purpose, let $G(\cdot)$ and $g(\cdot)$ be the firms' cumulative distribution function and corresponding density function that describe the equilibrium distribution of p_{i1} . As we now know that firms will always choose their price p_{i1} from the interval $[\tilde{p}, \tilde{p}/c]$, this implies that $G(p_{i1}) = 0$ for $p_{i1} < \tilde{p}$, $G(p_{i1}) = 1$ for $p_{i1} \ge \tilde{p}/c$, and $G(p_{i1}) \in [0, 1]$ otherwise. An equilibrium in mixed strategies will now be characterized by $G(\cdot)$ if firm *i* is indifferent between setting either price p_{-i1} according to $G(\cdot)$. Calculating firm i's equilibrium expected profits $\Pi_i(p_{i1}, p_{-i1})$ and using the equilibrium condition $\Pi(p_{i1}, p_{-i1}) = \Pi_i^*$ then leads to a second-order differential equation whose solution determines $G(\cdot)$ (see Supporting Information).¹⁹

Proposition 3. In case of a duopoly market for a homogeneous service with contractible customer characteristics, the firms' equilibrium market segments are determined by the pricing schedule

$$p_{ii}^* = c^{j-1} p_{i1}^* \text{ for } j = 1, ..., n, \ n \in \mathbb{N} \setminus \{1\}, \ i = 1, 2,$$
 (14)

where the price targeted for the most upper market segment, p_{i1}^* , is drawn from the interval $[\tilde{p}, \tilde{p}/c]$ according to a welldefined cumulative distribution function $G(\cdot)$. Each contract selling for a price of p_{ij}^* also stipulates a maximum level of $\theta_{ij}^{max} = 1$ for j = 1 and $\theta_{ij}^{max} = p_{ij-1}^*$ for j > 1.

Expected equilibrium duopoly profits are identical for both firms and are equal to:

$$\Pi_{1}^{*}(n) = \Pi_{2}^{*}(n) = c \left(\frac{c^{2} \left(1 - c^{2n-2}\right)}{1 - c^{2}} F(\tilde{p}/c) + \frac{1}{2} \left(\frac{c^{n} (1 - c)}{1 - \sqrt{1 - c^{2n}} (1 - c)} \right)^{2} \right) > 0, \quad (15)$$

with $\frac{\partial^2}{\partial p_{i1}^2} F(p_{i1}) = G(p_{i1}).$

All in all, our results show that service providers will structure their portfolio of prices in the same way as a monopolist who is, however, faced with noncontractible rather than contractible customer characteristics. For j > 1, market segments are again shrinking in size for each segment added to sell the service to customers who previously did not buy the service.

What is especially remarkable concerning the equilibrium profits as given in Proposition 3 is that, even though both firms are symmetric, sell a range of service contracts that are homogeneous for every given customer, and compete in prices, they earn positive expected profits in equilibrium. This is in stark contrast to an ad hoc economic intuition. As shown before, the main driver of our result is that, because service providers have an incentive to reject certain customers that lead to unprofitably high service costs, the targeted market segment shifts to the left with every reduction in price. In contrast, if firms sold consumer goods, a lower price would simply widen but not shift the customer base, and the undercutting process will indeed come to an end when prices equal marginal costs and firms earn zero economic profits.

5.2 | Noncontractible customer characteristics

Finally, we elaborate on how our results of the previous subsection transfer to situations where θ is noncontractible. For this purpose, let us again start with n = 2 and first consider the case where firm 1 believes that firm 2 targets the upper side of the market and sets its market segments adjacent to one another. Then, akin to the argumentation in Section 5.1.2, this will result in an undercutting spiral with adjacent market segments.

Now, suppose that firm 1 expects firm 2 to set its prices such that they target the most profitable customers with the highest levels of θ but create an internal rivalry. What is firm 1's best response now? There are two cases.²⁰ First, if the intersection of firm 2's market segments is sufficiently large (see Figure 2), it will be profitable for firm 1 to choose to forgo the customers with the highest θ and only steal firm 2's demand for its second contract by marginally undercutting



FIGURE 2 Illustration of firm 2's market demands in the absence of firm 1 for prices p_{21} and p_{22} involving a large internal rivalry, and both firms' market demands when firm 1 sets p_{11} and p_{12} as a best response

 p_{22} . In this case, it will, of course, be optimal for firm 1 to set its own second contract adjacent to its first one. Hence, contracts with a sufficiently large intersection can never be a best response to any of its rival's strategies.

Given this reduced set of pricing strategies, we can now go on to consider the second case. If the intersection of firm 2 is sufficiently low (see Figure 3a), then it might be profitable for firm 1 to steal firm 2's total demand. In that case, it will have to create the same internal rivalry by marginally undercutting both of firm 2's prices. Our familiar undercutting spiral will now again continue until the specific price \tilde{p} , where it will be profitable for, let us say firm 1, to set $p_{11} = \tilde{p}/c$ (see Figure 3b). Of course, creating an internal rivalry between its own two contracts will then never be optimal anymore as it would just imply that the lower price p_{12} steals demand from the higher price p_{11} . Consequently, the firm will set its market segments adjacent to one another, and we are now back in the same undercutting spiral with no internal rivalries as before.

As in Section 5.1.2, we now use these insights to iteratively eliminate pricing strategies that can never be a best response, which leads to an equilibrium in mixed strategies in which firms set their prices such as to create adjacent market segments. All in all, this discussion implies that the results we derived for the duopoly market with contractible customer characteristics continue to hold when customer characteristics are noncontractible, and firms will structure their contracts in the same way as a monopolist in Section 4.2.²¹

Proposition 4. In case of a duopoly market for a homogeneous service with noncontractible customer characteristics, the firms' equilibrium pricing schedules, market segments, and duopoly profits are as given in Proposition 3.

6 | FIRMS' EQUILIBRIUM PRICING STRATEGIES AND MARKET SEGMENTATION WHEN *n* TENDS TO INFINITY

In this section, we abolish the condition that the number of contracts that firms offer has to be finite. In the limit,



FIGURE 3 (a) Illustration of firm 2's market demands in the absence of firm 1 for prices p_{21} and p_{22} involving a small internal rivalry, and both firms' market demands when firm 1 sets p_{11} and p_{12} as a best response. (b) Illustration of firm 1's best response when the undercutting spiral has reached a point where it becomes attractive to again target customers with the highest level of θ

each firm then offers infinitely many different contracts. In principle, this allows firms to target each customer individually and, as such, employ personalized pricing. As $\lim_{n\to\infty} \overline{\theta}_{i1} > 1^{22}$ and $\lim_{n\to\infty} \underline{\theta}_{in} = 0$, the whole market in the monopoly as well as in the duopoly setting will be served. However, it is only in the monopoly setting with contractible customer characteristics that firms will actually utilize the possibility of engaging in personalized pricing. This is because the size of the monopolist's market segments is uniform and approaches zero when *n* tends to infinity. In all other cases, the firms' market segments are shrinking in size with each contract added to a given portfolio. It is then only the size of the very last market segment *n* that tends toward zero, whereas the size of all other market segments remains positive. Therefore, firms will not use their ability to employ personalized pricing but create market segments where multiple customers are charged the same price for the service.

The following corollary summarizes the results for a monopolist and follows from Propositions 1 and 2:

Corollary 1. If a monopolist can offer an arbitrarily large number of service contracts, it will not use its ability to offer personalized prices when customer characteristics are non-contractible. The monopolist's profits are then equal to

$$\lim_{n \to \infty} \Pi_M^*(n) = \frac{1-c}{4}.$$
 (16)

When customer characteristics are contractible, the monopolist will employ personalized pricing. The monopolist's profits are then equal to

$$\lim_{n \to \infty} \Pi_M^*(n) = \frac{1-c}{2},\tag{17}$$

which are higher than when customer characteristics are noncontractible.

In competition, the possibility of firms to offer personalized prices basically resembles firms competing in infinitely many Bertrand markets with one customer (or a homogeneous group of customers) marked by characteristics θ . Even more so than in the previous Section 5, intuition suggests that this would result in the so-called Bertrand paradox: The best response to a firm setting its price *p* above (marginal) cost is for the other firm to lower its price just slightly to $p - \varepsilon$ and thereby steal the customer θ . Hence, for each customer θ , both firms would set their price equal to the marginal cost of providing the service, and total equilibrium profits are zero. At second glance, however, this Bertrand paradox (zero profit) outcome is not an equilibrium outcome in our service model.

The corollary below follows from Propositions 3 and 4:

Corollary 2. If firms can offer an arbitrarily large number of service contracts, they will not make use of their ability to offer personalized prices irrespective of whether customer characteristics are contractible or not. Expected equilibrium duopoly profits are positive for both firms and equal to

$$\lim_{n \to \infty} \Pi_i^*(n) = c^3 \frac{F(1)}{1 - c^2}.$$
 (18)

Our finding is in sharp contrast to the standard Bertrand competition game. Obviously, the crucial difference between the standard Bertrand argument with zero profits and our equilibrium outcome with positive expected profits is how we take the limit: The situation of Bertrand competition assumes an identical number of characteristics and available product offers and takes the joint limit. In contrast, our modeling assumes a continuum of characteristics but a finite set of different product offers and then takes the limit of infinitely many offers per firm. How can we discriminate between these two equilibria? One justification for our modeling is the simple observation that, in reality, the number of different characteristics is generally much greater than the number of products a firm can offer. Many customers differ in their individual willingness-to-pay, but firms are restricted by the costs of administering too many different product offers. This suggests that, in the limit with infinitely many product offerings, the equilibrium with positive profits is more convincing than the equilibrium with zero profits.

Overall, we uncover a significant advantage of the observability of customer characteristics for service providers. We show that, even in highly competitive industries, firms can earn positive expected profits in equilibrium. What is more, our results suggest that monopoly profits and expected duopoly profits increase in n, which implies that firms that are able to employ personalized prices are better off than when being limited to only charge $n \in \mathbb{N} \setminus \{1\}$ different prices. Intuitively, this follows from two effects: Targeting more segments allows firms to increase their profits by offering the service to those customers that previously did not buy. This positive effect stemming from a broader customer base is reinforced by the additional possibility to profitably increase prices for all preexisting contracts.²³ Note that a necessary assumption for our result is that firms post and commit to their prices ex ante, which is in contrast to prior studies assuming that personalized prices are determined at the time of the sale (e.g., Choe et al., 2018; Ghose & Huang, 2009).

Last but not least, whereas throughout the paper we assume that *n* is exogenous and the same across firms, we also want to briefly discuss the situation where firms, in the first stage, simultaneously choose the number of contracts they want to offer, $n_i > 1$, $i \in \{1, 2\}$, as well as the contracts' prices. For this purpose, let us assume that setting up an additional contract involves setup costs or increases administration costs. We denote these costs by $k(n_i) > 0$ and, for simplicity, assume that they increase linearly by a factor Δ_k , with $\Delta_k = k(n_i + 1) - k(n_i) > 0$.

For illustrative purposes, we first discuss the case where the maximum number of contracts that firms can offer is 3, that is, $n_i \leq n^{max} = 3$. There are then four potential equilibrium configurations: Both firms choose to offer only two contracts; both first choose to offer three contracts; and two asymmetric scenarios, where one firm offers only two contracts, whereas the other firm offers three contracts.²⁴ Consistent with our analysis in the paper, we again focus on symmetric equilibria only. Equilibrium expected profits given that both firms set the same number of contracts are given in Proposition 3 and we know that $\Pi_i^*(3) > \Pi_i^*(2)$. We first want to know whether $n_1 = n_2 = 3 = n^{max}$ can be an equilibrium. If one firm, say firm 1, would decide to unilaterally deviate to $n_1 = 2$, it would, of course, receive lower expected (gross) profits. It will, hence, only have an incentive to reduce its number of contracts if this loss is more than recovered by avoiding the costs of offering three instead of only two contracts.

Let us, next, suppose that $n_1 = n_2 = 2$ and consider again firm 1's incentive to unilaterally deviate to $n_1 = 3$. Firm 1 will then be a quasi-monopolist for its third contract and will always receive additional gross profits. It will, hence, have an incentive to deviate whenever these additional profits exceed the additional costs. Up to now, our discussion implies that if offering an additional contract was costless, that is, $\Delta_k = 0$, the only equilibrium would be $n_1^* = n_2^* = 3$ or, more generally, $n_i^* = n^{max}$, i = 1, 2.

Moreover, note that deviating from $n_1 = n_2 = 2$ leads to a gain in the size of the de facto gross monopoly profit from contract 3. Let us denote this gain by $\Pi^G > 0$. Deviating from $n_1 = n_2 = 3$, on the other hand, leads to a loss in the size of the gross duopoly profit from contract 3. We denote this loss in absolute terms by $\Pi^L > 0$. Because of the higher expected market coverage under monopoly than under duopoly, it is immediately apparent that the gain from unilaterally deviating from $n_i = 2$ to $n_i = 3$ is higher than the loss from unilaterally deviating from $n_i = 3$ to $n_i = 2$, i.e. $\Pi^G > \Pi^L$. With increasing costs of contract administration, $\Delta_k > 0$, we get the following symmetric equilibrium configurations, $i \in \{1, 2\}$:

$$n_i^* = 3 \quad \text{if } \Delta_k < \Pi^L,$$

$$n_i^* = 2 \quad \text{if } \Delta_k > \Pi^G.$$
(19)

Finally, let us discuss the more general case where $n^{max} \in \mathbb{N} \setminus \{1\}$. Note that Π^L and Π^G generally depend on n_i . As n_i increases, $\Pi^G(n_i)$ as well as $\Pi^L(n_i)$ decrease because firms obtain additional demand from less profitable customers with a lower willingness-to-pay. Therefore, even if we consider a situation where Δ_k is relatively low, there will be a threshold value $\overline{n_i} < n^{max}$ such that, eventually, $\Delta_k > \Pi^G(n_i)$ and, hence, $n_i^* < n^{max}$.

7 | SUMMARY AND CONCLUDING REMARKS

In the context of a service market, we analyze the ability of firms to offer different prices for a homogeneous service. Our novelty consists of accounting for a specific property that demarcates various services from products in that not only the customers' willingness-to-pay but also firms' costs of providing the service after the contract is signed depend on customers' observable idiosyncratic characteristics. In our model, both demand and supply are therefore positively associated with these customer characteristics. In addition to characterizing firms' optimal pricing schedules and resulting market segmentation strategies in a monopoly as well as a competitive market, a central message of our paper is that expected profits in equilibrium are positive even though firms compete in homogeneous services. Strikingly, this result holds even if firms can offer de facto personalized prices in equilibrium and, hence, compete in infinitely many Bertrand markets.

Although our paper isolates this competitive mechanism in the context of firms offering a single service, our findings can be directly transferred to several other markets. The crucial features of our model are also present in markets where firms bundle the service, also called ancillary service or service add-on, with a base product (see, e.g., Geng & Shulman, 2015; Song & Li, 2018). For example, in our previously

Production and Operations Management

mentioned example of the enterprise software application market, software providers often do not sell their product and support services separately but choose to bundle them in one contract. Again, depending on the customer's characteristics (e.g., the scope of application), the customer's willingnessto-pay as well as the frequency and complexity of the firm's support and training services vary. Airlines often bundle their core service of transporting customers from the point of departure to the final destination, for example, with on-board service, baggage transfer, and check-in services, where characteristics are associated with the customers' utilization of the service. In this context, our analysis uncovers a significant potential for firms in excessively competitive industries to escape the zero economic profit outcome. We show that firms selling a homogeneous product will be well advised to bundle their products with an ancillary service to earn positive expected profits in equilibrium. In addition, our model also translates to situations where firms are able to customize their product according to a customers' characteristics and incur customer-specific costs of customization. This, for example, also relates to the enterprise software applications market, which experiences a trend toward customizing the software to suit clients' needs (Ghose & Huang, 2009). A customer's scope of application might then relate to the desired variability of the software's features and consequently determines the firm's programming and implementation cost. Hence, our insights apply to a broad range of markets comprising professional services, consumer goods where a service is offered as an add-on as well as customized products.

There are several important issues we have left out in the current paper. We briefly discuss three of them. First, our model assumes that customer heterogeneity can be directly observed, so that market segmentation is possible even though firms sell a single homogeneous service. In other situations, however, this assumption is not satisfied because a firm may learn about the characteristics of its customers over time so that it can use this information when designing the future relationship with those customers (see Choe et al., 2018; Dell'Ariccia et al., 1999; Dell'Ariccia & Marquez, 2004). It would be interesting to see whether the mechanism described here is still valid. Second, in Section 4.1, we argued in a footnote that θ might also be the result of an underlying decision problem where customers marked by specific exogenous characteristics $\tilde{\theta}$ choose their optimal utilization levels. Future research might focus on a sophisticated microfoundation of an endogenous utilization level and the implications thereof. Finally, we abstract from service quality differentiation. Although this is a significant simplification of reality, it allows us to highlight the implications of the distinct characteristics of services on firms' pricing and market segmentation strategies. Because introducing additional heterogeneity regarding the valuation for service quality is beyond the scope of the paper, we leave this aspect for the future.

ACKNOWLEDGMENTS

We are grateful for the clear and constructive feedback from Haresh Gurnani, two anonymous reviewers, and, especially, the senior editor, which led to a substantial improvement of our paper. We also appreciate the helpful comments and suggestions from Eberhard Feess and Christoph Kuzmics.

ORCID

Peter-J. Jost https://orcid.org/0000-0002-0685-3069 *Anna Ressi* https://orcid.org/0000-0003-0950-0151

ENDNOTES

- ¹For an overview, see, for example, Stole (2007) and Armstrong (2006).
- ² If the price difference of two versions cannot be justified by the cost differences, then this strategy is also subsumed under the concept of second-degree price discrimination as coined by Pigou (1920).
- ³This is reminiscent to Schulte and Pibernik (2017), where a monopolist offers a menu of contracts, which stipulate a guaranteed service level and a corresponding price and customers differ in their shortage costs if their order is not fulfilled immediately. The firm's cost implication of those customer heterogeneities are captured by an operating model with inventory rationing.
- ⁴This is in contrast to other studies finding that competitive third-degree price discrimination can intensify competition; see, for example, Thisse and Vives (1988), Shaffer and Zhang (1995), and Corts (1998).
- ⁵The case of unobservable characteristics is analyzed in the stream of literature that builds upon Rothschild and Stiglitz (1976).
- ⁶In oligopoly pricing models, mixed strategy equilibria often arise; see, for example, Padilla (1992), Deneckere et al. (1992), or Allen and Hellwig (1986, 1989, 1993).
- ⁷We exclude the possibility that firms offer only a single contract, n = 1, because the market dynamics in the duopoly case vary (but not qualitatively) for n = 1 and n > 1 (see footnote 14). Further, we restrict $n \in \mathbb{N} \setminus \{1\}$ and, hence, allow firms to only target a limited number of market segments because, in reality, service providers often face a variety of legal restrictions and administrative costs that limit the number of contracts that they can profitably administer. We relax this assumption in Section 6, where we allow firms to offer infinitely many contracts. We then also briefly discuss the case where the number of contracts is not exogenously given but chosen by firms.
- ⁸ In the relevant case of noncontractible characteristics, our timeline implies that the bargaining power with respect to the choice of the contract lies on the customers' side. Alternatively, we could assume that firms hold the bargaining power. This would imply that, in stage 2, firms offer customers the contract that promises the highest profit, and, in stage 3, customers can only accept or reject. In a nutshell, the equilibrium outcome is then always the same as when characteristics are contractible. We will provide a more detailed explanation in footnotes 11 and 19. In the course of our analysis, we will show that the assumption regarding the contractability of customer characteristics does not affect the results of the duopoly market and does also not significantly change the main results of the monopoly market.
- ⁹Note that in case θ is contractible this might be somewhat misleading as one would think of the firm's market segment as comprising all customers with $\theta \in [\underline{\theta}_{ij}, \theta_{ij}^{max}]$. Our definition of a firm's market segment as $[\underline{\theta}_{ij}, \overline{\theta}_{ij}]$ will, however, prove to be useful in our subsequent discussions.
- ¹⁰ In this case, $\hat{\theta}$ constitutes a reasonably accurate contractible proxy for a customer's noncontractable utilization level and, as such, firms' service cost. If the utilization level is contractible, $\hat{\theta}$ might also be interpreted as the outcome of an underlying maximization problem of a customer marked by a certain exogenous characteristic $\tilde{\theta}$. In this case, θ would be an endogenous utilization level, and θ_{Mj}^{max} directly defines a maximum utilization level for each price.
- ¹¹Remember that we ordered the prices such that $p_{ij} \ge p_{ij+1}$.
- ¹²Let us come back to footnote 7. The reason why shifting the bargaining power to the monopolist leads to the same outcome as when customer characteristics are contractible should now be apparent. In this situation, the firm has the power to offer customers lying at intersecting market segments the higher price because those customers will always accept. Hence, the analysis is equivalent to the previous Section 4.1.

- ¹³As $\overline{\theta}_{M1} > 1$, the first market segment will be cut at 1. Hence, whether this statement is also true for j = 1 depends on c because $\frac{\partial p_{M1}}{\partial r} > 0$.
- ¹⁴The basic intuition for the optimal market segmentation is similar to Jost (2016). Because both firms are symmetric, we concentrate on symmetric equilibria and do not discuss uniqueness.
- ¹⁵Note that this price p'_{i1} then exactly targets those customers on the upper side of the market that remain unserved given that price \tilde{p} is charged. We relegate the derivation of \tilde{p} to the next subsection.
- ¹⁶Rearranging $\tilde{p} < p_{M1}$ yields the inequality $-\frac{(1-c)(n(c^2+z-1)+c(z(1-n)-1))}{(1-c^2+z-1)+c(z(1-n)-1))} > 0$ (1-z(1-c))(1+(1-c)n)

0. This reduces to $z < \frac{c+n-c^2n}{c+n-cn}$, which is clearly always satisfied. Similarly, rearranging $\tilde{p}/c > p_{M1}$ yields the inequality $-\frac{(1-c)(n(1-z)-cz-cn(1-z))}{(1-z(1-c))(1+(1-c)n)} > 0$, which reduces to n(1-z) - cz - cn(1-z) < 0 and is also always satisfied (see the Supporting Information).

- ¹⁷ If n = 1, the undercutting process will also cease to continue at \tilde{p} and jump to \tilde{p}/c . From there, it will, however, not directly start anew but will first jump to p_{M1} and then continue.
- ¹⁸ If firm 2 starts out with strategies that involve a gap between its two market segments, the line of argumentation is similar. We relegate this discussion to the next Section 5.2.
- ¹⁹Note that, because firm *i*'s profit function is different for $p_{i1} \in [\tilde{p}, c]$ and $p_{i1} \in [c, \tilde{p}/c]$, two second-order differential equations characterize the optimal solution, one for each interval.
- ²⁰If firm 2 starts out with strategies that involve a gap between its two market segments, the line of argumentation is analogous.
- ²¹Let us again briefly come back to footnote 7, where we argued that shifting the bargaining power from customers to firms leads to the same outcome as when customer characteristics are contractible. In a competitive market, firms will not be able to utilize their bargaining power like a monopolist and charge customers lying at intersecting market segments the higher price. This is because they will be disciplined by their competitor, who would then steal those customers at the intersection with one of its cheaper contracts. Hence, our analysis and results do not change.
- ²²Note that, in competition, $\lim_{n\to\infty} \tilde{p} = c$.

²³Regarding monopoly prices and profits, the following holds: If customer characteristics are contractible, $\frac{\partial p_{Mj}^*}{\partial n} = \frac{j(1-c)^2}{(1+n(1-c))^2} > 0$ and $\frac{\partial \Pi_M^*}{\partial n} = \frac{(1-c)^2}{2(1+n(1-c))^2} > 0$. If customer characteristics are not contractible, $\frac{\partial p_{Mj}^*}{\partial n} = -\frac{c^{2n+j}(1-c^2)\ln(c)}{(2c+(1-c)c^{2n})^2} > 0$. In the duopoly case, the range from which p_{i1} is drawn increases in *n*, that is, $\frac{\partial \tilde{p}}{\partial n} =$ case, the range from when p_{i1} is drawn increases in n, that is, $\frac{1}{\partial n} = c \frac{d\tilde{p}/c}{\partial n} = -\frac{(1-c)c^{2n+2}\ln(c)}{\sqrt{1-c^{2n}((1-c)}\sqrt{1-c^{2n}+1)^2}} > 0$, while the derivative of the expected price p_{i1} yields $\frac{\partial p_{i1}}{\partial n} = \frac{1}{c} \left(\frac{\partial \tilde{p}}{\partial n} - G(\tilde{p}/c) \right) + G(\tilde{p}) \frac{\partial \tilde{p}}{\partial n} = 0$. The formal proof for the result that $\frac{\partial \Pi_i^*}{\partial n} > 0$ can be found in the Supporting Information.

²⁴Remember that we restricted the number of contracts to n > 1.

REFERENCES

- Allen, B., & Hellwig, M. (1986). Bertrand-Edgeworth oligopoly in large markets. Review of Economic Studies, 53(2), 175-204.
- Allen, B., & Hellwig, M. (1989). The approximation of competitive equilibria by Bertrand-Edgeworth equilibria in large markets. Journal of Mathematical Economics, 18(2), 103-127.
- Allen, B., & Hellwig, M. (1993). Bertrand-Edgeworth duopoly with proportional residual demand. International Economic Review, 34(1), 39-60.
- Armstrong, M. (2006). Recent developments in the economics of price discrimination. In R. Blundell, W. Newey, & T. Persson (Eds.), Advances in economics and econometrics: Theory and applications (2nd ed., pp. 97-141). Cambridge University Press.
- Armstrong, M., & Vickers, J. (2001). Competitive price discrimination. RAND Journal of Economics, 32(4), 579-605.
- Arora, N., Dreze, X., Ghose, A., Hess, J. D., Iyengar, R., Jing, B., Joshi, Y., Kumar, V., Lurie, N. H., Neslin, S. A., Sajeesh, S., Su, M., Syam, N. B.,

Thomas, J., & Zhang, Z. J. (2008). Putting one-to-one marketing to work: Personalization, customization, and choice. Marketing Letters, 19(3-4), 305-321.

- Baye, M. R., & Morgan, J. (1999). A folk theorem for one-shot Bertrand games. Economic Letters, 65(1), 59-65.
- Baye, M. R., & Morgan, J. (2002). Winner-take-all price competition. Economic Theory, 19(2), 271-282.
- Bertrand, J. (1883). Theorie Mathematique de la Richesse Sociale. Journal des Savants, 67, 499-508.

Bester, H., & Petrakis, E. (1996). Coupons and oligopolistic price discrimination. International Journal of Industrial Organization, 14(2), 227-242.

- Champsaur, P., & Rochet, J. C. (1989). Multiproduct duopolists. Econometrica. Journal of the Econometric Society, 57(3), 533-557.
- Chen, Z., Choe, C., & Matsushima, N. (2020). Competitive personalized pricing. Management Science, 66(9), 4003-4023.
- Choe, C., King, S., & Matsushima, N. (2018). Pricing with cookies: Behavior-based price discrimination and spatial competition. Management Science, 64(12), 5669-5687.
- Choudhary, V., Ghose, A., Mukhopadhyay, T., & Rajan, U. (2005). Personalized pricing and quality differentiation. Management Science, 51(7), 1120 - 1130.
- Corts, K. S. (1998). Third-degree price discrimination in oligopoly: Allout competition and strategic commitment. RAND Journal of Economics, 29(2), 306-323.
- Dastidar, K. G. (1995). On the existence of pure strategy Bertrand equilibrium. Economic Theory, 5(1), 19-32.
- De Fraja, G. (1996). Product line competition in vertically differentiated markets. International Journal of Industrial Organization, 14(3), 389-414.
- Dell'Ariccia, G., Friedman, E., & Marquez, R. (1999). Adverse selection as a barrier to entry in the banking industry. RAND Journal of Economics, 30(3), 515-534.
- Dell'Ariccia, G., & Marquez, R. (2004). Information and bank credit allocation. Journal of Financial Economics, 72(1), 185-214.
- Deneckere, R., Kovenock, D., & Lee, R. (1992). A model of price leadership based on consumer loyalty. Journal of Industrial Economics, 40(2), 147-156
- Geng, X., & Shulman, J. D. (2015). How costs and heterogeneous consumer price sensitivity interact with add-on pricing. Production and Operations Management, 24(12), 1870-1882.
- Ghose, A., & Huang, K. W. (2009). Personalized pricing and quality customization. Journal of Economics & Management Strategy, 18(4), 1095-1135
- Harrington, J. (1989). A re-evaluation of perfect competition as the solution to the Bertrand price game. Mathematical Social Sciences, 17(3), 315-328.
- Hoernig, S. H. (2002). Mixed Bertrand equilibria under decreasing returns to scale: An embarrassment of riches. Economic Letters, 74(3), 359-62.
- Hoernig, S. H. (2007). Bertrand games and sharing rules. Economic Theory, 31(3), 573-585.
- Holmes, T. (1989). The effects of third-degree price discrimination in oligopoly. American Economic Review, 79(1), 244-50.
- Johnson, J. P., & Myatt, D. P. (2006). Multiproduct Cournot oligopoly. The RAND Journal of Economics, 37(3), 583-601.
- Jost, P.-J. (2016). Competitive insurance pricing with complete information, loss-averse utility and finitely many policies. Insurance: Mathematics and Economics, 66, 11-21.
- Kaplan, T. R., & Wettstein, D. (2000). The possibility of mixed-strategy equilibria with constant-returns to scale technology under Bertrand competition. Spanish Economic Revue, 2(1), 65-71.
- Kreps, D. M., & Scheinkman, J. A. (1983). Quantity precommitment and Bertrand competition yield Cournot outcomes. Bell Journal of Economics, 14(2), 326-337.
- Obermiller, C., Arnesen, D., & Cohen, M. (2012). Customized pricing: Winwin or end run. Drake Management Review, 1(2), 12-28.
- Padilla, A. (1992). Mixed pricing in oligopoly with consumer switching costs. International Journal of Industrial Organization, 10(3), 393-411.

Production and Operations Management

Pigou, A. C. (1920). The economics of welfare (1st ed.). Macmillan.

2852

Ross, S. L. (1984). Differential equations (3rd ed.). John Wiley and Sons.

- Rothschild, M., & Stiglitz, J. (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, 90(4), 629–649.
- Routledge, R. (2010). Bertrand competition with cost uncertainty. *Economics Letters*, 107(3), 356–359.
- Schulte, B., & Pibernik, R. (2017). Profitability of service-level-based price differentiation with inventory rationing. *Production and Operations Management*, 26(5), 903–923.
- Shaffer, G., & Zhang, Z. J. (1995). Competitive coupon targeting. *Marketing Science*, 14(4), 395–416.
- Shaffer, G., & Zhang, Z. J. (2002). Competitive one-to-one promotions. Management Science, 48(9), 1143–1160.
- Song, B., & Li, M. Z. (2018). Dynamic pricing with service unbundling. Production and Operations Management, 27(7), 1334–1354.
- Stole, L. A. (1995). Nonlinear pricing and oligopoly. Journal of Economics & Management Strategy, 4(4), 529–562.
- Stole, L. (2007). Price discrimination and competition. In M. Armstrong & R. Porter (Eds.), *Handbook of industrial organization* (3rd ed., pp. 2221– 2299). North Holland.
- Szech, N., & Weinschenk, P. (2013). Rebates in a Bertrand game. Journal of Mathematical Economics, 49(2), 124–133.

- Thisse, J. F., & Vives, X. (1988). On the strategic choice of spatial price policy. *American Economic Review*, 78(1), 122–137.
- Valletti, T., & Wu, J. (2020). Consumer profiling with data requirements: Structure and policy implications. *Production and Operations Management*, 29(2), 309–329.
- Zhang, Z. J. (2009). Competitive targeted pricing: perspectives from theoretical research. In V. R. Rao (Ed.), *Handbook of pricing research in marketing* (pp. 302–318). Edward Elgar Publishing.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

How to cite this article: Jost, P.-J., & Ressi, A.

(2022). What can I do for you? Optimal market segmentation in service markets. *Production and Operations Management*, *31*, 2838–2852. https://doi.org/10.1111/poms.13720