DISSERTATION / DOCTORAL THESIS

Titel der Dissertation /Title of the Doctoral Thesis

"Essays on Imperfect Competition"

verfasst von / submitted by

Radostina Shopova-van Linthoudt, BSc MSc

angestrebter akademischer Grad / in partial fulfilment of the requirements for the degree of

Doctor of Philosophy (PhD)

Wien, 2023 / Vienna 2023

Studienkennzahl lt. Studienblatt / degree programme code as it appears on the student record sheet:	UA 794 370 140
Dissertationsgebiet lt. Studienblatt / field of study as it appears on the student record sheet:	Volkswirtschaftslehre
Betreut von / Supervisor:	Univ.-Prof. Dr. Maarten Janssen Assoz. Prof. Daniel Garcia, PhD

Acknowledgements

I would like to thank my PhD supervisors Maarten Janssen and Daniel Garcia for their support and great advice throughout my studies. I am grateful for the opportunity to learn from their experiences and to work with them on my first research project. I am grateful to Maarten for inspiring me to explore the field of Industrial Organization during my master's studies. I am grateful to Daniel for being so supportive and patient during my first research steps.

I would like to thank Matan Tsur for being the first to notice the economist in me and for encouraging me to study economics.

I am especially thankful to Massimo Motta and Sandro Shelegia for the opportunity to visit them at UPF Barcelona and for the useful feedback and advice during my stay there. I am also thankful for the invitation to be part of the Workshop on Digital Platforms, where I had the chance to meet and talk with many top researchers in the field.

I am thankful to Omar Besbes, an associate editor, three anonymous reviewers, Marc Dudey, Robert Somogyi, and participants at the Vienna Graduate School of Economics seminar, the Theory Committee Meeting of the VfS, the Oligo Workshop 2021, the CEPR Applied IO Conference 2021, and the EEA-ESEM Conference 2021 for useful comments during the writing and publishing process of the first chapter of this thesis.

I am thankful to the editor Yossi Spiegel and two anonymous referees, as well as Özlem Bedre-Defolie, Francis Bloch, Federico Etro, Philippe Gagnepain, Clement Minaudier, Wieland Müller, Maximilian Schader, Karl Schlag, Sandro Shelegia, and Nikhil Vellodi, as well as the participants at the VGSE seminars, the workshop at the PSE Summer School in IO 2021, the 2022 MaCCI Annual Conference, the 5th Doctoral Workshop on The Economics of Digitization, the 22nd CEPR/JIE Applied IO Conference and 18th CEPR/JIE School on Applied Industrial Organization, as well as CRESSE 2022 and EARIE 2022 for the valuable comments that lead to the publishing of my second chapter.

Many thanks go to my colleagues at the Vienna Graduate School of Economics for fruitful discussions, shared laughs, and moral support during my PhD studies. Despite the Covid crisis, we had a lot of great times and made unforgettable memories.

Most importantly, I want to thank my husband, Jan-Michael, for being the shining light in my life, and for all the support, love, and laughs together, which make every day and challenge so much easier to face - I love you so much!

I thank the Austrian Science Fund FWF for its financial support under project FG6.

Contents

Introduction

The first fundamental theorem of welfare states that perfectly competitive markets, characterized by completeness and rational agents (consumers and firms) who take prices as given, are Pareto efficient. However, in reality, markets deviate from this ideal benchmark and violate the assumption of perfect competition. Agents frequently make strategic decisions that influence prices and the well-being of other market participants. Understanding imperfectly competitive markets, the nature of strategic decision-making, and their impact on economic outcomes is the central focus of the field of Industrial Organization.

The early development of the field of Industrial Organization can be characterized by two influential waves: the empirical 'Harvard tradition' and the theoretical 'Chicago tradition' (Tirole (1988)). The first wave primarily focused on empirical analysis and descriptive studies of markets. However, as this wave progressed, it revealed gaps and unanswered questions, particularly when market outcomes contradicted the stylized facts. To address these challenges, the second theoretical wave emerged, building upon the empirical observations and aiming to provide rigorous theoretical analyses.

Technological advancements, the use of algorithms, and the emergence of the digital economy have significantly altered the landscape of various markets, raising new concerns about the potential exploitation of market power by large firms. Empirical research has provided some insights into the outcomes, yet theoretical analysis remains crucial for understanding the incentives and behaviors of market participants. This thesis consists of three chapters that contribute to the field of Industrial Organization by offering theoretical analyses of contemporary issues arising from information asymmetry, market power, and its potential abuse.

The first chapter, joint work with Daniel Garcia and Maarten Janssen, focuses on markets, such as those for airline tickets and hotel accommodations. In these markets, firms sell time-dated products and have private information about their own unsold capacities. We show that competition under private information may explain observed phenomena, such as increased price dispersion and higher expected prices toward the deadline. We also show that private information severely limits the market power of firms and that information exchange about capacity increases firms' profits. Finally, we inquire into the incentives to unilaterally disclose information or to engage in espionage about rival's capacity and show that they increase firms' profits compared to the private information setting and are particularly harmful to consumers.

The second chapter addresses the concern of regulators that by introducing their own private labels, dominant online marketplace operators distort competition in their own favor. The chapter addresses this concern by studying how online marketplaces differ from classic retailers with a wholesale arrangement. In online marketplaces, individual sellers set their own consumer prices, while the marketplace operator collects fees from their sales. I show that when introducing a private label, the marketplace operator does not have an incentive to distort competition and foreclose the outside seller. On the contrary, when introducing a private label, there is an incentive to decrease the fee charged to the outside seller and to vertically differentiate its own product in order to protect the seller's channel. However, relative to the wholesale model of classic retailers, online marketplace operators offer a lower quality with higher consumer prices, leading to less improvement in consumer surplus and potentially less harm to the outside seller.

The third chapter expands on the findings of the second chapter to horizontal differentiation with multiple sellers. In particular, a hybrid platform that runs a marketplace for third-party sellers and offers its own product in competition with these sellers tends to set a lower fee than a pure platform with the decrease in the fee being connected to the level of competition and market power of sellers. Hybrid mode is strictly more profitable than solely operating a pure platform, as it enhances competition, expands demand, and avoids double marginalization, which also benefits consumers and increases total welfare. The platform does not have an incentive to exclude third-party sellers and monopolize their market space.

This dissertation provides valuable policy insights. The first chapter emphasizes the significance of shared information for regulators, particularly in light of the increasing use of algorithms. It highlights the potential for these algorithms to facilitate coordination among firms, potentially leading to adverse consequences for consumers. The second and third chapters shed light on the positive effects of platforms operating in hybrid mode, while also providing insights into how third-party sellers may be impacted. Furthermore, these chapters explore the influence of different vertical organizational structures and levels of competition in the market, demonstrating their relevance for platform incentives in accommodating or excluding market competitors.

Chapter 1

Dynamic Pricing with Uncertain Capacities[1]

1.1 Introduction

In markets such as those for hotel accommodations, airline flights, shipping, generated electricity, or any other time-dated product, firms sell a fixed capacity with a deadline. A large literature in management and economics has devoted attention to this issue, focusing mostly on a monopolistic seller who sets prices to serve consumers who arrive gradually over time.[2] Starting with Dudey (1992) there is also a smaller literature dealing with competition (see also Martínez-de-Albéniz and Talluri (2011) and Farias et al. (2011)). While it is often the case that firms are uncertain about their competitors' unsold capacity, this literature assumes that capacity is commonly known. We close this gap by providing the first equilibrium analysis of dynamic competitive pricing of time-dated products where information about unsold capacity is private.[3]

To focus the analysis on the strategic implications of private information about capacities, we start with a stylized model where two firms compete in two periods and demand is known and constant in both periods. A firm can be either constrained or unconstrained. A constrained firm's capacity is just enough to serve the demand in one period, whereas an unconstrained firm can serve consumers in both periods. The model is simplified version of the model proposed by Dudey (1992) with the main difference that the firms have

[1]Joint work with Maarten Janssen and Daniel Garcia. The chapter has been published in the journal Management Science and has been circulated as a CEPR Discussion Paper No. 1567 as a working paper version. This chapter builds on the master thesis "A game of spies and uncertain capacities" by Shopova (2019).

[2]See, e.g., Talluri and Van Ryzin (2006) for an overview of early literature and, e.g., Hörner and Samuelson (2011) and Board and Skrzypacz (2016) for important recent contributions, where optimal contracts are determined depending on whether or not a firm can commit and where buyers are forward-looking.

[3]A paper by Lin and Sibdari (2009) proposes a heuristic pricing policy based on the Nash equilibrium prices derived from the game with complete information, but they do not consider the strategic equilibrium pricing decisions under private information.

private information about their capacity at the beginning of the game. Firms set prices in both periods, and consumers buy from the lowest-priced firm in the period in which they arrive.[4] We show that the market equilibrium in the second period critically depends on the belief that the firm that did not sell in the first period has about the capacity of its rival. This belief depends on the structure of the market equilibrium in the first period. If the equilibrium would display price signaling, firms charge different first-period prices depending on their capacity, allowing their rivals to infer their private information from the prices they charge. On the other hand, if firms would set their price independent of whether they are constrained or unconstrained, then the equilibrium is pooling and no inference can be made.

The contribution of our paper is twofold. From a positive perspective, we provide a new supply-side explanation that can reconcile several empirical observations. First, prices tend to increase and become more dispersed as the deadline approaches (see, e.g., McAfee and Te Velde (2006) and Clark and Vincent (2012)).[5] Second, there is no clear, monotonic relation between price dispersion and whether demand is peak or off-peak (see, e.g., Puller et al. (2009)). This second observation is hard to reconcile with other supply-side explanations (e.g. Dana (1999) and related literature). Our simple model of a market in which firms are uncertain about their rivals' capacity suffices to explain these patterns without introducing heterogeneous behavior on the demand side.[6] Our explanation emphasizes that firms want to limit information revelation in periods long before the deadline to generate more profits in periods closer to the deadline.

From a normative perspective, we show that firms' profit margins are significantly lower (and consumer surplus significantly higher) if they are uncertain about their rivals' capacities. Conversely, we show that if firms manage to exchange information, they can increase their profits — at the expense of consumers. This finding may help explain the recent adoption of information sharing protocols, such as the online publication of real-time seat maps.

We now explain these results in more detail. Given some market outcome in the first period, both firms know which firm sold one unit and at which price. This information enables the non-selling firm to form a posterior belief about the residual capacity of the selling firm and, therefore, estimate the probability that it will face competition in the second period. Following the analysis in Janssen and Rasmusen (2002) price dispersion is an essential feature in the second period as the non-selling firm is uncertain about whether

[4]We also consider several extensions to this framework, including multiple periods, strategic consumers, random arrival and captive consumers.

[5]It is not always true that transaction prices are higher closer to the deadline. For example, an Airbnb host may provide a last-minute discount if her listing has not been booked closer to the available date (Huang (2022)). As we explain later, this feature is consistent with some of the equilibria of our model.

[6]In the context of their model, McAfee and Te Velde (2006) conclude that "changing demand is a salient feature of the data, and models that assume that the shape of demand is constant over time are empirically invalid." We show that uncertain capacities naturally result in market outcomes that resemble those in markets with changing demand. Furthermore, we show that our results still apply if valuations increase towards the deadline.

the competitor can actively compete. The expected equilibrium prices and profits are determined by this uncertainty, as it is the only source of market power.

Turning to the first period, we characterize the unique pooling equilibrium where both firms charge the same price. This price equals the opportunity cost of selling in the first period for a constrained firm, given its prior belief about its rival type. In addition, there exists a continuum of semi-separating equilibria where constrained and unconstrained firms randomize their pricing decisions in both periods. In all these semi-separating equilibria, first-period prices reveal some information about the rival's capacity but are not fully informative. Firms make more profits in the first period than in the unique pooling equilibrium, but their second-period profits are lower, driven by a lower average posterior belief. It turns out that the effect on second-period profits dominates so that the pooling equilibrium Pareto-dominates these semi-separating equilibria, *i.e.*, firms suffer from revealing information through first-period pricing. In some of these semi-separating equilibria, the effect of information revelation is so strong that the expected prices in the second period are lower than first-period prices.[7] Still, in all equilibria, there is less price dispersion in periods further away from the deadline as firms try to hide their private information. Similarly, price dispersion is highest in markets with most uncertainty and lowest when firms are relatively certain that capacity is either too high (off-peak demand) or too low (peak demand). Most importantly, the pooling equilibrium yields considerably lower profits than what Dudey (1992) predicts in his model where firms' capacities are public information.

We extend this basic framework to accommodate several features of real-world markets. First, we show that all our results are unchanged if consumers are patient and strategically choose to wait for prices to fall. Second, our results also survive with minor modifications if consumers' valuation increases as we approach the deadline or if their arrival process is stochastic. Third, we show that the pooling equilibrium we identify in the baseline model survives if we extend the number of periods in which trade occurs. Finally, we consider an extension of the model whereby some consumers have a (strong) preference for each of the firms and show that the mixed (semi-separating) equilibria we identified in the baseline model extend nicely to this setting. These results make clear that the basic mechanisms identified in the baseline model are robust to several natural extensions.

In real-world markets, firms may voluntarily disclose information regarding their unsold capacities or, alternatively, gather information about rival's capacities and one may wonder whether private information regarding capacities will remain if firms can engage in such activities. Airlines frequently announce (disclose) how many seats they have left at the current price they offer. Alternatively, firms may learn the unsold capacity of their

[7] This multiplicity allows the model to reconcile a wide range of findings, including those discussed in footnote 3.

rivals through industrial espionage, for example, by scraping online data sources.[8] Industrial espionage is in many ways the reverse of voluntary disclosure as the firm that takes the initiative and bears the cost gets informed about the rival's capacity. An important difference is that, unlike the public nature of disclosure, espionage cannot be observed so that it is the *expectation* of industrial espionage that drives competitors' behavior rather than the act itself. We show that for any positive disclosure or spying cost, no matter how small, private information remains an important element in understanding markets with time-dated products. Not surprisingly, if the cost of disclosing or engaging in espionage is sufficiently large, no firm wants to engage in voluntary disclosure or in industrial espionage. More interestingly, if this cost is smaller, only one firm decides to disclose or acquire information, and it is the firm whose information is not revealed that benefits most. Moreover, *ex ante* market power of firms is higher than when firms have to act not knowing the private information of their rivals and even higher than when both firms exchange information![9] One interesting aspect of the results on espionage is that even if both firms know their rival still has unsold capacity in the last period, they do not engage in marginal cost pricing as firms may (second-order) believe that their competitor believes that they are sold out.[10]

1.1.1 Related Literature

As indicated at the start of the paper, to the best of our knowledge there does not exist a paper providing an equilibrium analysis of dynamic pricing of time-dated products where capacities are private information. Of course, there is a large literature in revenue management studying different aspects of dynamic pricing problems. Restricting ourselves to the literature introducing competition between firms, the seminal paper by Dudey (1992) has been extended in various directions. Dasci and Karakul (2009) consider two customer segments with different reservation values. Uncertainty about the size of demand has been addressed by Martínez-de-Albéniz and Talluri (2011). Gallego and Hu (2014) analyze firms producing a mix of substitutable and complementary goods, while Xu and Hopp (2006) study inventory choices prior to competition. Strategic consumers are considered in Levin et al. (2009) and Anton et al. (2014), while some of these papers also introduce more than two firms competing with each other. To focus on the implications of private

[8]In Canada, WestJet seems to have engaged in industrial espionage when it acquired "access to a special reservation website for the employees and retirees of Air Canada" The Globe and Mail reported in July 2004. The article reports on a lawsuit initiated by Air Canada and continues that the "website contains confidential information about the number of passengers booked on all flights at Air Canada and its subsidiary Zip for up to 352 days in the future". In May 2006, CBC News reports that a settlement was reached in which "WestJet apologized to Air Canada" and that "the lawsuit centered on allegations that WestJet management used the password of a former Air Canada employee to access a website maintained by Air Canada to download "detailed and commercially sensitive" information."

[9]There is some evidence that firms engage in communication about capacities, albeit on a more long-term basis. See, in particular, Aryal et al. (2022).

[10]In other words, Bertrand competition requires that firms have *common knowledge* about unsold capacity levels of all firms.

information about capacities, our baseline model abstracts away from all these factors, without implying that these factors are unimportant in real-world markets. In section 1.4.3 we show how our results are robust to re-introducing quite a few of these factors.

Two recent papers (Somogyi et al. (2023), Montez and Schutz (2021)) study the implications of private information about capacities in a static pricing environment. Somogyi et al. (2023) find that capacity-constrained firms price less aggressively than unconstrained firms as they prefer to focus on any left-over demand in case their rival is also capacity-constrained. Although we find that a similar consideration is relevant in some of our equilibria, the dynamic game allows for much richer equilibrium patterns, allowing us to explain the stylized facts and to derive welfare conclusion about information sharing. Montez and Schutz (2021) consider a model of price competition where the pre-committed capacities are unobserved. They show that the fact that capacity choices are typically not known to competitors drastically affects the prediction that Cournot type of behaviour should be expected in markets where firms produce in advance. Our model does not allow for endogenous capacity choice, yet has dynamic pricing instead, and shares the message that not knowing capacity choices of rivals negatively affects firms' market power.[11]

The paper is also clearly related to the extensive literature on disclosure (see, e.g., Dranove and Jin (2010) for an overview of the quality literature) and the much smaller literature on industrial espionage (see, e.g., Solan and Yariv (2004) and Barrachina et al. (2014)). These topics, of course, only cover part of our paper, and none of these papers investigate the incentives to disclose or spy on capacities.

The paper is also related to the literature on information exchange (see, e.g., Gal-Or (1985) and Shapiro (1986) for early contributions) and the competition policy issues related to information exchange (see, e.g., the Horizontal Guidelines by the EU Commission (2011) or the OECD (2010) report). It is common wisdom that information exchange can have both a positive and negative impact on the competitiveness of the market, depending on the type of information exchanged. We show that in markets for time-dated products information exchange about unsold capacity is unambiguously anti-competitive and that, in particular, with information exchange firms can achieve collusive outcomes even if they choose prices unilaterally.

Finally, there is an interesting relation between our paper and the literature on sequential auctions with budget constraints (see, e.g., Pitchik (2009) and Balseiro et al. (2015)). An alternative way to interpret our model is that two players compete in two sequential auctions for two objects with a common value and that they have a privately known budget which is either equal to the value of one object or the value of two objects. The sequential auction literature differs in that it is important *how much* budget a bidder has left to compete in a later auction and not only whether a bidder can compete. This difference allows us to fully characterize the set of equilibria and derive their properties.

[11] Similarly, Fabra and Llobet (2023) study a uniform price auction where firms are privately informed about their (random) capacities. Interestingly, they also find that uncertainty about rivals' capacities leads to less excise of market power and higher consumer surplus.

The rest of the paper is organized as follows. The next section presents the baseline model with private information, while Section 1.4 presents the results of the baseline model and the extensions. In between, Section 1.3 briefly analyzes the complete information model as a benchmark. Section 1.5 develops the arguments pertaining to disclosure, while Section 1.6 deals with espionage. Section 1.7 concludes. Proofs that are essential for a proper understanding of the results are given in the main body of the paper; other proofs are relegated to Appendix 1.A. Detailed discussions on equilibria under espionage are in Appendix 1.B, while some additional material can be found in Appendix 1.C.

1.2 The Basic Model and Solution Concept

The baseline model builds upon the work of Dudey (1992). We consider a homogeneous goods market where two firms compete in prices over time. The demand side is composed of myopic consumers with unit demand, and we normalize their willingness to pay to 1. Consumers enter the market in the first or the second period, observe the prices charged in the market at that time, and buy at the lowest price if that is below their willingness to pay. Otherwise, the consumer leaves the market and does not come back. In case the observed prices are the same, all consumers buy from one of the firms (with each selling with equal probability).[12] For ease of notation, we normalize the market size in each period to 1.

Each firm has an initial capacity which is private information. To make things interesting, a firm is either constrained if it can only sell to one consumer, or unconstrained if it can cover the demand in both periods. Firms' production cost is normalized to zero. The prior probability a firm is constrained is denoted by α and is independent across firms.[13]

$t = 0$

- Nature draws firms' type;
- Firms learn their type.

$t = 1$

- Firms set p_{i1};
- A consumer arrives and buys at cheapest firm iff $p_1 < 1$;
- Firms observe prices and infer who sold.

$t = 2$

- Firms update beliefs about rival's capacity;
- Active firms set p_{i2};
- A consumer arrives and buys at cheapest firm iff $p_2 < 1$;
- Profits are realized.

Figure 1.1: Timeline of the events in the market.

[12] Note that under an alternative tie-breaking rule whereby both firms sell half of their capacity in the first period if they set identical prices, the continuation game is outcome-equivalent to the one we study here but more complicated to analyze. Details are given in Appendix 1.C.1.

[13] For the case of asymmetric priors, see Footnote 30.

The game unfolds as follows. In period 1, depending on their capacity each firm i chooses a price p_{i1}, $i = 1, 2$. The consumer that entered the market in period 1 observes both prices and buys at the lowest price, provided that price is not larger than 1. Firms observe both first-period prices[14] and know from whom the consumer bought in the first period (if at all). At the beginning of the second period, they update their beliefs about whether or not the rival firm has unsold capacity left based on the price of the rival firm. All firms with unsold capacity set a price p_{i2}, $i = 1, 2$ in the second period. The period 2 consumer observes both prices p_{i2} and also buys at the lowest price if that is not larger than 1. Firms choose their prices to maximize the sum of profits in both periods. This timeline is depicted in Figure 1.1.

We solve the model using Perfect Bayesian Equilibrium as a solution concept (see, e.g., Fudenberg and Tirole (1991)).

Definition. *A Perfect Bayesian Equilibrium (PBE) of this game is a combination of (possibly mixed) strategies $(p^c_{i1}, p^u_{i1}, p_{i2}(p_{i1}, p_{j1})), i = 1, 2, j \neq i$, where p^c_{i1} and p^u_{i1} are the first-period prices of firm i in case it is constrained, respectively unconstrained, and a set of beliefs $\theta(p_{j1})$ that the non-selling firm i has about the rival firm $j \neq i$ being constrained based on the first-period price p_{j1} such that*

- *at the beginning of period 2, the non-selling firm i updates its belief $\theta(p_{j1})$ using Bayes' Rule when possible*;
- *for every pair of prices (p_{11}, p_{21}) each firm i that still has unsold capacity left in the second period chooses price p_{i2} so as to maximize second-period profit given the competitor's second-period strategy and the updated belief $\theta(p_{j1})$*;
- *each firm i chooses a pair of first-period prices (p^c_{i1}, p^u_{i1}) so as to maximize the sum of expected profits in both periods giving the first-period prices (p^c_{j1}, p^u_{j1}) chosen by the rival and the equilibrium price reactions $p_{i2}(p_{i1}, p_{j1}), i = 1, 2, j \neq i$.*[15]

When appropriate, we also show that the PBE we analyze also satisfy the Intuitive Criterion.

As discussed in the Introduction, we acknowledge that this model abstracts from many features that are important in real-world markets, but we first want to study the strategic implications of private information about capacities in its simplest possible form. Subsection 1.4.3 analyzes several extensions of this baseline model.

[14] In many markets where firms and consumers make online transactions, it is easy for firms to observe competitors' prices: if consumers can observe both prices, it is typically also possible for firms to observe them. If we dispose of this assumption, the analysis becomes more involved as firms must form conjectures about the prices charged by competitors. In Appendix 1.C.2 we show, however, that the pooling equilibrium characterized in Section 1.4 remains an equilibrium when first-period prices are hidden from competitors.

[15] Note that the second-period prices do not depend on whether a firm is constrained or not as the firm can only sell one unit in the last period. The only thing that is still relevant at that stage is whether a firm still has a unit of capacity left to sell.

1.3 Revisiting the Complete Information Results

Before analyzing the private information game, in this section, we briefly revisit the results for the complete information game. This is essentially the Dudey (1992) model but simplified to two periods and firms either having one or two units to sell. When there is common knowledge that both firms are constrained, they will set monopoly prices when they have unsold capacity left. When there is common knowledge that both firms are unconstrained, they will set prices equal to marginal cost in both periods. Finally, if it is commonly known that one firm is constrained, while the other is unconstrained, we have the following result.

Proposition 1.1. *If it is common knowledge that one firm is constrained while the other is not, then there exists a unique subgame perfect equilibrium in undominated strategies where the constrained firm sells at* $p_1^c = 1$ *in the first period, and the unconstrained firm sells for* $p_2^u = 1$ *in the second period. In addition, there exists a continuum of subgame perfect equilibria in weakly dominated strategies, indexed by* $x \in (0, 1)$*, where the constrained firm sells at* x *in the first period and the unconstrained firm sells at* 1 *in the second period.*

To understand why there are multiple equilibria, it is best to provide the subgame perfect equilibrium strategies and to recall the result of Blume (2003) who shows that in the Bertrand model of competition with homogeneous goods, there exists an equilibrium where one firm sets the price equal to the marginal cost, while the other firm randomizes uniformly over a very small interval above this price. With this result in mind, suppose that the constrained firm sets $p_1 = x$ for some $x \in (0, 1]$, while the unconstrained firm uniformly randomizes its first-period price in the interval $(x, x + \varepsilon)$ for some small $\varepsilon > 0$. The constrained firm makes an equilibrium profit of x, while the unconstrained firm makes a profit of 1 as it can sell at the monopoly price in the second period. If the constrained firm does not sell in the first period, firms engage in marginal cost pricing in the second period.

It is straightforward to verify that there are no profitable deviations. If the constrained firm sets a higher price, there is a large probability that it will not sell in the first period and will face Bertrand competition in the second period. Thus, this deviation results in a profit close to 0. Selling at a lower price in the first period is also not profitable. If the unconstrained firm undercuts the constrained firm in the first period, it will also face Bertrand competition in the second period, making the deviation unprofitable.

The unconstrained firm does not have a strict incentive to set first-period prices smaller than 1, however. Moreover, if the constrained firm deviates and sets $p_1 \in (x, 1)$, the unconstrained firm would be better off setting a price larger than p_1. Thus, setting a price $p_1 < 1$ is a weakly dominated strategy for the unconstrained firm. Therefore, for the benchmark model of complete information, we continue focusing on the Dudey (1992) monopoly outcome as that is the unique equilibrium in undominated strategies. In Section 1.6 we will refer, however, to the weakly dominated equilibria to explain that equilibria of the espionage game do not converge to the Dudey (1992) monopoly outcome if the cost of espionage approaches 0. For future reference, the ex-ante expected profits to the firms under public information is equal to $1 - (1 - \alpha)^2 = \alpha(2 - \alpha)$: firms make a profit of 1 in all cases except when both are unconstrained.

1.4 The Implications of Information Asymmetry

To understand the importance of private information, it is instructive to understand in more detail how private information about capacity undermines Dudey's monopoly result. If the Dudey outcome was supported by an equilibrium in the game with uncertainty about capacity, it must be that constrained firms set a price of 1 in the first period, and unconstrained firms set a higher price. Given these first-period prices, in the second period, the firms would set prices of 1 if at least one firm set a price of 1 in the first period and a price equal to marginal cost if both firms set a larger price. In such a case, the unconstrained firm has incentives to deviate and imitate the constrained firm in the first period. There are two benefits of doing so. First, the unconstrained firm has the chance of being able to sell in the first period at the maximum possible price, increasing its expected first-period profits. Second, this deviation leads the rival to set the monopoly price of 1 in the second period because it expects to be a monopolist. As a result, the second-period profit of the deviating firm is approximately equal to the monopoly profit. Hence, the unconstrained firm accrues extraordinary rents by imitating the constrained firm. It follows that the Dudey equilibrium requires a level of coordination between constrained and unconstrained firms that is impossible to achieve under private information.

The above argument extends more generally. In particular, no separating equilibrium with constrained and unconstrained firms setting different prices in the first period exists. Using the fact that if first-period prices reveal the information of the selling firm, then second-period prices are either equal to 0 or 1, the main argument shows that incentive compatibility cannot hold: either the unconstrained or the constrained firm wants to imitate the first-period price of the other type. Thus, our first result follows:

Proposition 1.2. *A separating equilibrium does not exist.*

Proof. The proof is given in the Appendix.

We will next construct a pooling equilibrium where both types choose the same first-period price. As a generalization of this equilibrium plays an important role also in the next sections, we discuss its construction in some detail for the more general case where the posterior probability that the rival is constrained given the rival's first-period price is denoted by $\theta \in (0, 1)$.[16] We first focus on the second period following such a pooling

[16]In the special case of the pooling equilibrium we analyze in this section, the first-period price does not reveal any new information to the firms so that Bayes' Rule yields $\theta(\alpha) = \alpha$. In later sections where we analyze other equilibria, θ will generally not be equal to α. To stress that this posterior probability may depend on the price that is charged, in later sections we also write $\theta(p)$.

outcome in the first period and let the cumulative price distributions in the second period be denoted by F^S and F^N for firm S (who *sold* in period 1) and firm N (who did *not sell* in period 1), respectively. In the second period, it is common knowledge that firm N still has a unit to sell, whereas firm S is sold out with probability θ. Thus, as in Janssen and Rasmusen (2002), with probability θ firm N is a monopolist in period 2, while it faces a competitor with probability $1-\theta$. It is clear from their analysis that a Nash equilibrium in pure strategies does not exist and the same argument applies to our setting. By setting a price $p_2 \leq 1$ in the second period, firm N has an expected profit of $\pi_2^N(p_2) = \theta p_2 + (1-\theta)\left(1 - F^S(p_2)\right)p_2$: with probability θ it is a monopolist and always sells, while with the remaining probability, the firm only sells if the competitor sets a larger price. As firm N gets an expected profit of θ when setting $p_2 = 1$, it is easy to see that to make firm N indifferent, the unconstrained firm S must randomize according to

$$F^S(p_2) = 1 - \frac{\theta(1-p_2)}{(1-\theta)p_2}, \tag{1.1}$$

with $p_2 \in [\theta, 1)$. Similarly, the expected second-period profit of an unconstrained firm S equals $\left(1 - F^N(p_2)\right)p_2$. To make an unconstrained firm S indifferent between any $p_2 \in [\theta, 1)$ it must be that

$$F^N(p_2) = 1 - \frac{\theta}{p_2} \tag{1.2}$$

with a mass point of θ at 1.[17] Thus, importantly, provided they still have unsold capacity left both firms make an expected profit of θ in the second period and to do so they randomize their pricing decisions in that period if the uncertainty concerning their rival's capacity is not fully resolved. As in a pooling equilibrium $\theta = \alpha$ a constrained firm is not willing to accept a price below α in the first period for its sole unit, as by deviating and setting $p_1 > 1$ and letting the other firm sell first it can expect at least α in the second period.

To finalize the construction of a candidate pooling equilibrium, we now argue that it must be that the first-period equilibrium price p_1^* equals α. It is clear that p_1^* cannot be smaller than α as a constrained firm would want to deviate: by setting a higher price it does not learn anything and expects a profit of α in the second period. On the other hand, by having $p_1^* > \alpha$, the constrained firm would get an expected pay-off of $\pi^{*1} = \frac{1}{2}p_1^* + \frac{1}{2}\alpha$ (as it would sell with probability $\frac{1}{2}$ in period 1 and gets an expected profit of α in period 2 if it did not sell in period 1), and would like to undercut p_1^*. Thus, it must be that $p_1^* = \alpha$.

We now argue that no type of firm wants to deviate from this candidate pooling equilibrium. It is clear that neither type wants to deviate upwards as they then forego

[17] Note that $F^N(p_2)$ is different from the distribution function found in Janssen and Rasmusen (2002) as they only deal with symmetric information, whereas in our setting it is important to realize that the selling firm knows that there is competition in period 2 (whereas the firm that does not sell does not know).

the possibility of selling in the first period and the above argument on second-period pricing implies their expected profit in that period equals α as the firm that sells is constrained with probability α.[18] In addition, it is clear that the constrained firm does not want to deviate downwards. Whether an unconstrained firm wants to deviate downwards depends on how one specifies the out-of-equilibrium beliefs. If it sticks to the pooling price, it expects an overall equilibrium profit of $\frac{3}{2}\alpha$. If it undercuts the first-period price by setting $p_1 < \alpha$ and induces an out-of-equilibrium belief θ', it expects a pay-off of $p_1 + \theta'$. Thus, it is not profitable to deviate if $\theta' \leq \frac{\alpha}{2}$. Such a belief is actually very reasonable. For example, the Intuitive Criterion implies that $\theta' = 0$ as only the unconstrained type can possibly benefit from undercutting.

Thus, we have proved the following:

Proposition 1.3. *There exists a unique symmetric*[19] *pooling equilibrium, where both types of firms charge* $p_1^* = \alpha$. *In the second period firms' prices satisfy* $F^S(p_2) = 1 - \frac{\alpha(1-p_2)}{(1-\alpha)p_2}$, *and* $F^N(p_2) = 1 - \frac{\alpha}{p_2}$ *for* $p_2 < 1$ *and a mass point of* α *at* 1.

The ex-ante expected profit in this pooling equilibrium is equal to $\alpha[\alpha + \frac{3}{2}(1-\alpha)] = \alpha\frac{3-\alpha}{2}$. This expression can be easily understood by looking at a firm that charges α in both periods. The ex-ante probability of selling is then given by 1 if the firm is constrained (which happens with probability α), while if the firm is unconstrained it sells with probability 0.5 in period 1 and sells for sure in period 2 (which happens with probability $1-\alpha$). Note this profit is strictly smaller than $\alpha(2-\alpha)$, which, as we have seen in the previous section, is the profit under complete information. This is an important step in our overall claim that under private information, firms have less market power than under full information.

1.4.1 Semi-separating Equilibria

We will now argue that, in addition to the pooling equilibrium, there exists a continuum of semi-separating equilibria. In a semi-separating equilibrium, both types of firms charge prices in the same support, albeit they may do so with different probabilities. As a result, firms imperfectly learn from first-period prices as different first-period prices are associated with different posterior beliefs and, consequently, different outcomes in the second period.

Letting $Q(p)$ denote the probability that a firm charging p sells in the first period and noticing that $\theta(p)$ is determined on the equilibrium path by the probability distributions that different types use, the expected profit of a constrained firm is equal to

[18] The second-period profit is determined by the belief of the deviating non-selling player regarding the probability that the non-deviating selling player is constrained. Clearly, given the selling player chose his equilibrium price in the first period, Bayes' Rule implies this belief should be equal to α.

[19] There also exist two asymmetric pooling equilibria, in which independent of their type one of the firms charges α and the other mixes in the interval $(\alpha, \alpha + \varepsilon)$. This equilibrium is outcome equivalent to the pooling equilibrium.

$$\Pi^c(p) = Q(p)p + \int_{\underline{p}}^{p} \theta(p')d(1 - Q(p')), \tag{1.3}$$

where $\underline{p}$ is the lowest price in the interval of prices charged in the second period. This expression can be understood as follows. If the rival firm charges some price $p' > p$, which happens with probability $Q(p)$ the constrained firm will sell and exit. Otherwise, if the rival firm charges some price $p' < p$, the constrained firm may still sell in the second period. In that period, it follows from the discussion above Proposition 1.3 that its profit is equal to the posterior belief $\theta(p')$ of the selling rival being initially constrained and sold out.

The expected profit of an unconstrained firm is equal to

$$\Pi^u(p) = Q(p)(p + \theta(p)) + \int_{\underline{p}}^{p} \theta(p')d(1 - Q(p')). \tag{1.4}$$

This expression can be understood by realizing that the unconstrained firm's profit includes an additional term reflecting the value of holding an extra unit. If the unconstrained firm does not sell in the first period, its continuation profit is equal to that of a constrained firm that did not sell. If, however, an unconstrained firm sells in the first period, it can still sell in the second period. From the discussion above Proposition 1.3 it follows that the second-period profit in this case equals the posterior belief $\theta(p)$.

In a semi-separating equilibrium, for all prices in the support, the profit of the constrained and the unconstrained firm must be constant. This implies that $Q(p)\theta(p)$ must be equal to a constant, and so the posterior belief must be an increasing function of the price (since the probability of selling $Q(p)$ decreases in p). Since the profit of this additional unit must be positive, the unconstrained firm will only charge prices that sell with positive probability in the first period. As a result, the upper bound of the price distribution must have an atom.

Using these insights and the randomization conditions, we are able to characterize the equilibrium strategies. The results are summarized in the following proposition:

Proposition 1.4. *For every $\overline{p} \in (\alpha, 1]$ there exists a semi-separating equilibrium such that both types of firms randomize their first-period prices in the interval $[\alpha, \overline{p})$ with a mass point at $\overline{p}$, while the second-period prices are in the interval $[\theta(\alpha), 1]$, with $\theta(\alpha) = \frac{-\alpha}{W_{-1}\left(-\frac{2\alpha}{e^2\overline{p}}\right)}$, where W_{-1} is the negative branch of the Lambert-W function.*[20]

The constrained type's equilibrium profit equals α, while the unconstrained type's profit equals $\alpha\left(1 - \left(W_{-1}(\frac{2\alpha}{e^2\overline{p}})\right)^{-1}\right)$. As W_{-1} is decreasing in $\overline{p}$, $\theta(\alpha)$ is also decreasing in $\overline{p}$ and these equilibria are Pareto-ranked and inferior, from the firms' perspective, to the pooling equilibrium.

Proof. The formal proof for this result can be found in the Appendix.

[20] Recall that the Lambert-W function is such that $W(xe^x) = x$.

The result follows from three basic observations. First, no equilibrium exists in which both types of firms choose prices on distinct supports, as (similar to a separating equilibrium) the unconstrained firm would have an incentive to deviate and mimic the constrained type. Second, as discussed in the beginning of this subsection lower prices must induce lower posteriors to keep firms indifferent across different first-period prices in the interval $[\alpha, \overline{p})$. Third, the mass point at the upper bound of the distribution must induce a posterior belief exactly equal to the price. This is because the firm must be indifferent between selling at the price $\overline{p}$ and losing the tiebreak to another firm selling at the same price.

Depending on $\overline{p}$, the semi-separating equilibria can reveal more or less information. The smaller $\overline{p}$, the less information is revealed, with the pooling equilibrium (which can be regarded as the limit of the semi-separating equilibria for $\overline{p}$ converging to α) revealing no information. Equilibria with higher $\overline{p}$ have a wider range of first-period prices and a more dispersed distribution of second-period beliefs. This higher dispersion is associated with a higher likelihood that the unconstrained firm sells in the first period, and, therefore, a lower average second-period price. That is, information revelation through market outcomes is detrimental to firms' profits since it necessarily involves a sort of miscoordination: unconstrained firms are more likely to sell early. In the proof in the Appendix we show that a greater $\overline{p}$ induces lower ex-ante profits, as the impact of a lower expected posterior in the second-period profit outweighs the (stochastic) increase in first-period prices brought about by a higher upper-bound.

We now establish some results that allow us to make comparisons across equilibria. For this reason, let us refer to E1 as the semi-separating equilibrium associated with the upper bound $\bar{p}_1$ and E2 as the equilibrium associated with the upper bound $\bar{p}_2$, with $\bar{p}_2 > \bar{p}_1$. The following comparisons can be made:

Proposition 1.5. *When compared with E1, E2 displays*

1. *a more informative first-period price distribution,*
2. *a higher probability that the unconstrained firm sells in the first period,*
3. *a lower expected second-period price.*

These results are illustrated in Figure 1.2 where we depict the ex-ante expected profit and the odds-ratio measuring the likelihood that the constrained firm sells relative to the prior in each equilibrium for $\alpha = 0.2$. The Figure shows that the odds ratio is quickly decreasing in $\overline{p}$. If $\overline{p} = \alpha$, the odds ratio is 1 as both types are equally likely to sell, but if $\overline{p} = 1$, it is approximately equal to only 0.33 so the unconstrained firm is disproportionately much more likely to sell. This results in firms assessing it is much more likely that there is competition in the second period, resulting in lower expected second-period prices and profits. The Figure also shows that as a result the ex-ante expected profit *over both periods* is decreasing from 0.56 to 0.47 when $\overline{p}$ ranges from α to 1, amounting to a 16% decrease.

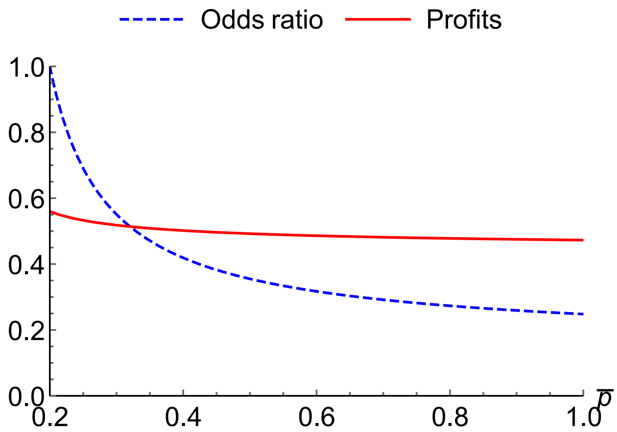

Figure 1.2: Ex ante expected profit (red, solid) and the relative likelihood of a constrained firm selling in the first period (blue, dashed) in different equilibria as a function of $\bar{p}$, calibrated for $\alpha = 0.2$.

1.4.2 Empirical and Welfare Implications

Now that we have characterized all equilibria of the game with private information, we discuss the implications. First, from the above, it is clear that Dudey (1992) overstates the market power of firms engaging in dynamic competitive pricing: the pooling equilibrium under private information yields already lower profits, and the semi-separating equilibria result in even lower profits. Thus, the ratio of ex-ante equilibrium profits under private information and complete information equals at most $\frac{\alpha(3-\alpha)}{2\alpha(2-\alpha)}$. For small values of α this ratio approaches 75%, shedding light on the incentives for information exchange regarding firms' capacities: firms can increase their margins significantly if they exchange information about capacities so that rivals' capacities are common knowledge. Because total sales are constant across equilibria, the ex-ante expected prices are lower while consumer surplus is higher under private information. Second, if we measure equilibrium price dispersion by the range of prices that may be charged along the equilibrium path, price dispersion increases over time. This holds in the pooling equilibrium (since there is no price dispersion in the first period) and in every semi-separating equilibrium. In a semi-separating equilibrium, the range of equilibrium prices in the first period is equal to $[\alpha, \overline{p}]$, and equal to $[\theta(p_1), 1]$ in the second period, where $\theta(p_1)$ is the posterior belief that the firm that sold in the first period is constrained given that it sold at p_1. Because in any semi-separating equilibrium $\overline{p} \leq 1$, the result on price dispersion follows if $\theta(p_1) \leq \alpha$ for any $p_1 \leq \overline{p}$. But as the expected pay-off of a constrained firm must be equal to α in any semi-separating equilibrium, and $\theta(p_1)$ is also the expected second-period profit, it must be that $\theta(p_1) \leq \alpha$ as otherwise a constrained firm will prefer setting $p_1 > \alpha$ as it could guarantee itself a profit larger than α in the second period. The same results hold for other measures of dispersion, such as variance, both in the pooling and semi-separating equilibria.

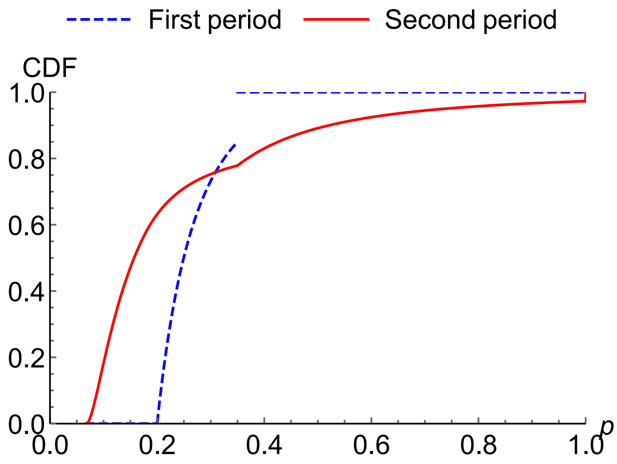

Figure 1.3: First- (blue, dashed) and second-period (red, solid) price distributions, calibrated for $\alpha = 0.2$ and $\overline{p} = 0.35$.

To illustrate the result for a semi-separating equilibrium, Figure 1.3 presents the first- and second-period price distributions for the case where $\alpha = 0.2$ and where we focus on an equilibrium with $\overline{p} = 0.35$.[21] The Figure clearly shows not only that the interval of prices that can be charged along the equilibrium path is larger in the second period than in the first, but also that first-period prices are more concentrated if other measures of price dispersion are used. In fact, one can also show that a property the figure displays, namely that the first- and second-period price distributions intersect only once, is always true so that the second-period price distributions have more mass both to the left and to the right of this intersection point. Thus, based on our model, an explanation for the observed increase in price dispersion towards the deadline is that in the first period(s) of competition firms try not to signal their capacity levels by choosing similar prices independent of their types. The pressure to sell (just) before the deadline, together with the remaining uncertainty regarding firms' types results in more price dispersion in later periods as suggested by the empirical evidence discussed in the Introduction.

Third, our model explains that flights that are expected to be peak do not have more dispersion than those that are expected to be off-peak (Puller et al. (2009)) as follows. Whether or not a flight is expected to peak is captured by the prior probability α a firm is constrained: the higher α, the more likely it is the flight is peak. If α is large, then there is no reason for firms to strongly compete and prices in both periods are close to the monopoly price with little price dispersion. Conversely, if α is small, then firms compete severely in both periods and all prices are close to marginal cost: even if the second-period distribution ranges from α to 1 almost all probability mass is close to α. Instead, when α is intermediate, the second-period price distribution displays high variance and the first-period price lies at the lower bound of this distribution. Thus, our model predicts that there is no clear monotonic relationship between price dispersion and peak and off-peak

[21] More precisely, for every $p_1 \in [\alpha, \overline{p}]$ there is a different price distribution that is characterised by $\theta(p_1)$. When we talk in this paragraph about the second-period price distribution, we mean the average second-period price distribution, where the average is taken over all possible $\theta(p_1)$'s.

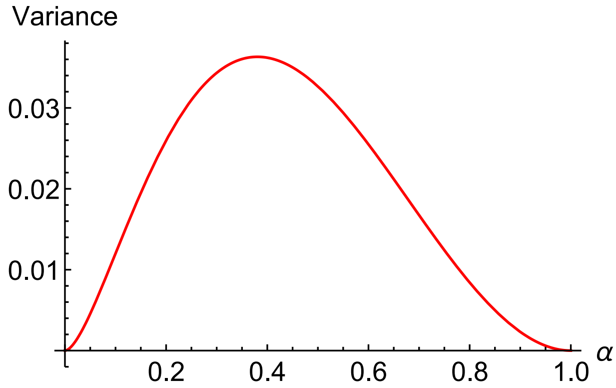

Figure 1.4: The variance of transaction prices in both periods in the pooling equilibrium for different levels of α.

flights. Figure 1.4 depicts the overall price dispersion across periods as a function of α in the pooling equilibrium.

Fourth, empirical evidence also suggests that transaction prices, *i.e.*, the prices at which consumers buy the good, often (but not always (cf., footnote 5)) tend to be higher towards the deadline. This is obviously true in the pooling equilibrium as the first-period price equals α, while second-period prices are distributed over the interval $[\alpha, 1]$. For semi-separating equilibria with higher profits, where $\overline{p}$ is close to α, the same holds true. However, other semi-separating equilibria with lower profits exhibit the opposite feature. Thus, our model can account for both increasing and decreasing prices towards the deadline. However, as firms' profits are inversely related to $\overline{p}$ one may argue that first-period prices are lower as firms are likely able to coordinate on the equilibria that are most profitable for them. Interestingly, there is an intimate connection between firms' profits and the evolution of transaction prices over time. A simple statistic summarizing the evolution of transaction prices is $\mathbf{E}[p_2 \mid p_1]$, *i.e.*, the expected second-period price given the first-period price. It turns out that the expected second-period price is $(2-\theta(p_1))\theta(p_1)$,[22] which depends on the posterior belief. In the pooling equilibrium, this posterior is equal to α and so the expected price is $(2-\alpha)\alpha > \alpha$. That is, prices follow a sub-martingale and we know that the pooling equilibrium is the best possible equilibrium from the firms' perspective. Instead, in the worst semi-separating equilibrium, where $\overline{p} = 1$, each price in $[\alpha, 1)$ is associated with a posterior $\theta(p_1) \leq \frac{1}{2}p_1$ and so $(2-\theta(p_1))\theta(p_1) < p_1$. In addition, since $\theta(1) = 1$, we have that $\mathbf{E}[p_2 \mid 1] = 1$. As a result, for all $p_1 \in [\alpha, 1]$, $\mathbf{E}[p_2 \mid p_1] \leq p_1$ so that prices follow a super-martingale in the worst possible equilibrium

[22] This can be seen as follows. First, with probability $\theta(p_1)$, only one firm has unsold capacity left in the second period and in that case the expected transaction price is simply the expected price given that the CDF of prices is given by (2), *i.e.*, $\theta(p_1)(1-\ln(\theta(p_1)))$, where the natural logarithm comes from integrating the term p^{-1} that arises in the distribution function. term With the remaining probability $1-\theta(p_1)$ there are two firms with unsold capacity in the second period. Using (1) and (2) the CDF of the minimum price in that case is given by $1 - \frac{\theta(p_1)(1-p_2)}{(1-\theta)p_2}\frac{\theta(p_1)}{p_2}$. The second-period transaction price follows by adding these different expressions.

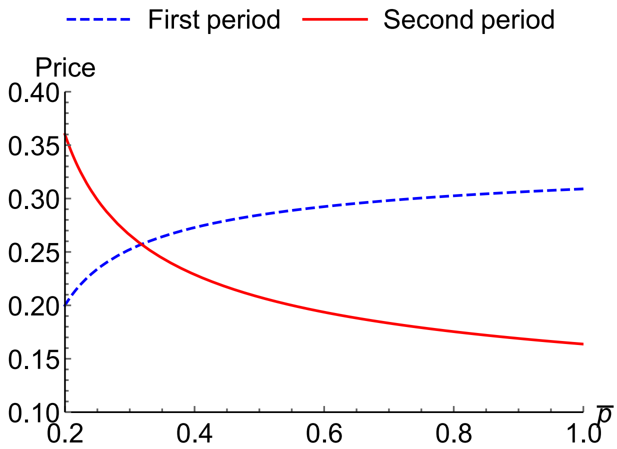

Figure 1.5: Ex ante expected first- (blue, dashed) and second-period (red, solid) transaction prices in different equilibria as a function of $\bar{p}$, calibrated for $\alpha = 0.2$.

from the firms' perspective.[23] This is illustrated in Figure 1.5, depicting expected first- and second-period transaction prices as a function of $\overline{p}$ for $\alpha = 0.2$. The figure shows that in the pooling equilibrium the expected second-period price is close to two times the expected first-period price, but that this ratio changes quickly when $\overline{p}$ increases and that when $\overline{p}$ is close to 1 expected second-period transaction prices can be in the order of four times smaller than first-period prices! Thus, semi-separating equilibria where $\overline{p}$ is close to 1 may explain some of the observations of last-minute discounts on Airbnb in Huang (2022).

Finally, there is a novel empirical implication of our model that, to the best of our knowledge, has not been tested yet, namely, that prices are positively correlated over time, even conditional on type, *i.e.,* higher first-period prices induce higher beliefs and, as a result, higher second-period prices. This implication follows from the fact that, for a given $\overline{p}$, $\theta(p_1)$ is increasing in p_1: higher prices are more likely to be set by constrained firms as relative to an unconstrained firm they do not benefit as much from selling in the first period as when they do, they cannot sell again in the second period. As $\theta(p_1)$ is the lower-bound of the distribution of second-period prices, a higher $\theta(p_1)$ means a first-order stochastic dominance shift of second-period prices.

1.4.3 Extensions

The model we have analyzed so far is stylized and abstracts from many features that are present in real-world markets and that have been studied before in the revenue management literature. The simplicity of the baseline model allows us to explain the main mechanism in detail and to argue that important real-world observations can already be explained by focusing on a new feature of markets with time-dated products, namely

[23]Across semi-separating equilibria, if the first-period price is competitive (below the mass point), the continuation price is expected to fall further. Instead, if the market price is not competitive (at the mass point) the continuation price is expected to raise even further.

asymmetric information about capacities. A legitimate concern, of course, is whether the explanation is robust to alternative, more realistic modeling assumptions. To this end, we now discuss several extensions to the baseline model. Depending on the extension, we show that the qualitative properties of the pooling and/or the semi-separating equilibria continue to hold in more complicated models that take important features of real-world markets into account. In addition, if the Dudey (1992) analysis extends without major complications, we compare the equilibrium profits under complete and private information to conclude that firms' profits are higher under complete information.

1.4.3.1 Random Arrival of Consumers

In the baseline model, it is common knowledge that a consumer arrives in every period. Whenever a firm fails to sell, it correctly infers that its rival sold, which allows it to keep track of the number of units sold. In this section, we relax this assumption so that, in every period, a consumer is present only with probability $\lambda \in (0, 1)$. We analyze both the model with complete information about capacities and our model with information asymmetry. We show that uncertainty about whether the competitor sold to a consumer further strengthens the importance of our analysis and our conclusion that private information reduces market power. In particular, firms' mixing strategies in the second period now also extend to the analysis where there is perfect information about capacities affecting Proposition 1.1. In addition, it remains true that the profits in the pooling equilibrium with private information about capacities are lower. In this subsection, we focus on the pooling equilibrium, while the analysis of semi-separating equilibria is in Appendix 1.C.3.

Consider first the (Dudey) model with complete information about capacities whereby in every period a consumer arrives with probability $\lambda \in (0, 1)$ and let us focus on the interesting case in which only one firm is constrained. If, as in Dudey's equilibrium, the constrained firm charges the lowest price in the first period, then there is a probability $1-\lambda$ that it does not sell and, therefore, still competes in the second period. This uncertainty creates an incentive to undercut in the second period that translates into an equilibrium price distribution with support in $[\lambda, 1]$. Given that the expected second-period profits equal λ^2, the constrained firm is pushed to sell at λ^2 in the first period. Interestingly, the same equilibrium applies if both firms know that they are constrained since a firm that did not sell in the first period does not know whether or not the rival sold. Only when both firms know that they are unconstrained, will the equilibrium remain the same as before with firms pricing at marginal cost. Overall, demand uncertainty drives the prices down, but the Dudey equilibrium structure is preserved, a result consistent with the findings of Martínez-de-Albéniz and Talluri (2011). We can compute the expected ex-ante profit by considering a firm that prices at λ^2 in both periods and it is given by $\lambda^2 \left[\alpha^2 + 2\alpha(1-\alpha)\right] = \lambda^2\alpha(2-\alpha)$.

Suppose then that firms are also uncertain about the capacity of the rival. Given that a firm did not sell in the first period, it updates the probability that it is a monopolist in the second period. Using Bayes' Rule this probability is $\frac{\lambda\alpha}{2-\lambda}$. Thus, the pooling equilibrium we constructed before for the case where $\lambda = 1$ can be easily extended by having both constrained and unconstrained firms choose $p_1 = \frac{\lambda\alpha}{2-\lambda}$ and randomize over prices larger than $\frac{\lambda\alpha}{2-\lambda}$ in the second period. Since all prices in the support yield the same profit, we can compute the expected ex-ante profit by assuming that the firm charges $\frac{\lambda\alpha}{2-\lambda}$ in every period. If it is constrained, the probability it sells its one unit is given by $\frac{\lambda}{2} + (1 - \frac{\lambda}{2})\lambda$, while if it is unconstrained the expected number of units it sells is given by $\frac{\lambda}{2} + \lambda$. Thus, the overall profit is given by $\frac{\lambda^2\alpha}{2-\lambda}\frac{3-\alpha\lambda}{2}$.

Ex-ante expected profits under information asymmetry about capacities are still smaller than those under complete information since $2 - \alpha > \frac{3-\alpha\lambda}{2(2-\lambda)}$. As the RHS is increasing in λ one could even argue that the effect of asymmetric information about capacities is even stronger if consumer arrivals are stochastic.

1.4.3.2 Multiple Periods

Another dimension to extend our model is to consider multiple periods. The two-period model we have considered so far is rather particular in that in the first period it is common knowledge that both firms are present while, being the last period, further dynamic considerations are irrelevant in the second period. With longer time horizons, the strategic considerations of firms include (*i*) uncertainty regarding the presence of their rival and (*ii*) the willingness to forgo current profits to improve future outcomes. To make progress towards a multi-period analysis, we need to separate these two strategic considerations. We do so by informing each firm whether their rival is still active at the beginning of each period. This is necessary to support an asymmetric pooling equilibrium, whereby firms alternate as the lowest-pricing firm.

Consider a T-period version of the model with a constant consumer arrival rate probability of $\lambda \in (0, 1)$ and allow firms to observe whether their rival is still active before setting the price. Constrained firms have a single unit, while unconstrained firms have at least T units. An asymmetric pooling equilibrium exists whereby, in every period in which both firms are active, the firm with the highest current posterior charges a price p_t (independent of its capacity), and its rival charges prices $p \in (p_t, p_t + \epsilon)$. Given this rule and some prior belief θ, the posterior belief conditional on the event that the firm who charged the lowest price remains active follows from Bayes' Rule and is given by

$$\theta' = \frac{\theta(1-\lambda)}{\theta(1-\lambda) + 1 - \theta}.$$

In such a pooling equilibrium, the constrained type of both firms must be indifferent in every period between selling at p_t or deferring to its rival. In the last period, the continuation profit is, by definition 0. In any previous period, let Π_t denote the continuation profit conditional on the rival firm remaining active and denote by V_t the continuation profit for a constrained type that becomes the only active firm in period

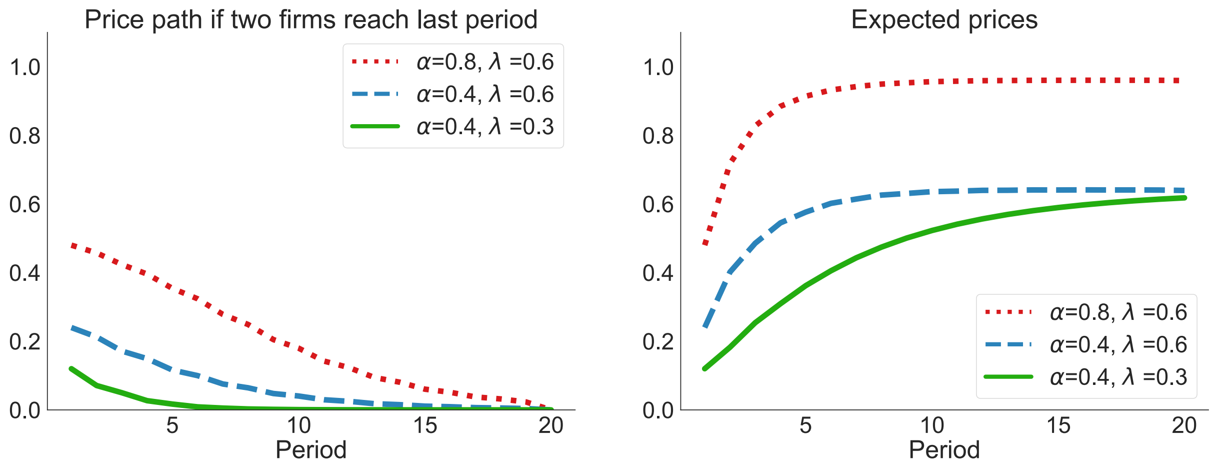

Figure 1.6: Evolution of prices for $T = 20$ periods for different values of the prior probability α and probability of consumer arrival λ.

t (i.e. $V_t = 1 - (1 - \lambda)^{T-t}$). In a pooling equilibrium, the constrained firm must be indifferent between being the lowest-pricing seller or not. If the firm is not, then it enjoys monopoly rents if its rival leaves and competitive rents if the rival stays. If the firm is the lowest-pricing seller and a customer arrives it sells at price p_t, while if no consumer arrives both firms remain active in the next period. Thus, the indifference condition reads: $\lambda\theta_t V_{t+1} + (1 - \lambda\theta_t)\Pi_{t+1} = \lambda p_t + (1 - \lambda)\Pi_{t+1}$. Iterating forward, we know that the same firm will be indifferent between selling or not next period. Hence,

$$\lambda\theta_t V_{t+1} + (1 - \lambda\theta_t)p_{t+1} = \lambda p_t + (1 - \lambda)p_{t+1},$$

which now characterizes the sequence of prices given the process of θ_t and $p_T = 0$. It can be shown that there are no incentives to deviate for either type of firm.

As pointed out by Martínez-de-Albéniz and Talluri (2011) and Gallego and Hu (2014), prices tend to decrease as long as firms fail to sell a unit. This is visible from the left plot of Figure 1.6, which depicts the resulting price sequence in case two firms reach the last time period. Still, the expected per-period price increases with time, as depicted by the right plot of Figure 1.6 for different values of λ and the initial belief α.

1.4.3.3 Multiple Firms

We now revert back to a two-period analysis, but extend our model to include $n \geq 3$ firms and have k consumers entering the market in each period. To keep the analysis close to the baseline case, we assume that firms can serve at most one consumer per period. Let $\tilde{n}$ denote the number of active firms in period 2 and $q(k)$ the probability the number of active firms is smaller than or equal to k. We focus on the existence of pooling equilibria for the case where $k \in [n/2, n]$.[24]

[24] If instead $k > n$, then $\tilde{n} \leq k$ with probability 1 and all active firms would sell at $p_2 = 1$, while if $2k < n$, then $\tilde{n} > k$ with probability 1 and firms would engage in Bertrand competition in the second period.

Given some outcome in the first period, the set of firms in period 2 can be partitioned into two groups: selling and non-selling firms. Non-selling firms are commonly known to be active in the second period while there is uncertainty about non-selling firms' participation. Let θ_i denote the (posterior) probability that firm i is active from the perspective of firm $j \neq i$. In the terminology of Armstrong and Vickers (2022), θ_i is the reach of firm i and because types are drawn independently, this market corresponds to the independent reach case (see also McAfee (1994)). Given the model features, in a pooling equilibrium k firms are (believed to be) inactive in the second period with probability α, while the remaining $n-k$ firms are active with probability 1. Denote this profile of posteriors by $\theta(k)$.

Similar to the analysis at the beginning of Section 1.4.1, firms with $\theta_i = 1$, must be willing to charge the monopoly price ($p_2 = 1$), at which firms sell, if and only if, $\tilde{n} \leq k$. This pins down the profit level of non-selling firms to be equal to $q(k)$, which conditional on being active is the same expected second-period profit all firms get.

Moving to the first period, we ask whether a pooling equilibrium exists. If it exists, it must be at the price $p_1 = q(k, \theta(k))$, where we make the dependence of q on $\theta(k)$ explicit. As before, the only reason why pooling may not be an equilibrium is that unconstrained firms prefer to undercut and ensure they sell in the first period. Their equilibrium profit is

$$\pi^u = \frac{k}{n} q(k, \theta(k)) + q(k, \theta(k)).$$

Their deviation profit is instead equal to $q(k, \theta(k)) + q(k, \theta(k-1))$, as unconstrained firms are the only types that possibly could have an incentive to undercut and, thus, will reveal themselves to be unconstrained if they deviate. Hence, the pooling equilibrium exists, if and only if,

$$\frac{k}{n} q(k, \theta(k)) \geq q(k, \theta(k-1)).$$

It is straightforward to verify that if k is the largest integer smaller than or equal to $(n+1)/2$, this condition is trivially satisfied (e.g. if $k = 1$ and $n = 2$ as in the baseline model). Instead, if $k = n - 1$, this condition is only met for α sufficiently low. In general, for every $k \leq n$, there exists some $\alpha(k, n)$ such that the condition is met for $\alpha < \alpha(k, n)$, and for these parameter values the pooling equilibrium is robust.

1.4.3.4 Demand and Consumer Heterogeneity

Next, we consider three extensions related to the modeling of demand. So far, we have studied a simple demand structure where demand is constant and certain over time, while consumers are myopic and demand the good in the period they arrive in the market. The revenue management literature has considered more sophisticated demand structures. For example, Gallego and Hu (2014) study heterogeneous consumer preferences and time-varying demand. Levin et al. (2009) consider forward-looking consumers, who can decide to wait and delay purchase. Martínez-de-Albéniz and Talluri (2011) and Anderson and

Schneider (2007) study dynamic pricing with stochastic demand. More recently, Dana Jr and Williams (2022) study dynamic capacity-then-price competition, when demand elasticity decreases over time. In this subsection, we incorporate some of these features in our model to show that the qualitative results remain robust.

Increasing willingness to pay. First, in our baseline model, we assume that consumers' valuation for the product does not change over time. The vast empirical literature on dynamic pricing in airline industries (see, e.g., McAfee and Te Velde (2006), Lazarev (2013), Williams (2022)) explains increasing prices towards the deadline with increasing willingness to pay of consumers. While demand becomes less price sensitive closer to the end of the selling period in the hotel industry as well (Garcia et al. (2022)), for the cruise industry Joo et al. (2020) show that consumers are more price-elastic closer to the end of the selling period, but prices increase due to the large amount of demand.

To accommodate the possibility of increasing willingness to pay, we now assume the second period consumer has a higher valuation $v > 1$, while we also comment on how our analysis below is robust to consumer heterogeneity.

It is not difficult to see that for a given posterior probability θ the second-period mixed pricing strategies are now given by

$$F^S(p_2) = 1 - \frac{\theta(v - p_2)}{(1-\theta)p_2} \quad \text{and} \quad F^N(p_2) = 1 - \frac{\theta v}{p_2},$$

which are defined over the interval $[\theta v, v]$.

Moving to the first period, it can be verified that for $\alpha v < 1$ the pooling equilibrium can be sustained with the pooling price being equal to $p = \alpha v$. If v is larger and $\alpha v > 1$ it is not surprising that a pooling equilibrium cannot be sustained anymore as firms have an incentive to concentrate almost exclusively on consumers that arrive late.

To construct the equilibrium for this case, consider that in the first period the constrained firm chooses a price of 1 with probability q and higher prices with the remaining probability. The unconstrained firm sets a price of 1 in the first period as it has no incentive to set a price above one. Thus, the posterior belief is given by $\theta = \frac{\alpha q}{1-\alpha+\alpha q}$. Hence, the expected profits of the constrained firm from charging the monopoly price and any higher price are given by:

$$\begin{aligned} &\pi^C(p_1 = 1) = \left[\frac{1}{2}(\alpha q + 1 - \alpha) + \alpha(1-q)\right] \times 1 + \left[\frac{1}{2}(\alpha q + 1 - \alpha)\right]\theta v \\ &\stackrel{!}{=} \pi^C(p_1 > 1) = (\alpha q + 1 - \alpha)\theta v. \end{aligned}$$

Solving for the mass point on the monopoly price gives $q = \frac{1+\alpha}{\alpha(v+1)}$, which for $\alpha v > 1$ yields $q < 1$ and $\theta v > 1$. The latter inequality also guarantees that the constrained firm does not want to deviate and undercut in the first period, while - like in our baseline

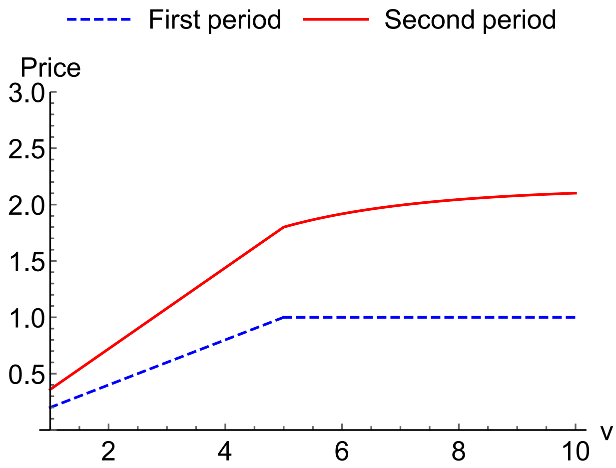

Figure 1.7: Expected first- (blue, dashed) and second-period (red, solid) price with increasing willingness to pay. Calibrated for $\alpha = 0.2$.

model - the unconstrained firm does not want to undercut because it triggers Bertrand competition in the second period.[25]

Figure 1.7 represents the expected (selling) prices in the first and second periods as a function of v. Not surprisingly, the difference in prices is increasing in v, while the Figure confirms that even for $v = 1$, the expected price is larger in the second period.

It is clear that if $\alpha v < 1$, the equilibrium profits equal $\alpha v(\alpha + \frac{3}{2}(1-\alpha))$, while under complete information about capacities, the ex-ante expected profits equal $\alpha^2 + \alpha(1-\alpha)(1+v)$.[26] Thus, for any α if v is not too large, the profits under complete information remain larger than under private information.

Loyal consumers. In our baseline model, all consumers consider the products offered by the firms as identical. One way to depart from this homogeneous goods market assumption is to consider that a fraction $\sigma \in (0, 1/2)$ of consumers is loyal (captive) to each of the firms and does not consider purchasing from the rival. A fraction $1 - 2\sigma$ of consumers considers both firms and purchases from the cheapest. We further assume that there is a single consumer in each period.

As is common in models with loyal consumers, only equilibria with mixed first-period prices exist, even if capacities are known. In this case, provided that $\sigma < 1/3$, the equilibrium is qualitatively similar to the case of no loyals studied in Section 3. The

[25]We can now consider a modification of this model by supposing that only a fraction q of consumers have a valuation $v > 1$ in period 2, while the rest have a valuation 1. If $qv < 1$, the monopoly price is 1, so the results from the baseline model are unaffected. If $\alpha qv < 1 < qv$, there exists a pooling equilibrium with $p_1 = \alpha qv$ and the only difference is that the second-period distribution has a hole in $(1, 1/q)$. Finally, if $\alpha qv > 1$, a semi-separating equilibrium exists where the constrained firm randomizes with prices above 1, and the price distribution in the second period may also have a hole in its support.

[26]When capacities are known, if one firm is known to be constrained and the other is unconstrained, the profits are 1 and v, respectively, while if both firms are constrained the equilibrium involves the firms choosing a first-period price of 1 with probability $\frac{1}{v}$ and a higher price with the remaining probability and to choose a second-period price equal to v. Interestingly, two constrained firms cannot benefit from the higher second-period valuation as they have to be indifferent between selling in the first and second periods.

unconstrained firm charges (stochastically) higher prices and, therefore, firms coordinate on a high margin equilibrium.

We focus instead on the model with unknown capacities. To see why a pooling equilibrium does not exist, note that in such an equilibrium the constrained firm must be indifferent between selling in the first period and selling in the second period. But with a fraction of loyal consumers, there cannot exist a mass point at any price below 1, as constrained firms would immediately deviate to 1 as they can sell to their loyal consumers and still get the same pay-off in the second period in case they did not sell. Semi-separating equilibria exist, however, and although more complex, the equilibrium characterization shares many features with the semi-separating equilibria described in Section 1.4.1. The reason for the additional complexity is that if σ is low enough, unconstrained firms are not willing to charge prices too close to 1 since these prices sell with low probability and, therefore, provide relatively low profits. Hence, in equilibrium, constrained firms are the only ones who charge prices close enough to 1. We refer the reader to Appendix 1.C.4 where we provide a detailed characterization.

Patient consumers. Finally, we have so far assumed that consumers are myopic and only consider buying when they enter the market. There has been considerable interest in the revenue management literature in the presence of forward-looking consumers (see, e.g. Levin et al. (2009), Besbes and Lobel (2015)): if prices are decreasing over time, consumers would rather wait before buying. As long as equilibrium prices are (stochastically) increasing over time, which is the case in the pooling equilibrium as well as in a range of semi-separating equilibria, price-taking consumers do not have an incentive to wait and buy later.

A more subtle point is raised in markets where there are a few large consumers that may internalize their impact on future prices when deciding to wait. In our setting, a consumer who waits induces a static game in the second period in which both firms compete for two consumers. Note that even in this case, competition between firms will not result in Bertrand competition as the constrained firms cannot serve the whole demand, and uncertainty about whether or not the competitor is constrained remains. What is an equilibrium of this second period is that a constrained firm charges a price of 1, while an unconstrained firm randomizes over the interval $(\alpha, 1)$. The expected price in the second period is then larger than α, implying that a large strategic consumer, facing a price of α in the first period of the pooling equilibrium does not want to postpone buying until the second period so that the assumption of myopic consumers is innocuous.[27]

Thus, where consumers are price takers or have an impact on future prices if they postpone their purchase decision, the pooling equilibrium is robust to this extension of patient, strategic buyers.

[27] Note also that this means that the dynamic market is more competitive than its static counterpart.

1.5 Disclosure

Having explored in detail the implications of private information, in this section and the next we explore to what extent private information continues to play an important role when firms can affect the information structure by either voluntarily disclosing information or by engaging in industrial espionage. Thus, in this section we allow firms to voluntarily disclose their private information about capacity and investigate whether such disclosure increases their profits.[28] As disclosure requires investing in technology to be able to do so,[29] disclosure is a long-term decision that is taken before firms learn their type. If a firm discloses the rival knows the capacity of the disclosing firm precisely.

It is clear that when the disclosure cost is too high, neither firm wants to disclose and the equilibrium analysis of the previous section remains valid. Therefore, in this section, we mainly focus on disclosure costs being small enough so that at least one firm will want to deviate from the equilibrium of the previous section and disclose. To this end, we start the analysis by assuming that one firm discloses its private information, while the other does not, and consider the impact disclosure has on the pricing strategies of firms. We then evaluate the overall profit the disclosing and non-disclosing firms make. We conclude that when the disclosing cost is small enough there only exist asymmetric equilibria where one firm discloses. Thus, despite the existence of voluntary disclosure, private information remains an important ingredient in understanding markets with time-dated products.

Consider first the situation where the disclosing firm is constrained. The continuation game is essentially identical to the complete information game in which one firm is constrained and the other is unconstrained. As we described earlier, equilibrium requires the disclosing firm to sell in the first period at a price p^* and its rival to sell in the second period at a price $p = 1$. Unlike the complete information game, equilibria where the first-period price is smaller than α do not exist as due to its ignorance about the rival's capacity, the disclosing firm has an expected profit of at least α. Similar to the complete information game, there is a continuum of equilibria in weakly dominated strategies and the only equilibrium in undominated strategies is the Dudey equilibrium with $p^* = 1$. It is on this equilibrium that we will focus.

The more interesting situation is where the disclosing firm is unconstrained. *If* it sells in the first period, it is common knowledge that both firms have capacity left in the second period and Bertrand competition results in marginal cost pricing. But this would leave the non-disclosing firm without profit in any period, which cannot be an equilibrium outcome. It follows that in equilibrium both types of the non-disclosing firm sell in the first

[28] In a separate literature, Cui et al. (2019) and Calvo et al. (2020) study empirically the effect of disclosing to consumers that capacity is running out. As discussed in the beginning of subsection 1.4.3 we study markets where consumers do not have to be afraid that capacity is running out so that disclosing capacity is mostly directed towards competitors.

[29] For example, in the airline industry it only makes sense to disclose information about unsold capacity if an airline does so for all its flights. This requires developing the software to track unsold capacity across several sales channels and display this information on the website.

period and set the same price. Using the second-period results of the pooling equilibrium discussed in the previous section reveals that both firms make an expected profit of α in the second period and the unconstrained non-disclosing firm can expect an overall profit of 2α. This outcome can be sustained with the following strategies. Following Blume (2003) the disclosing firm uniformly randomizes over the interval $(\alpha, \alpha + \varepsilon)$ in the first period, while both types of the non-disclosing firm charge α. In the second period, the disclosing firm has a posterior belief of α that its rival is constrained and both firms choose prices according to (1.1) and (1.2) as in the second period of the private information game.[30]

Combining the two cases, it is easy to see that the disclosing firm makes an expected ex ante profits of $\alpha p^* + (1-\alpha)\alpha - d$, where d is the disclosure cost:[31] with probability α it is constrained and makes a profit of $p^* \in [\alpha, 1]$ in the first period, whereas with the remaining probability it is unconstrained and makes an expected profit of α in the second period. On the other hand, the non-disclosing competitor makes an expected ex-ante profits of $\alpha \cdot 1 + (1-\alpha)(\alpha^2 + (1-\alpha) \cdot 2\alpha) = \alpha(3 - 3\alpha + \alpha^2)$: if the disclosing firm is constrained it makes a profit of 1 in the second period, while if the disclosing firm is unconstrained its overall profits are either α or 2α, depending on whether or not it is itself constrained.

Interestingly, the non-disclosing firm's profit exceeds the maximal pay-off it can get in the complete information game and, hence, it would never want to deviate. The reason is that it makes more profit if the rival is unconstrained. In that case, the disclosing firm lets the rival sell in the first period and also creates the opportunity for the non-disclosing firm to sell in the second period as the disclosing firm is uninformed about the rival's capacity. In addition, the non-disclosing firm earns more profits than the disclosing firm, regardless of the equilibrium that is being played. Therefore, equilibria exist in which only one firm discloses if the disclosure cost is sufficiently small ($d < \frac{\alpha(1-\alpha)}{2}$) so that the disclosing firm makes more profit than in the game under private information.

We summarize the above discussion in the Proposition below.

Proposition 1.6. *If $2d < \alpha(1-\alpha)$, then the disclosure game has a unique equilibrium in undominated strategies in which one firm discloses and obtains $\alpha(2-\alpha) - d$, while its rival obtains $\alpha(3 - 3\alpha + \alpha^2)$. If $d > \alpha(1-\alpha)(1-\zeta)$, where $\zeta \in (0, \frac{1}{2})$[32] all equilibria involve no disclosure and are outcome equivalent to the incomplete information game. For intermediate values, asymmetric disclosure and no disclosure equilibria co-exist.*

[30] This environment is strategically equivalent to a market with asymmetric priors. In such a case, the firm with a lower prior sells in the first period at a price equal to its prior.

[31] The disclosure cost is potentially multi-faceted. For instance, as disclosure is intimately related to transparency, it may impact consumer values.

[32] In the proof of Proposition 1.4 in the Appendix, we show that the lowest possible equilibrium pay-off of a semi-separating equilibrium under private information is given by $\alpha^2 + (1-\alpha)\alpha(1+\zeta)$, where $\zeta = -\left(W_{-1}\left(-\frac{2\alpha}{e^2}\right)\right)^{-1}$ and W_{-1} is the negative branch of the Lambert function. As the highest possible equilibrium pay-off of a disclosing firm equals $\alpha + (1-\alpha)\alpha - d$ there cannot exist a disclosure equilibrium if $\alpha + (1-\alpha)\alpha - d < \alpha^2 + (1-\alpha)\alpha(1+\zeta)$, which gives the condition stated in the Proposition.

The operating profits of disclosing, given by $\alpha(2-\alpha)$, are strictly larger than the maximal pay-off in the private information game analyzed in the previous section. Thus, for a small enough disclosure cost, it is optimal for one firm to disclose. The multiplicity of equilibria arises for intermediate disclosure cost and stems from the multiplicity of equilibria studied in the previous section. As the private information game has multiple semi-separating equilibria with different pay-offs, the decision whether or not to disclose depends on a firm's expectations regarding which equilibrium they will play if it does not disclose.

Notice that all of these equilibria are less competitive than the incomplete information benchmark. More importantly, the disclosure equilibrium induces higher prices than in the complete information outcome! Therefore, allowing *voluntary* information disclosure in this market is anti-competitive, and even more so than mutual information exchange, and accordingly severely harms consumers. From a competition policy perspective, our analysis indicates that there may be good reasons to forbid firms to reveal their capacities in a dynamic pricing setting where capacity constraints are important. The disclosing firm mainly gains in all cases where it is capacity constrained, while the additional loss of not having private information when it is unconstrained is insufficient to dominate this gain. The non-disclosing firm is mainly free-riding on the disclosure decision of its rival and benefits especially if it is unconstrained.

1.6 Industrial Espionage

We now consider the incentives of firms to spy on the rival and to secretly learn its capacity. Industrial espionage is, in a sense, the reverse of voluntary disclosure as the one who incurs the spying cost c is the one who learns the capacity of the rival. If there exist equilibria with industrial espionage, then the firm expecting to be spied upon also expects the rival to know its capacity. As in the previous section, we model espionage as a long-term endeavor that is decided upon before pricing decisions are made and it delivers certainty about whether or not the rival is constrained. Importantly, and unlike the previous section, the act of spying cannot be observed so it is the expectation of being spied upon rather than the act itself that affects a firm's behavior.

It is not difficult to see that for any $c > 0$, there cannot be an equilibrium where both firms spy for sure. If they would do so, the equilibrium outcome of the complete information game would result in the firms setting prices equal to marginal cost if and only if both turn out to be unconstrained. But then one of the firms can obtain the same operating profits without incurring the spying cost, by setting a first-period price of 1 if constrained and a larger price if unconstrained. If the rival is constrained, a firm can then anyway obtain a profit of 1 (possibly in the second period), while if the rival is unconstrained it also obtains a profit of 1 if constrained. This result should not come as a surprise: there is something secretive about spying, and one may think that a firm may not want to act in such a way that the rival expects it to engage in industrial espionage for sure. As an implication, some form of private information will prevail also when allowing for industrial espionage.

Similar to the previous section, when the spying cost is large an equilibrium without industrial espionage exists, whereas if the spying cost is smaller, but not too small, an asymmetric equilibrium exists where one firm spies and the other does not. We provide the details of the threshold values where these equilibria exist in Appendix 1.B. Equilibrium play in this equilibrium mimics that of the asymmetric equilibrium in the disclosure game - one firm is informed about her rival's type and the other firm expects this to be the case. It follows from Section 1.5 that the ex-ante expected candidate equilibrium profits are equal to $\alpha(3-3\alpha+\alpha^2)-c$ and $\alpha(2-\alpha)$ for the spying (informed) and non-spying (uninformed) firm, respectively. Since the non-spying firm has a lower equilibrium pay-off, it may be tempted to secretly spy on its rival and use this information in the second period to increase its profit. This puts a lower bound on the spying cost for this pure strategy asymmetric equilibrium to exist.[33]

Important new considerations apply when the spying cost is small and we now turn our attention to constructing an equilibrium in this range where one firm spies for sure, while the other firm spies with some probability $0<\beta<1$. We refer to the latter firm as the mixing firm. If this mixing firm is constrained, this event is common knowledge and this firm will sell at a price not lower than the prior.

When the mixing firm is unconstrained, instead, equilibrium play dictates that the rival (spying) firm sells in the first period. Continuation play will then depend on the capacity of the rival and the information gathered by the mixing firm. What is of special interest in the second period is that if the mixing firm spied, it charges the monopoly price if its rival is constrained and charges lower, but still strictly positive, prices otherwise. If both firms spy and are unconstrained, both firms know that there is enough unsold capacity in the market to result in severe competition, *i.e.*, marginal cost pricing in both periods. However, the firm that is supposed to always spy does not know that the mixing firm (spied and) knows that it is unconstrained. As both the constrained and unconstrained spying firm sell at the same first-period price (and learn nothing about the rival), an unconstrained spying firm believes with probability $1-\beta$ that the mixing rival still has its posterior of it being constrained equal to the prior α, and thus is setting positive prices, therefore the spying unconstrained firm reacts correspondingly by charging positive prices as well. *Thus, even if both firms are unconstrained and know that the rival is unconstrained, they may escape marginal cost pricing as it is not common knowledge that both are unconstrained.*

In Appendix 1.A we show the following proposition holds, which essentially says that the candidate equilibrium we outlined above is indeed an equilibrium:

[33] Note that the non-disclosing (informed) firm in the previous section, would never want to deviate as it makes the most profit, and by disclosing she does not get more information itself. As the roles are reversed, this does not hold for the non-spying (uninformed) firm.

Proposition 1.7. *If $0 < c < \alpha(1-\alpha)^2$, then there exists an asymmetric equilibrium where one firm engages in industrial espionage for sure, while the other randomizes between spying and not spying. The randomizing firm has the same ex-ante expected profit as under complete information, while the spying firm's profits are smaller. As the cost of spying approaches 0, the spying probability of the mixing firm approaches 1, but the ex-ante expected equilibrium profit of the firm that always engages in industrial espionage converges to α, which is smaller than that under the Dudey outcome.*

A reader familiar with Dudey (1992) may have expected that if the cost of espionage approaches 0 and both firms spy almost surely, equilibria converge to the Dudey monopoly result. The Proposition shows this is not necessarily the case. However, the asymmetric equilibrium we constructed here does converge to one of the equilibria of the complete information game (see Section 1.3). In this equilibrium the firm that always engages in industrial espionage has an expected profit of $\alpha + \frac{c}{1-\alpha}$, which is increasing in the cost of spying. The reason is that the rival firm's probability of spying decreases and because of this the spying firm makes a higher expected profit in the second period, which also translates into a higher first-period price in case the rival is unconstrained.

In this equilibrium, if c approaches 0 the ex ante expected pay-off of the spying firm converges to α, which is considerably below the pay-off of $\alpha + (1-\alpha)\alpha$ of the mixing firm, which is also the ex-ante expected pay-off of the complete information game. Calculating the average pay-off of a firm in this equilibrium as $\alpha + \frac{1}{2}(1-\alpha)\alpha$, it is clear that this is exactly equal to the firms' ex ante expected profit in the pooling equilibrium under private information, $\alpha\frac{3-\alpha}{2}$. As for small c equilibrium profits are increasing in c, firms prefer the spying equilibrium to the one under pure private information, while consumers are worse off.

An open question remains as to whether other equilibria supporting different outcomes exist. In particular, a symmetric equilibrium may exist whereby both firms choose to spy with some positive probability. In such an equilibrium, firms will mix in overlapping intervals depending on both their type and their information. As the cost of spying converges to zero, it may well be that these price distributions collapse to the prices in the Dudey outcome.

1.7 Conclusion

This is the first paper that performs an equilibrium analysis of dynamic competitive pricing in markets for time-dated products (such as hotel rooms, airline flights, shipping, and generated electricity), where firms have private information about their unsold capacities. We focus mostly on a simple framework with two firms, two periods, and one consumer arriving in each period, but we consider several extensions such as multiple periods, consumers arriving randomly, and having higher valuations closer to the deadline. Despite its simplicity, the equilibrium analysis is intricate and the results yield interesting insight into existing empirical observations. We show that the existence of private information considerably restricts the market power that firms have in such situations and that information exchange of private information regarding unsold capacity is anti-competitive and detrimental to consumer welfare. We also show that the model can explain observed pricing patterns. In particular, our model provides an explanation for increasing prices and price dispersion as the deadline approaches. We extend the analysis by considering the private incentives to voluntarily disclose private information or to engage in industrial espionage. Surprisingly, from a consumer welfare perspective, one-sided voluntary disclosure is even worse than mutual information exchange as the disclosing firm is able to get the full information pay-offs, but it is especially the non-disclosing firm that benefits from the additional information regarding the rival's capacity while maintaining uncertainty regarding its own capacity. This raises the question of whether voluntary disclosure of unsold capacity should be forbidden from a competition policy perspective. Industrial espionage is in some sense similar to "reverse voluntary disclosure", with similar results, however, the difference being that it is a secretive act that creates private rather than public information. We summarize the profit and consumer welfare comparisons across the different information scenarios we have analyzed in the following Table 1.1.

From a policy perspective, our work is also related to the recent literature studying potential channels in which (pricing) algorithms may impact competition (Harrington (2018) and Calvano et al. (2020)). As indicated above, one part of our analysis shows that it is not only essential for firms to coordinate their pricing that they are informed about their rivals' capacity but that it is known by their rivals that they know. When it is commonly known that competitors use similar algorithmic tools, one possible implication of our research is that the use of these tools may lead to a substantial increase in prices and a reduction in consumer surplus.

Our analysis points to many angles for future research. One possibility is to cover different extensions in one model. Another is to investigate our conjecture at the end of Section 1.6 and to inquire into the existence of a symmetric equilibrium where both firms randomize their decision to engage in industrial espionage, and to study whether the equilibrium converges to the Dudey equilibrium outcome as the cost of industrial espionage becomes negligible. Further direction would be to allow for the possibility that

Equilibrium type	Public Information (Dudey)	Private Information (Pooling)	Disclosure (Small cost)	Espionage (Small cost)
Industry profits	3	1	4	2
Consumer surplus	2	4	1	3

Table 1.1: Ranking of industry profits and consumer surplus across the main four equilibria with 4 denoting the best and 1 the worst outcome.

the product is sold out. With consumers being strategic and firms potentially selling out, it is interesting to see whether disclosure policies of unsold capacities may also influence consumer decisions. Third, In our basic model and in the different extensions we have assumed that consumers know that they will be able to buy the product if they accept the price, i.e., capacity is not a limiting factor. When consumers are myopic and do not consider waiting for future periods to buy this is not an issue as they either face firms selling their product or they do not. However, when consumers are strategic, as in one of the extensions analyzed, consumers may decide to buy now as they are afraid the product will be sold out later. Finally, we have assumed that demand is known and it would be interesting to introduce elements of uncertainty regarding demand and demand learning into our analysis (see, e.g., Bertsimas and Perakis (2006), Liu and Zhang (2013), Gallego and Hu (2014) and Cheung et al. (2017)).

1.A Proofs of Propositions

Proof of Proposition 1.2. Consider a separating equilibrium, possibly in mixed strategies, where p^u is a price in the support of the first-period prices set by an unconstrained firm, while p^c is in the support of the first-period prices set by a constrained firm. To be an equilibrium, no type should have an incentive to imitate the price of the other type. Note that both deviations we consider are on-the-equilibrium path, such that out-of-equilibrium beliefs do not play a role. The two deviations give the following two incentive compatibility constraints:

$$\begin{aligned}\pi^u(p^u) &= \left[\alpha(1-F_C(p^u)) + (1-\alpha)(1-F_U(p^u))\right] p^u + \alpha F_C(p^u)\\ &\geq \left[\alpha(1-F_C(p^c)) + (1-\alpha)(1-F_U(p^c))\right] p^c + 1 - (1-\alpha)F_U(p^c) = \pi^u(p^c)\end{aligned}$$

and

$$\begin{aligned}\pi^c(p^c) &= \left[\alpha(1-F_C(p^c)) + (1-\alpha)(1-F_U(p^c))\right] p^c + \alpha F_C(p^c)\\ &\geq \left[\alpha(1-F_C(p^u)) + (1-\alpha)(1-F_U(p^u))\right] p^u + \alpha F_C(p^u) = \pi^c(p^u),\end{aligned}$$

where π^u and π^c represent the profit of an unconstrained, respectively a constrained firm, while $F_U(p)$ and $F_C(p)$ represent the probability an unconstrained, respectively a constrained firm sets prices smaller than p in the first period. Note that for the equilibrium to be separating, these distributions have different supports. The imitation pay-offs $\pi^u(p^c)$ and $\pi^c(p^u)$ consists of two terms. The first term reflects the profits the unconstrained, respectively the constrained, firm makes in the first period, taking into account the probabilities the competitor is (un)constrained and the prices they may set. For the second-period profit of $\pi^u(p^c)$, we notice that by imitating the first-period price of a constrained firm, the competitor believes at the beginning of the second period that the firm is constrained. Therefore, the firm in question will be able to sell at a price just below 1 in all cases, apart from when the competitor is unconstrained and sold in period 1, which happens with probability $(1-\alpha)F_U(p^c)$. In that case, both firms believe that both firms still have unsold capacity and engage in Bertrand competition. Similarly, for the second-period profit of $\pi^c(p^u)$, the competitor believes at the beginning of the second period that the firm is unconstrained and the firm in question will therefore only be able to sell at a price of 1 when the competitor is constrained and sold in period 1, which happens with probability $\alpha F_C(p^u)$.

These two constraints can be rewritten as

$$\begin{aligned}&1-(1-\alpha)F_U(p^c) - \alpha F_C(p^u)\\ &\leq \left[\alpha(1-F_C(p^u)) + (1-\alpha)(1-F_U(p^u))\right] p^u - \left[\alpha(1-F_C(p^c)) + (1-\alpha)(1-F_U(p^c))\right] p^c\\ &\leq \alpha(F_C(p^c) - F_C(p^u)),\end{aligned}$$

but this requires that $1 \leq \alpha F_C(p^c) + (1-\alpha)F_U(p^c)$. This is only possible if $F_C(p^c) = F_U(p^c) = 1$ for all prices p^c in the support of the first-period prices set by a constrained firm. This would imply that the constrained firm chooses p^c with probability 1 and that $p^c = 1$. In such a candidate equilibrium unconstrained firms randomize over the interval $[\alpha, 1)$ as the two types can definitively not pool on a price of 1 as firms would like then to undercut the price, while they should randomize in period 1 as they do not know whether the other firm is constrained or not. The candidate equilibrium results in both types making an expected profit of α, but the unconstrained would then be better off deviating to 1, making an expected profit of $3\alpha/2$.

Proof of Proposition 1.4. We show this by construction. Consider any semi-separating equilibrium with a support of first-period prices on an interval $(\underline{p}, \bar{p})$. Clearly, $\underline{p} \geq \alpha$ as prices below α are weakly dominated for constrained firms. Similarly, $\underline{p} \leq \alpha$, as otherwise the constrained firm will never price at $\bar{p}$, which yields an expected profit of α. Hence, $\underline{p} = \alpha$. Naturally $\bar{p} \leq 1$.

For every price $p \in [\alpha, \bar{p}]$, the (posterior) belief $\theta(p)$ can be derived from the equilibrium strategies of each firm. For prices $p < \alpha$, simply let $\theta(p) = 0$ while for prices $p > \bar{p}$ we let $\theta(p) = 1$. We now use the randomization conditions for both types of firms. Given that $Q(p) = \Pr(\tilde{p} > p)$ is the probability that the rival's price exceeds p, for every price $p \in [\alpha, \bar{p})$, we have:

$$\Pi^c(p) = Q(p)p + \int_\alpha^p -Q'(\tilde{p})\theta(\tilde{p})d\tilde{p}.$$

In words, a constrained firm charging p sells with probability $Q(p)$. If it fails to sell (because the competitor's price $\tilde{p}$ is smaller than p) it obtains a continuation profit equal to the posterior type of the selling firm, which depends on the price the competitor charges ($\theta(\tilde{p})$). Taking derivatives we obtain:

$$0 = Q'(p)p + Q(p) - Q'(p)\theta(p) = Q'(p)(p - \theta(p)) + Q(p).$$

The constrained firm optimally trades off selling its unit in the first period, which brings p, but results in a loss of continuation profit of $\theta(p)$, with obtaining more revenue on the infra-marginal units.

The unconstrained firm's profit reads:

$$\Pi^u(p) = Q(p)(p + \theta(p)) + \int_\alpha^p -Q'(\tilde{p})\theta(\tilde{p})d\tilde{p} = \Pi^c(p) + Q(p)\theta(p).$$

That is, the unconstrained firm's profit reflects the fact that if it sells in the first period it can still sell in the second and obtain a continuation profit that depends on the price it charges (as it becomes the selling firm). Since Π^c is constant by construction,

$$Q'(p)\theta(p) + Q(p)\theta'(p) = 0.$$

Replacing this condition in the randomization condition of the constrained firm yields:

$$\frac{\theta'(p)}{\theta(p)}(p-\theta(p))-1=0,$$

an ordinary differential equation (ODE) in $(\theta(p),p)$ on $[\alpha,\bar{p})$. This differential equation admits the solution:

$$\theta(p)=-\frac{p}{W(-cp)},$$

for some constant c, where W is the Lambert-W function.

A crucial observation is that the unconstrained firm must have a mass point at the upper-bound $\bar{p}$ of the support, for otherwise $Q(\bar{p})=0$ and $Q(\bar{p})\theta(\bar{p})=0$ which is a contradiction with the fact that the unconstrained firm is indeed willing to charge such a price. The indifference conditions for the mass point now read:

$$\lim_{p\to\bar{p}}Q(p)p+\int_{\alpha}^{\bar{p}}-Q'(\tilde{p})\theta(\tilde{p})d\tilde{p}=\frac{1}{2}(Q(\bar{p})\bar{p}+Q(\bar{p})\theta(\bar{p}))+\int_{\alpha}^{\bar{p}}-Q'(\tilde{p})\theta(\tilde{p})d\tilde{p},$$

for the constrained firm and

$$\lim_{p\to\bar{p}}Q(p)\theta(p)=\frac{1}{2}Q(\bar{p})\theta(\bar{p}),$$

for the unconstrained firm. The first condition requires that $\theta(\bar{p})=\bar{p}$. In other words, a constrained firm must be indifferent between selling at price $\bar{p}$ and losing a tiebreak to another firm selling at the same price. The second condition requires that $\lim_{p\to\bar{p}}\theta(p)=\frac{1}{2}\theta(\bar{p})$. Notice that $\theta(p)$ is discontinuous at $p=\bar{p}$ since the quantity also jumps discretely at the mass point as a result of the tie-breaking rule. Using these conditions we can solve for c above.

$$\theta(p)=\frac{p}{-W_{-1}\left(-\frac{2p}{e^2\bar{p}}\right)},$$

where $W_{-1}(x)$ is the negative branch of the W-Lambert function (which is only defined for $p<\bar{p}$).

The expected profits can then be simply obtained by noticing that $Q(\alpha)=1$. Hence,

$$\Pi^u(\alpha)=\alpha+Q(\alpha)\theta(\alpha)=\alpha+\frac{\alpha}{-W_{-1}\left(-\frac{2\alpha}{e^2\bar{p}}\right)}.$$

This expression is decreasing in $\bar{p}$ with $\alpha+\frac{\alpha}{-W_{-1}\left(-\frac{2\alpha}{2e^2}\right)}\leq\Pi^u\leq 3\alpha/2$.

To verify that this constitutes an equilibrium we need to impose that the expected posterior equals the prior. Namely,

$$\int_{\alpha}^{\bar{p}}-Q'(p)\theta(p)dp+Q(\bar{p})\bar{p}=\alpha.$$

That this condition holds can be seen as follows.

$$\int_{\alpha}^{\bar{p}} -Q'(p)\theta(p)dp = \int_{\alpha}^{\bar{p}} -Q'(p)\frac{\theta(\alpha)}{Q(p)}dp \tag{A.1}$$

$$= \theta(\alpha)(\ln(Q(\alpha)) - \ln(Q(\bar{p}))) = -\theta(\alpha)\ln(Q(\bar{p})). \tag{A.2}$$

The first step uses the randomization condition of the unconstrained firms, the second condition follows from integration and the third step uses the fact that $Q(\alpha) = 1$. As the boundary conditions of the ODE yield $Q(\bar{p}) = \frac{2}{\bar{p}}\theta(\alpha)$, it follows that

$$\int_{\alpha}^{\bar{p}} -Q'(p)\theta(p)dp + Q(\bar{p})\bar{p} = -\theta(\alpha)\ln\left(\frac{2\theta(\alpha)}{\bar{p}}\right) + \frac{2\theta(\alpha)}{\bar{p}}\bar{p} \tag{A.3}$$

$$= -\theta(\alpha)\ln\left(\frac{2\theta(\alpha)}{\bar{p}}\right) + 2\theta(\alpha). \tag{A.4}$$

Recall that $\theta(\alpha) = \frac{-\alpha}{W_{-1}\left(-\frac{2\alpha}{e^2\bar{p}}\right)}$. A useful property of the Lambert-function is that $\ln\big(W(z)\big) = \ln(z) - W(z)$.[34] Hence,

$$\ln\left(\frac{2\theta(\alpha)}{\bar{p}}\right) = \ln\left(\frac{2\alpha}{\bar{p}}\right) - \ln\left(\frac{2\alpha}{e^2\bar{p}}\right) + W_{-1}\left(-\frac{2\alpha}{e^2\bar{p}}\right) \tag{A.5}$$

$$= 2 + W_{-1}\left(-\frac{2\alpha}{e^2\bar{p}}\right). \tag{A.6}$$

Using this, we can further simplify

$$\int_{\alpha}^{\bar{p}} -Q'(p)\theta(p)dp + Q(\bar{p})\bar{p} = -\theta(\alpha)W_{-1}\left(-\frac{2\alpha}{e^2\bar{p}}\right) = \alpha. \tag{A.7}$$

Proof of Proposition 1.5. We first establish (i). By definition, the expected posterior about the type of each firm is the prior in both E1 and E2. The distribution of posteriors in equilibrium i is

$$F_i(\theta) = 1 + \frac{\alpha}{\theta W_{-1}(-\frac{2\alpha}{e^2\bar{p}_i})},$$

for $\theta < \bar{p}_i/2$, $F_i(\theta) = F_i(\bar{p}_i/2)$ for $p \in (\bar{p}_i/2, \bar{p}_i)$ and $F_i(\theta) = 1$ otherwise, where $F_i(\bar{p}_i/2)$ is simply given by substituting $\bar{p}_i/2$ for θ in the above expression. It is easy to see that $F_1(\theta) < F_2(\theta)$ for all $\theta < \bar{p}_1$ and $F_1(\theta) \geq F_2(\theta)$ for all $\theta \geq \bar{p}_1$. This establishes that the distribution of posteriors in E2 is more dispersed than under E1 and since they have the same mean, F_2 is a mean-preserving spread of F_1. This trivially guarantees that $F2$ is more informative under the disperse order than $F1$.

[34] To see this, notice that for $z_1 > 0$ that $w_1 e^{w_1} = z_1$ so that $W(z_1) = w_1$ and $w_2 e^{w_2} = -z_1$ so that $W(-z_1) = w_2$. We then have that $w_1 e^{w_1} = -w_2 e^{w_2}$. Taking logs on both sides we get $w_1 + \ln w_1 = w_2 + \ln(-w_2)$, or $W(z_1) + \ln W(z_1) = W(-z_1) + \ln(-W(-z_1))$. As $W(z_1) + \ln W(z_1) = \ln z_1$ we have that $\ln(-W(-z_1)) = \ln z_1 - W(-z_1)$.

We now show (ii) and (iii). The expected second-period profit is $\Pi(\theta) = \theta + (1-\theta)\theta = (2-\theta)\theta$, which is increasing and concave in θ. Let θ_1 and θ_2 be the random variables associated with a given firm's posterior in both equilibria. By Rothschild and Stiglitz (1976), $\theta_2 = \theta_1 + \varepsilon$ for some ε with zero mean. Let $G_i(\theta)$ denote the distribution of the minimum of two draws from $F_i(\theta)$. Notice then that,

$$\begin{aligned}\int \Pi(\theta_2)dG_2(\theta_2) &\leq \int\int \Pi(\theta_1+\varepsilon)dG_1(\theta_1)dH(\varepsilon)\\ &\leq \int \Pi(\theta_1)dG_1(\theta_1).\end{aligned}$$

The first step follows from the fact that the minimum of θ_2 is not higher than the minimum of θ_1 plus a random draw of ε, establishing (ii); while the last step follows from Jensen's inequality, proving (iii).

Proof of Proposition 1.7. We start the analysis in the second period of the price competition and consider the case where the spying firm has successfully sold a unit in the first period and believes that the rival is informed with probability β and uninformed with probability $(1-\beta)$.[35] Applying the insights gained in Section 1.4, the spying firm expects that the mixing firm, when it does not spy and is uninformed, randomizes over some interval $[\hat{p}, 1]$ with some cumulative distribution function $G^{MU}(p)$ which has a mass point at the upper bound. Let us denote the mass by ω. As a best response, the spying firm must also randomize on the interval $[\hat{p}, 1)$. However, an informed mixing firm (that has spied and knows the rival's type) chooses different prices, depending on the rival's type. If the rival is constrained, then the mixing firm naturally chooses the monopoly price as the spying firm has sold in the first period. If, however, the spying firm is unconstrained, then the mixing firm knows that and will undercut this interval in order to maximize profits. As there cannot be an equilibrium with the informed mixing firm choosing a second-period price for sure (as the spying firm will have an incentive to undercut), the spying firm should not only mix over $[\hat{p}, 1]$ but also over a lower interval. Thus, the mixing strategy of the spying firm will be represented by some cumulative distribution function $G^S(p)$ over the interval $[p_L, 1]$, for some $p_L < \hat{p}$. In addition, when informed the mixing firm will randomize over the interval $[p_L, \hat{p}]$ with $G^{MI}(p)$.

In what follows, we derive the unknowns $(\omega, \hat{p}, p_L)$, and some properties of the functions $G^S(p)$, $G^{MI}(p)$ and $G^{MU}(p)$. First, by undercutting the mass point of ω at 1, the spying firm has an expected profit of $(1-\beta)\omega$. Thus, $G^{MI}(p)$ and $G^{MU}(p)$ are characterized such that any price in the interval $[p_L, 1]$ yields the spying firm the same profit. Second, as by pricing at p_L firms know for sure they will sell, it must be that $p_L = (1-\beta)\omega$. Third, as by setting $\hat{p}$ the spying firm makes a second-period profit of $(1-\beta)\hat{p}$ it should be that

[35]The latter implicitly assumes that the first-period pricing strategy of the mixing firm is such that no information regarding the mixing firm's type is revealed. However, as the mixing firm is not supposed to sell in the first period anyway, the pricing strategy in this period is pay-off irrelevant and we may as well assume that both types choose the same first-period pricing strategy.

$\hat{p} = \omega$. Fourth, by not spying and setting a price equal to 1 an unconstrained mixing firm expects to get a second-period profit of α. As the non-spying firm expects to get the same profit over the whole domain $[\hat{p}, 1)$, it should be that $(\alpha + (1-\alpha)(1-G^S(\hat{p})))\hat{p} = \alpha$. Fifth, as by spying and setting a price equal to p_L in case the rival is unconstrained and choosing a price of 1 if the rival is constrained, the mixing firm gets an ex-ante expected pay-off of $\alpha + (1-\alpha)\alpha + (1-\alpha)^2(1-\beta)\omega - c$ when it actually spies, while the ex ante expected pay-off of not spying equals $\alpha + (1-\alpha)\alpha$ it should be that $(1-\alpha)^2(1-\beta)\omega - c = 0$. Sixth, as when spying the mixing firm should be indifferent between setting any price in the interval $[p_L, \hat{p}]$, it should be that $(1-\beta)\omega = (1-G^S(\hat{p}))\hat{p} = (1-G^S(\hat{p}))\omega$ so that $G(\hat{p}) = \beta$.

Combining the last three points, yields expressions for β and ω in terms of exogenous parameters. It follows that for any $c > 0$, $\omega > \alpha > 0$ and from the fifth point that β converges to 1 and ω converges to 1 if c approaches 0. This implies that in this asymmetric equilibrium the expected second-period profit of the spying firm converges to 0. In terms of first-period pricing strategies, as the mixing firm should not be willing to deviate from its strategy of pushing the rival to sell when it spies, the first-period price set by the spying firm if the rival is unconstrained should be equal to $(1-\beta)\omega$ and the interval of prices where the mixing firm uniformly randomizes over is $[(1-\beta)\omega, (1-\beta)\omega + \varepsilon]$. It follows that these prices also converge to 0 if c approaches 0 and that in the limit the spying firm will only make profit if the mixing firm is constrained. Thus, if c approaches 0 the ex-ante expected pay-off of the spying firm equals α, which is considerably below the pay-off of $\alpha + (1-\alpha)\alpha$, which is the pay-off of the mixing firm and also the ex-ante expected pay-off of the complete information game. Interestingly, it is also lower than the profit under private information.

1.B Equilibria when the Cost of Espionage is not Small

In this Appendix, we provide more details of two equilibria of the game with industrial espionage discussed in Section 1.6. In particular, we establish two propositions dealing with large and intermediate spying costs, where the next two subsections first investigate for which values of c one can sustain the equilibrium outcomes of the previous two sections when firms have the possibility to spy on each other.

1.B.1 Equilibria without Spying

It is clear that for a large enough spying cost, firms will not engage in industrial espionage. To determine boundary costs where spying is not profitable, we compare the candidate equilibria under spying with the pooling equilibria under private information. This keeps the analysis focused without qualitatively affecting the results. In order to establish the cutoff value of c such that no spying constitutes an equilibrium, we need to derive the optimal strategy of a firm that decides to spy on its rival when it is not expected to do so. If the deviating firm discovers that its rival is unconstrained, it has no profitable deviation and will charge α in both periods (if it has available capacity in the second period), obtaining $\frac{3}{2}\alpha$ if unconstrained and α if constrained. If, on the other hand, it discovers the rival is constrained the optimal pricing depends on the value of α. Clearly, it could guarantee itself monopoly profits in the second period by letting the rival sell for sure in the first period (by deviating to a higher price). If she herself is, however, unconstrained, it may stick to the equilibrium (pooling) pricing strategy, selling with probability 0.5 in the first period, yielding an overall expected profit of 2α, or, if it does not sell, which also happens with probability 0.5, make monopoly profits in the second period. clearly, the overall profit of $\frac{1+2\alpha}{2}$ is larger than 1 if $\alpha > 0.5$. The associated profit of spying is therefore

$$(1-\alpha)\alpha\frac{3-\alpha}{2} + \alpha\left[\alpha \cdot 1 + (1-\alpha)\cdot \max\{1, \frac{1+2\alpha}{2}\}\right] - c,$$

where the first term reflects the profit if the rival is unconstrained, and the second - if the rival turns out to be constrained.

Thus, it is not optimal to deviate if

$$c > \alpha(1-\alpha)\left[\max\left\{1, \frac{1+2\alpha}{2}\right\} - \frac{\alpha}{2}\right] \equiv c_N$$

yielding the following result:

Proposition A.8. *If $c > c_N$, there exists an equilibrium where firms choose not to spy and price as in the pooling equilibrium of Proposition 1.3.*

In case that the continuation equilibrium is semi-separating, the value of c_N would be modified slightly since (i) the equilibrium profit of a non-spying unconstrained firm falls, and (ii) the equilibrium profit of an unconstrained firm that learns that its rival is unconstrained also falls. It can be shown that the net benefit from spying increases as $\overline{p}$ increases. In the limit case in which $\overline{p} = 1$, the corresponding cutoff value is $(1-\alpha)\alpha$.

1.B.2 Asymmetric Equilibria for Intermediate Spying Cost

Next, we consider whether, and if so under what conditions, an asymmetric equilibrium exists where one firm spies and the other does not. The equilibrium play in such an equilibrium mimics the one under disclosure with the roles reversed. It follows from Section 1.5 that the ex-ante expected candidate equilibrium profits are equal to $\alpha(3 - 3\alpha + \alpha^2) - c$ and $\alpha(2-\alpha)$ for the spying and non-spying firm, respectively.

To see for which values of c these strategies constitute an equilibrium, we should verify whether the spying and non-spying firms have incentives to deviate. If a firm is supposed to spy and it does not, it will optimally choose a first-period price of α, yielding α if constrained and (for out-of-equilibrium beliefs satisfying the Intuitive Criterion) $\alpha(2-\alpha)$ if unconstrained. It follows that the firm that is expected to spy will do so if $c \leq \alpha(1-\alpha)$. Similarly, if the firm that is supposed not to spy deviates and spies, it may tailor its price to the information. Given equilibrium play, this ability is useless if it has only one unit, as it would have obtained monopoly profits regardless. If, instead, the deviating firm is unconstrained it will allow its rival to sell in the first period and reap the monopoly rents in the second period, increasing its expected pay-off by $\alpha(1-\alpha)^2$. It then follows that

Proposition A.9. *If $\alpha(1-\alpha)^2 \leq c \leq \alpha(1-\alpha)$, an asymmetric equilibrium exists where one firm spies for sure and the other does not. The operating profit of the spying firm is larger than that of the non-spying firm, but the difference is smaller than the spying cost.*

Interestingly, as $c_N < \alpha(1-\alpha)$, there exists a range of c values such that the asymmetric equilibrium characterized in this section co-exists with the equilibrium where no firm spies. It is clear that industrial espionage by one firm increases the market power of both firms and that it is therefore anti-competitive and bad for consumer welfare. Also of interest is that the non-spying firm always makes a profit of $\alpha(2-\alpha)$, which is the ex-ante equilibrium profit if the uncertainty is revealed before prices are chosen. A firm does not need to know its rival's capacity: it is enough that the rival knows its capacity! Finally, as the equilibrium only exists for $c \geq \alpha(1-\alpha)^2$ it is clear that in this equilibrium the profit $\alpha(3-3\alpha+\alpha^2)-c$ of the spying firm is smaller than the profit $\alpha(2-\alpha)$ of the non-spying firm.

Proof of Proposition A.9. In this proof, we provide more detail about why deviating from the equilibrium strategies is not optimal. First, if the spying firm, *i.e.*, the firm that according to the asymmetric equilibrium is supposed to spy, deviates and decides not to spy, the non-spying firm does not observe this deviation and sticks to its first-period pricing, so that the spying firm only knows that with probability α the rival is constrained and charges 1 and with probability $1-\alpha$ randomizes uniformly above α. The deviation profits clearly depend on the continuation pricing strategy and on the beliefs of the rival. To give the candidate equilibrium maximal chance, we look for every possible pricing strategy at reasonable beliefs of the rival that create the lowest deviation pay-off. The candidate equilibrium can be sustained as an equilibrium if all of the possible pricing strategies yield a lower pay-off than the equilibrium pay-off for some reasonable belief of the rival.

If, after deviating and not spying, the firm sets a first-period price of α independent of whether or not it itself is constrained, its deviation profit is (at least) equal to $\alpha+\alpha(1-\alpha)^2$. To see this, note that with this pricing strategy, the firm sells and the first-period profit equals α. A firm can only make a profit in the second period if it is unconstrained. As α is the equilibrium first-period price if the non-spying firm is unconstrained, the unconstrained, non-spying firm has to believe that the rival plays according to the candidate equilibrium so the expected second-period profit equals α. However, for the constrained non-spying firm α is an out-of-equilibrium price, and it may well believe that its rival is unconstrained and in that case, the expected second-period profit equals 0. Thus, the deviating firm IS only makes a second-period profit of α if both firms are unconstrained, yielding an overall deviation profit of $\alpha + \alpha(1-\alpha)^2$.

All other first-period prices can generate lower pay-offs for reasonable beliefs of the non-spying firm. After a first-period price by the spying firm larger than 1, the non-spying firm may believe that its rival is unconstrained and in that case the spying firm will only make an expected profit of α (in the second period). On the other hand, after a first-period price by the spying firm that is larger than $\alpha + \varepsilon$ and smaller than 1, the non-spying firm may again believe its rival is unconstrained and in that case the spying firm will only make an expected profit of α (in the first period). Thus, the spying firm does not want to deviate if $\alpha(3 - 3\alpha + \alpha^2) - c \geq \alpha + \alpha\ (1-\alpha)^2$ or $c \leq \alpha(1-\alpha)$.

Next, consider that the non-spying firm deviates and decides to spy, which is not observed by its rival. As the spying firm knows the rival's type and lets a non-spying constrained rival sell at a price of 1 anyway, the non-spying firm does not gain anything from spying. If, however, the non-spying firm is unconstrained, learning the capacity of its rival can increase its expected profit. If it learns that its rival is constrained, the non-spying firm can still push its rival to sell in the first period (as the spying rival believes a non-spying firm would do), but then set the monopoly price in the second period. In case the spying firm is unconstrained as well, the deviating non-spying firm cannot do better than choosing the equilibrium prices. Hence, the ex-ante expected deviation profit of the non-spying firm equals $\alpha \cdot 1 + (1-\alpha)(\alpha \cdot 1 + (1-\alpha) \cdot \alpha) - c = \alpha(3 - 3\alpha + \alpha^2) - c$. This is smaller than the equilibrium profit of $\alpha(2-\alpha)$ if $c \geq \alpha(1-\alpha)^2$.

1.C Extensions and Supplementary Materials

The following Appendix contains supplementary materials on the extensions. It is organized in the following way. Section 1.C.1 discusses the equilibrium outcome under an alternative tie-break assumption, when firms set the same price. Section 1.C.2 relaxes the assumption that firms are able to observe each others' prices. Section 1.C.3 contains additional materials on the extension of Random Arrival of Consumers in Section 1.4.3.1. Similarly, Section 1.C.4 expands on Section 1.4.3.4 about Loyal Consumers. Finally, Section 1.C.5 displays additional graphs for different model parameters' values.

1.C.1 Alternative Tie-breaking Rules

In this part, we show that the equilibrium outcome of the game with hidden capacities does not depend on the specific tie-breaking rule that we assumed in the main model. In particular, we are interested in the alternative tie-breaking rule that allocates the demand evenly among all firms that charge the lowest price instead of randomly allocating all demand to one firm.

First, as mentioned in footnote 12 of the main text observe that the pooling equilibrium outcome can be supported by an asymmetric equilibrium in which one of the firms, say firm 1, charges $p = \alpha$, while its rival randomizes in an interval $(p, p + \varepsilon)$, for $\varepsilon > 0$ sufficiently small. In this case, there are no ties in equilibrium (so the tie-breaking rule - whatever it is - does not apply) and the same equilibrium outcome emerges. In particular, a constrained firm 2 has no incentive to deviate as its expected (second-period) profit equals α and any deviation in the first period cannot yield more profit. Like in Section 4, whether an unconstrained firm 2 has an incentive to deviate depends on the out-of-equilibrium beliefs if firm 1 observe its rival sets a first-period price smaller than α. As only unconstrained firms may have such an incentive to deviate, it is natural that firm 1 believes a deviating, undercutting firm is unconstrained, resulting in zero expected second-period profits, implying that a deviation by an unconstrained firm 2 is not profitable.

Second, the pooling equilibrium in pure strategies can also be supported by any tie-breaking rule. Indeed, consider the pooling equilibrium in which both firms charge α in the first period and then randomize in the second period in $[\alpha, 1]$. Moreover, consider that given the first-period prices firm i's first-period market share equals q_i. Notice that the on-path continuation game following such a split of demand in the first period is such that with probability $1 - \alpha$ firm i is unconstrained while with probability α firm i has capacity to serve exactly $1-q_i$ share of the market. It follows that the continuation profit of a constrained firm is $(1-q_i)\alpha$, and this firm will randomize in the interval $[\alpha, 1]$. Thus, the overall expected pay-off of a constrained firm i is equal to α and independent of q_i. Deviating in either period cannot improve upon this expected profit. The unconstrained firm setting the lowest price in the support sells for sure and obtains α, which is also its maximum continuation profit. Thus, the overall expected pay-off of a constrained firm i equals $\alpha(1+q_i)$. For the same reason as above, an unconstrained firm also does not want to deviate in the first period. Thus, for any tie-breaking rule the pooling equilibrium remains an equilibrium.

Finally, consider the semi-separating equilibrium strategies described in Section 1.4.1. If a tie occurs, it must be at the highest equilibrium price $\bar{p}$. If a firm sells quantity q_i at $\bar{p}$, it expects to make $q_i\bar{p} + (1 - q_i)\theta(\bar{p}) = \bar{p}$, regardless of q_i. Hence, the same first-period equilibrium strategies constitute an equilibrium (together with the natural adaptation of the second-period strategy).

1.C.2 Hidden Prices

In certain markets it may be difficult for a firm to observe its rival's price. Since prices act as signals, it is natural to ask whether the results would significantly differ in an environment in which prices are unobserved. We now argue that the pooling equilibrium outcome we characterized in the beginning of Section 1.4 is the natural outcome for such markets as it is the unique equilibrium outcome.

First, just like in the case of alternative tie-breaking rules, it is easy to see that the asymmetric equilibrium that induces the pooling equilibrium outcome remains an equilibrium regardless of the information about prices. Second, it is straightforward to check that the symmetric pooling equilibrium breaks down if prices are hidden. As deviations are not observed if they do not result in another firm selling in the first period, an unconstrained firm can undercut its competitor and pretend that it simply was lucky in the first period and was selected by the tie-breaking rule. For this very same reason, a separating equilibrium cannot be supported. Third, a similar argument can be used to show that no semi-separating equilibrium can exist in this case. Recall that in a semi-separating equilibrium, the expected profit of selling the second unit in the second period must be the same, regardless of the price chosen in the first period. This requires that higher first-period prices induce higher continuation profits. But any firm that charges a marginally lower price in the first period can simply mimic the second-period behavior of a firm that charged a higher first-period price, making such a deviation profitable (as it is unobserved).

1.C.3 Random Arrival of Consumers

In this Appendix we show that not only the pooling equilibrium can be extended to random arrivals of consumers, but also the semi-separating equilibria. Given that in every period a consumer is present only with probability $\lambda \in (0,1)$, the second-period equilibrium profit of the highest-pricing firm is $\lambda^2\theta(p)$, where $\theta(p)$ is the posterior probability that it will be a monopolist in the second period. This is also the equilibrium profit of her rival (in case it remains active). In the first period, then, a constrained firm that charges p expects

$$\Pi^c(p) = Q(p)\left(\lambda p + (1-\lambda)\theta(p)\lambda^2\right) + \int_{\underline{p}}^{p} -Q'(x)\lambda^2\theta(x)dx,$$

while the unconstrained firm expects a profit of $\Pi^c + Q(p)\theta(p)\lambda^2$. As in the main model, this implies that $Q(p)\theta(p)$ must be a constant. The first order condition reads:

$$Q'(p)\lambda\left(p - \theta(p)\lambda^2\right) + Q(p)\lambda(1 + (1-\lambda)\lambda^2\theta'(p)) = 0$$

since $Q'(p)\theta(p) + Q(p)\theta'(p) = 0$, we obtain

$$p - \theta(p)\lambda^2(2-\lambda) - \frac{\theta(p)}{\theta'(p)} = 0$$

which yields

$$\theta(p) = \frac{-p}{\lambda^2(2-\lambda)W_{-1}\left(\frac{pc}{\lambda^2(2-\lambda)}\right)}.$$

At the upper bound, we must have that the constrained firm is indifferent between selling or not: $\lambda\bar{p} = \lambda^2\theta(\bar{p})$, so $\theta(\bar{p}) = \bar{p}/\lambda$. The unconstrained firm must also be indifferent. Hence, $2\theta(\bar{p}^-) = \theta(\bar{p}) = \frac{\bar{p}}{\lambda}$. Hence,

$$\lambda^2(2-\lambda)W_{-1}\left(\frac{\bar{p}c}{\lambda^2(2-\lambda)}\right) = \frac{\bar{p}}{2\lambda}.$$

Solving this equation yields

$$c = \frac{1}{2\lambda} + \lambda^2(2-\lambda)\log(2\lambda p).$$

Using the same arguments derived in the Proof of Proposition 4, we obtain that for every $\bar{p} \in [\frac{\lambda\alpha}{2-\lambda}, \lambda]$, there exists a semi-separating (pooling if $\bar{p} = \lambda\alpha$), such that the first period prices lie in $[\frac{\lambda\alpha}{2-\lambda}, \bar{p}]$ and the second-period prices in $[\theta(p)\lambda, 1]$.

1.C.4 Captive Consumers

Complete Information about Capacities. Each firm has the same probability $\sigma \leq \frac{1}{2}$ of capturing the whole demand of the arriving consumer. There are 3 cases: (i) both firms constrained, (ii) one firm constrained, one unconstrained, (iii) both unconstrained.

Case (i): Even in this simple case with two constrained firms, which can split the market, the equilibrium from the perfect competition case does not carry on. In other words, an equilibrium where both firms price at the monopoly price in the first period does not exist. If it did, both firms ex-ante expect $\pi = \frac{1}{2}1 + \frac{1}{2}(1-\sigma) = \frac{2-\sigma}{2}$ as the firm that does not sell today, expects $1-\sigma$ tomorrow. Deviations upwards are not profitable, but a firm deviating downwards expects $\pi^C(1^-) = (1-\sigma)1^- + \sigma(1-\sigma) = 1-\sigma^2 > \frac{2-\sigma}{2}$. In this sense, there also does not exist an equilibrium in pure strategies, where the firms price at a price $p^* < 1$ as the same incentives to undercut arise.

Looking for an equilibrium, where both firms mix implies that the firms have an expected profit of

$$\pi^C(p) = [\sigma + (1-2\sigma)(1-F^C(p))]p + [\sigma + (1-2\sigma)F^C(p)](1-\sigma)$$

The firms need to be indifferent between selling at p or pricing at the monopoly price. This implies that their mixing strategy is given by $F^C(p) = \frac{1-2\sigma+2\sigma^2-(1-\sigma)p}{(1-2\sigma)(1-\sigma-p)}$ with support $p \in [\frac{1-2\sigma+2\sigma^2}{1-\sigma}, 1]$. Both firms' ex-ante expected profit is equal to $\pi^C = 1-\sigma+\sigma^2$.

In the limits of $\sigma \in [0, \frac{1}{2}]$ the lower bound of the price distributions converges to the monopoly price. Hence, as $\sigma \to 0$ the equilibrium converges to the Dudey equilibrium, as firms split their sales among the periods. For $\sigma > 0$ there is paradoxically stronger competition, even though, competition gets relaxed. The reason lies in the continuation profit $1-\sigma < 1$. The Dudey equilibrium requires that firms are able to agree for equal probability of sale in the first period at monopoly price, however, this does not translate to equal split among periods, as there is always the chance that the consumer tomorrow is captive of the sold-out rival. This effect gets mitigated as σ becomes large and $\sigma \to \frac{1}{2}$, as then the market is exogenously evenly split between them in each period, leading to an optimal strategy to maximize the price and increase profits.

Case (ii): The Dudey equilibrium, where the constrained firm can sell in the first period at the monopoly price, no longer exists. The reason lies in the fact, that even if both firms are active in the second period, they expect strictly positive profits of σ. This gives an incentive for the unconstrained firm to undercut the monopoly price and sell today and expect a positive profit tomorrow.

There does not exist an equilibrium, where the unconstrained firm pushes the constrained firm to sell first at a price lower than the monopoly price. The reason is that the unconstrained firm has a share of captives to which it can always sell, thus, would charge them the monopoly price. A constrained firm would then have an incentive to increase its price to the monopoly price as well. Therefore, the monopoly price 1 should always be in the support of the equilibrium strategies of both firms.

Following the intuition of Dudey, where the unconstrained firm prefers to charge higher prices and let the constrained firm sell out first, we look for an equilibrium, where the unconstrained firm puts a mass ω at the top of the price distribution. The ex-ante expected profits of both firms are given by

$$\begin{aligned}\pi^C(p) &= [\sigma + (1-2\sigma)(1-F^U(p))]p + [\sigma + (1-2\sigma)F^U(p)]\sigma \\ \pi^U(p) &= [\sigma + (1-2\sigma)(1-F^C(p))](p+\sigma) + [\sigma + (1-2\sigma)F^C(p)](1-\sigma)\end{aligned}$$

The optimal mixing strategy of the constrained firm, keeping the unconstrained firm indifferent between charging the monopoly price and any price in the support $p \in [\frac{\sigma+(1-2\sigma)^2}{1-\sigma}, 1)$ is by mixing according to $F^C(p) = 1 - \frac{\sigma(1-p)}{(1-2\sigma)(p-(1-2\sigma))}$. The unconstrained firm mixes over the same interval according to $F^U(p) = \frac{p-(1-2\sigma)}{p-\sigma} - \frac{\sigma(1-p)}{(1-2\sigma)(p-\sigma)}$ with a mass point $\omega = \frac{1-3\sigma}{1-\sigma}$ at the top. For $\sigma \in [0, \frac{1}{3}]$ this equilibrium retains some main characteristics from the perfect competition case. Namely, the unconstrained firm charges higher prices expecting to be a monopolist in the second period, while the constrained firm charges lower prices in order to sell first and avoid competition later. However, this is an equilibrium as long as $\sigma \leq \frac{1}{3}$.

For $\sigma \in (\frac{1}{3}, \frac{1}{2})$ the unconstrained firm prefers to compete for demand in both periods, as competition is relatively weak and the continuation profit is relatively high. In this case, the constrained firm prefers to aim and serve its captives at the maximum price, therefore it is the one setting a mass at the top. The unconstrained firm mixes according to $F^U(p) = 1 - \frac{\sigma(1-p)}{(1-2\sigma)(p-\sigma)}$ over the support $p \in [2\sigma, 1)$, while the constrained firm mixes according to $F^C(p) = \frac{(1-\sigma)(p-2\sigma)}{(1-2\sigma)(p-(1-2\sigma))}$ with a mass $\omega = \frac{3\sigma-1}{2\sigma}$ at the top. Interestingly, here as well, as $\sigma \to \frac{1}{2}$ the lower bound of the distribution converges to the monopoly price reflecting the decreasing competition and the convergence to an even split of the market like in the case with two constrained firms.

Case (iii) with two unconstrained firms is equivalent to repeated imperfect Bertrand competition, where both firms mix over the interval $[\frac{\sigma}{1-\sigma}, 1]$ in each period according to $F(p) = 1 - \frac{\sigma(1-p)}{(1-2\sigma)p}$ and expect a profit of 2σ.

Incomplete Information about Capacities. Let $\pi(\theta)$ denote the equilibrium continuation profit (of both firms) when the firm who sold is believed to be constrained with probability θ and let $\tilde{\pi}(p) := \pi(\theta(p))$. In particular $\pi(\theta) = \sigma + (1-2\sigma)\theta$ is affine in θ.

Proposition A.10. *There exists some $\bar{\sigma} \in (0, 1/2)$ such that*

- *If $\sigma < \bar{\sigma}$, constrained firms randomize in $(\underline{p}, 1)$, and unconstrained firms randomize in $(\underline{p}, \bar{p})$, and $\theta(p)$ is increasing and convex in $(\underline{p}, \bar{p})$ and $\theta(\bar{p}) = 1$.*
- *If $\sigma > \bar{\sigma}$, both types of firms randomize in $(\underline{p}, 1)$, and $\theta(p)$ is increasing and convex in the whole support, with $\theta(\underline{p}) \in (\alpha, 1)$.*

Proof. The presence of captives implies that the upper bound of the price distribution must necessarily be 1 and that there are no mass points. To see why consider any upper bound lower than 1 and notice that the constrained type must be indifferent between selling now at that price or not (for otherwise, it would undercut). However, in this case it can profitably deviate to a higher price that has the same continuation profit in case it does not sell but a higher profit now. Captives make this deviation strictly profitable. The same argument shows that no mass point can exist in the interior of the interval. To see that no mass point can exist at the upper bound, notice that captives imply that the continuation profit is bounded above by 1 and so constrained firms would prefer to undercut the mass point than expose themselves to losing to a rival with the same price.

It follows that there are only two types of equilibria, determined by some cutoff value $\bar{\sigma}$ (see below). First, if $\sigma > \bar{\sigma}$, the equilibrium involves both types of firms randomizing in an interval $(\underline{p}, 1]$, with $\underline{p}$ being the lowest price that a constrained firm is willing to charge. Namely,

$$(1-\sigma)\underline{p} + \sigma\pi(\alpha) = \sigma + (1-\sigma)\pi(\alpha).$$

In words, the constrained firm must be indifferent between charging $\underline{p}$ and selling to any consumer that is not captive with its rival now, and charging the monopoly price and selling to its own captives. In both cases, the continuation profit in the case of not selling is the same because not selling at the highest or the lowest price is uncorrelated with the type of the opponent. This yields,

$$\underline{p} = \frac{(1-2\sigma)\pi(\alpha) + \sigma}{1-\sigma}.$$

The same argument used as in the Proof of Proposition 1.4 in the case of homogeneous consumers shows that the continuation profit following any price in the support p is given by

$$\tilde{\pi}(p) = -\frac{p}{W_{-1}(-cp)},$$

where c is pinned down by the boundary condition $\tilde{\pi}(1) = \bar{\pi}$. In particular, $c(\bar{\pi}) = \frac{1}{\bar{\pi}\exp(\frac{1}{\bar{\pi}})}$, where we make explicit the dependence of c on the specific continuation profit at the upper bound of the price distribution. Finally, we can use the isoprofit condition of unconstrained firms to obtain

$$\tilde{\pi}(\underline{p})(1-\sigma) = \tilde{\pi}(1)\sigma = \bar{\pi}\sigma.$$

Combining these two yields the following equation:

$$1 + \frac{1-2\sigma}{\sigma}\pi(\alpha) = -\bar{\pi}W_{-1}\left(-c(\bar{\pi})\frac{(1-2\sigma)\pi(\alpha)+\sigma}{1-\sigma}\right).$$

This equation has a solution in $\bar{\pi} \in (\pi(\alpha), \pi(1))$ for every $\sigma \geq \bar{\sigma}$, where $\bar{\sigma}$ is implicitly defined as

$$1 + \frac{1-2\bar{\sigma}}{\bar{\sigma}}\pi(\alpha) = -\pi(1)W_{-1}\left(-c(\pi(1))\frac{(1-2\bar{\sigma})\pi(\alpha)+\bar{\sigma}}{1-\bar{\sigma}}\right).$$

Numerical results show that $\bar{\sigma}$ approaches 0 as $\alpha \to 0$ and $\bar{\sigma}$ approaches $1/2$ as $\alpha \to 1$.

If $\sigma < \bar{\sigma}$, unconstrained firms cannot be lured to charge $p = 1$ no matter how high the continuation profits. In this case, the equilibrium requires that $\theta(p) = 1$ for $p \geq \bar{p}$, for some $\pi(1) < \bar{p} < 1$. Constrained firms randomize in $(\bar{p}, 1)$ trading off higher selling probability against a higher markup. This yields the following randomization condition:

$$pQ(p) + \int_{\underline{p}}^{\bar{p}} -Q'(y)\pi(y)dy + \pi(1)(Q(\bar{p}) - Q(p)) = \sigma + \int_{\underline{p}}^{\bar{p}} -Q'(y)\pi(y)dy + \pi(1)(Q(\bar{p}) - \sigma).$$

Notice that the continuation profit is flat for $p > \bar{p}$. Hence:

$$pQ(p) = \sigma + \pi(1)(Q(p) - \sigma)$$

This equation can be readily solved to obtain

$$Q(p) = \frac{\sigma(1-\pi(1))}{p - \pi(1)}.$$

As before, in $(\underline{p}, \bar{p})$ both types of firms randomize. Their continuation profits are given by:

$$\tilde{\pi}(p) = \frac{-p}{W_{-1}(-cp)}.$$

We can solve for $c(\bar{p})$ using the fact that $\tilde{\pi}(\bar{p}) = \pi(1)$, resulting in

$$c(\bar{p}) = \frac{1}{\exp\left(\frac{\bar{p}}{\pi(1)}\right)\pi(1)}.$$

Finally, the equation that determines $\bar{p}$ is analogous as the case discussed before. Namely,

$$\tilde{\pi}(\underline{p})Q(\underline{p}) = \frac{-\underline{p}}{W_{-1}(-c(\bar{p})\underline{p})}(1-\sigma) = Q(\bar{p})\pi(1).$$

1.C.5 Additional Graphs

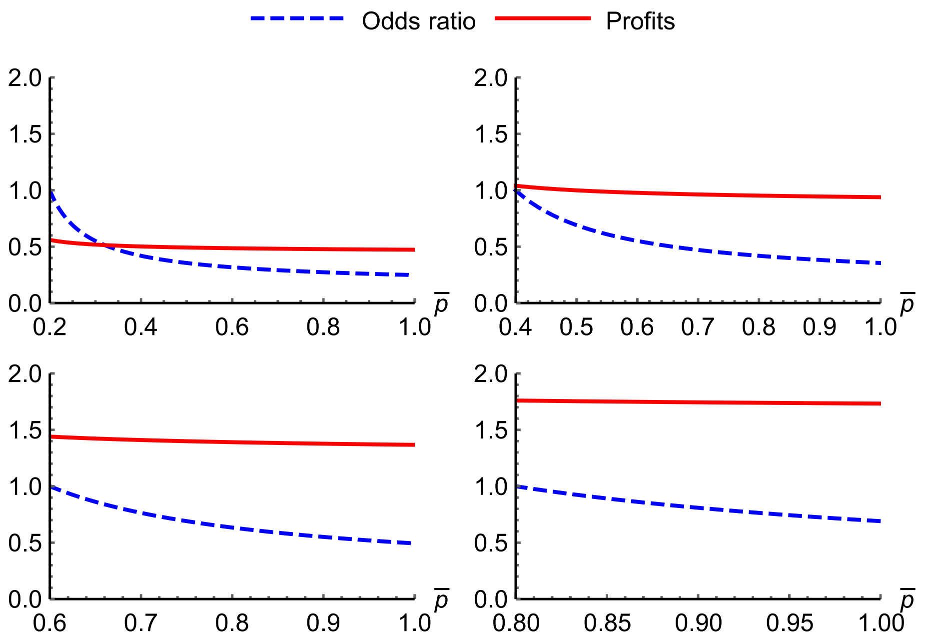

Figure 1.C.1: This plot establishes the ex-ante expected profit (red, solid) and the relative likelihood of a constrained firm selling in the first period (blue, dashed) in different equilibria (as a function of $\bar{p}$). Calibrated for (from left to right) $\alpha = \{0.2, 0.4, 0.6, 0.8\}$.

In section 1.4.1 we derived a continuum of semi-separating equilibria, which differed in the interval of prices $[\alpha, \bar{p}]$ over which firms randomized their strategies. In particular, two semi-separating equilibria differ in the upper bound $\bar{p}$. As stated in Proposition 1.5 and illustrated by Figure 1.2, a higher $\bar{p}$ is associated with lower expected profits, as there is a lower probability a constrained firm sells in the first period. In addition, Figure 1.C.1 displays the probability of a constrained firm selling in the first period (blue, dashed) and the ex-ante expected profit (red, solid) for different α. The graph shows that when the probability a constrained firm is in the market is low (first graph on the left for $\alpha = 0.2$), the odds ratio decreases quickly with $\bar{p}$. When the probability of a firm being constrained increases (the next three graphs), the odds ratio decreases at a slower pace, and profits increase as a result of the increased market power of firms. Nevertheless, profits decrease

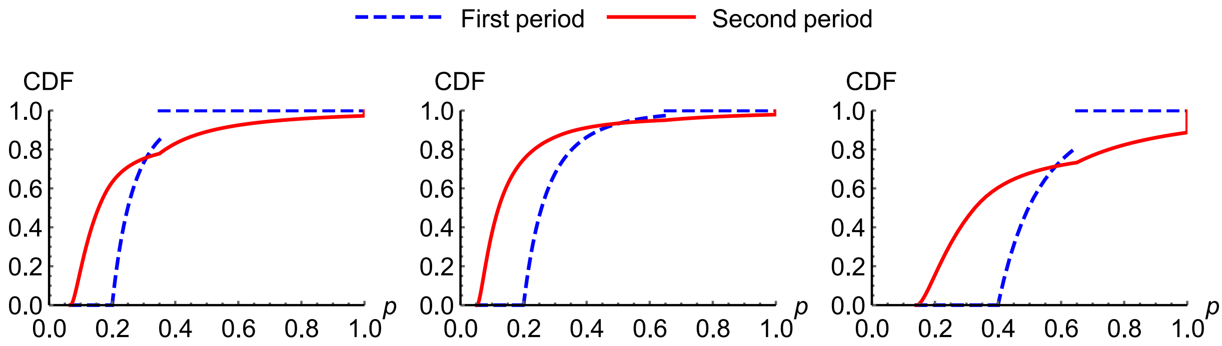

Figure 1.C.2: First- (blue, dashed) and second-period (red, solid) price distributions. Calibrated for (from left to right) $\{(\alpha = 0.2, \overline{p} = 0.35), (\alpha = 0.2, \overline{p} = 0.65), (\alpha = 0.4, \overline{p} = 0.65)\}$.

in $\bar{p}$ with the percentage difference between the pooling equilibrium and the equilibrium with $\bar{p} = 1$ decreasing as well from 16% to 1.6%.

Figure 1.C.2 complements Figure 1.3 by depicting the price distributions in the first (blue, dashed) and second (red, solid) period when α and $\bar{p}$ have different values. The Figure strengthens our result and shows that for different combinations of α and $\bar{p}$ the second-period price distribution is always wider, as firms mix over a larger set of prices. The price distributions intersect only once, as the second-period prices are more dispersed due to the remaining uncertainty of the presence of a rival in the respective period.

Finally, Figure 1.C.3 expands Figure 1.5 for different α. It shows the expected prices in the first and second periods for different semi-separating equilibria as a function of the upper bound of the distribution $\bar{p}$, for different values of α. In all semi-separating equilibria it holds that, as long as $\bar{p}$ is closer to α the second-period price is higher than the first-period price. When $\bar{p}$ becomes large and approaches 1, the probability that a

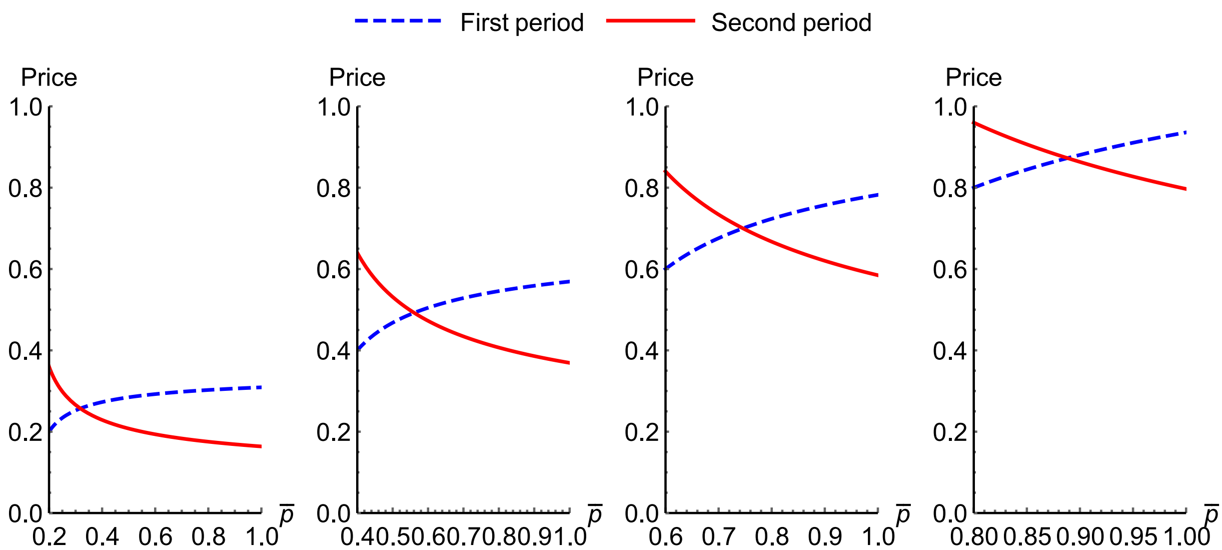

Figure 1.C.3: This plot shows the ex-ante expected first- (blue, dashed) and second-period (red, solid) transaction prices in different equilibria (as a function of $\bar{p}$). Calibrated for (from left to right) $\alpha = \{0.2, 0.4, 0.6, 0.8\}$.

constrained firm sells in the first period decreases (as depicted by Figure 1.C.1 above), which results in lower posteriors and lower prices in the second period. Naturally, as α increases, so do the prices in both periods.

Bibliography

Anderson, S. P., & Schneider, Y. (2007). *Price competition with revenue management and costly consumer search* (Technical report). University of Virginia, Department of Economics. Charlottesville, VA.

Anton, J. J., Biglaiser, G., & Vettas, N. (2014). Dynamic price competition with capacity constraints and a strategic buyer. *International Economic Review*, *55*(3), 943–958.

Armstrong, M., & Vickers, J. (2022). Patterns of competitive interaction. *Econometrica*, *90*(1), 153–191.

Aryal, G., Ciliberto, F., & Leyden, B. T. (2022). Coordinated capacity reductions and public communication in the airline industry. *The Review of Economic Studies*, *89*(6), 3055–3084.

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

Barrachina, A., Tauman, Y., & Urbano, A. (2014). Entry and espionage with noisy signals. *Games and Economic Behavior*, *83*, 127–146.

Bertsimas, D., & Perakis, G. (2006). Dynamic pricing: A learning approach. In *Mathematical and computational models for congestion charging* (pp. 45–79). Springer.

Besbes, O., & Lobel, I. (2015). Intertemporal price discrimination: Structure and computation of optimal policies. *Management Science*, *61*(1), 92–110.

Blume, A. (2003). Bertrand without fudge. *Economics Letters*, *78*(2), 167–168.

Board, S., & Skrzypacz, A. (2016). Revenue management with forward-looking buyers. *Journal of Political Economy*, *124*(4), 1046–1087.

Calvano, E., Calzolari, G., Denicolo, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, *110*(10), 3267–97.

Calvo, E., Cui, R., & Wagner, L. (2020). Disclosing product availability in online retail. *Manufacturing & Service Operations Management*. https://doi.org/10.1287/msom.2020.0882

Cheung, W. C., Simchi-Levi, D., & Wang, H. (2017). Dynamic pricing and demand learning with limited price experimentation. *Operations Research*, *65*(6), 1722–1731.

Clark, R., & Vincent, N. (2012). Capacity-contingent pricing and competition in the airline industry. *Journal of Air Transport Management*, *24*, 7–11.

Cui, R., Zhang, D. J., & Bassamboo, A. (2019). Learning from inventory availability information: Evidence from field experiments on amazon. *Management Science*, *65*(3), 1216–1235.

Dana, J. D. (1999). Equilibrium price dispersion under demand uncertainty: The roles of costly capacity and market structure. *The RAND Journal of Economics*, *30*(4), 632–660.

Dana Jr, J. D., & Williams, K. R. (2022). Intertemporal price discrimination in sequential quantity-price games. *Marketing Science*, *41*(5), 966–981.

Dasci, A., & Karakul, M. (2009). Two-period dynamic versus fixed-ratio pricing in a capacity constrained duopoly. *European Journal of Operational Research*, *197*(3), 945–968.

Dranove, D., & Jin, G. Z. (2010). Quality disclosure and certification: Theory and practice. *Journal of Economic Literature*, *48*(4), 935–63.

Dudey, M. (1992). Dynamic edgeworth-bertrand competition. *The Quarterly Journal of Economics*, *107*(4), 1461–1477.

EU Commission. (2011). Guidelines on the applicability of article 101 of the treaty on the functioning of the european union to horizontal co-operation agreements. *Official Journal of the European Union*, *54*, 1–72.

Fabra, N., & Llobet, G. (2023). Auctions with privately known capacities: Understanding competition among renewables. *The Economic Journal*, *133*(651), 1106–1146.

Farias, V., Monsch, M., & Perakis, G. (2011). Tractable notions of equilibrium in revenue management. *MSOM Annual Conference*.

Fudenberg, D., & Tirole, J. (1991). *Game theory* (Vol. 393). MIT press.

Gallego, G., & Hu, M. (2014). Dynamic pricing of perishable assets under competition. *Management Science*, *60*(5), 1241–1259.

Gal-Or, E. (1985). Information sharing in oligopoly. *Econometrica*, *53*(2), 329–343.

Garcia, D., Tolvanen, J., & Wagner, A. K. (2022). Demand estimation using managerial responses to automated price recommendations. *Management Science*, *68*(11), 7918–7939.

Harrington, J. E. (2018). Developing competition law for collusion by autonomous artificial agents. *Journal of Competition Law & Economics*, *14*(3), 331–363.

Hörner, J., & Samuelson, L. (2011). Managing strategic buyers. *Journal of Political Economy*, *119*(3), 379–425.

Huang, Y. (2022). Pricing frictions and platform remedies: The case of airbnb. *Available at SSRN 3767103*.

Janssen, M., & Rasmusen, E. (2002). Bertrand competition under uncertainty. *The Journal of Industrial Economics*, *50*(1), 11–21.

Joo, M., Gauri, D. K., & Wilbur, K. C. (2020). Temporal distance and price responsiveness: Empirical investigation of the cruise industry. *Management Science*, *66*(11), 5362–5388.

Lazarev, J. (2013). The welfare effects of intertemporal price discrimination: An empirical analysis of airline pricing in us monopoly markets. *New York University*.

Levin, Y., McGill, J., & Nediak, M. (2009). Dynamic pricing in the presence of strategic consumers and oligopolistic competition. *Management Science*, *55*(1), 32–46.

Lin, K. Y., & Sibdari, S. Y. (2009). Dynamic price competition with discrete customer choices. *European Journal of Operational Research*, *197*(3), 969–980.

Liu, Q., & Zhang, D. (2013). Dynamic pricing competition with strategic customers under vertical product differentiation. *Management Science*, *59*(1), 84–101.

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

McAfee, R. P. (1994). Endogenous availability, cartels, and merger in an equilibrium price dispersion. *Journal of Economic Theory*, *62*(1), 24–47.

McAfee, R. P., & Te Velde, V. (2006). Dynamic pricing in the airline industry. *Handbook on economics and information systems*, *1*, 527–67.

Montez, J., & Schutz, N. (2021). All-pay oligopolies: Price competition with unobservable inventory choices. *The Review of Economic Studies*, *88*(5), 2407–2438.

OECD. (2010). Information exchanges between competitors under competition law. *Policy Roundtables*.

Pitchik, C. (2009). Budget-constrained sequential auctions with incomplete information. *Games and Economic behavior*, *66*(2), 928–949.

Puller, S. L., Sengupta, A., & Wiggins, S. N. (2009). *Testing theories of scarcity pricing in the airline industry* (tech. rep.). National Bureau of Economic Research.

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.

Shapiro, C. (1986). Exchange of cost information in oligopoly. *The Review of Economic Studies*, *53*(3), 433–446.

Solan, E., & Yariv, L. (2004). Games with espionage. *Games and Economic Behavior*, *47*(1), 172–199.

Somogyi, R., Vergote, W., & Virag, G. (2023). Price competition with capacity uncertainty-feasting on leftovers. *Games and Economic Behavior*, *140*, 253–271.

Talluri, K., & Van Ryzin, G. J. (2006). *The theory and practice of revenue management* (Vol. 68). Springer Science & Business Media.

Williams, K. R. (2022). The welfare effects of dynamic pricing: Evidence from airline markets. *Econometrica*, *90*(2), 831–858.

Xu, X., & Hopp, W. J. (2006). A monopolistic and oligopolistic stochastic flow revenue management model. *Operations Research*, *54*(6), 1098–1109.

Chapter 2

Private Labels in Marketplaces[1]

2.1 Introduction

Large dominant platforms, like Amazon for example, introduce their own private label products in direct competition with the third-party sellers that they host. They operate a marketplace according to the agency vertical model - the platform owner dictates the fees that it collects from the sales of third-party sellers, while sellers set their product prices. Regulatory authorities both in the EU and US have raised the concern that when a dominant marketplace operator with considerable market power competes directly with outside sellers, it may have an incentive to skew competition in favor of its own products by raising the fees and potentially limiting access to the consumer side for outside sellers.

However, the purpose of private labels is not necessarily to foreclose competitors but rather to offer a lower-quality version of their products aimed at consumers with a lower willingness to pay. The literature on private labels in classic retailing shows that they improve consumer surplus by offering an additional variety and strengthening competition while also squeezing out the profits of the seller of the original product.[2] Classic retailers operate in the wholesale model where the third-party seller sells its product at a wholesale price to the intermediary, while the latter resells it to consumers. This raises the question of whether private labels in the agency model affect the market outcome differently and whether a marketplace operator has an incentive to manipulate competition, given its control over the fees that sellers pay.

I show that a marketplace operator introducing a lower-quality private label does not have an incentive to distort competition in favor of its own product and foreclose the seller of the original higher-quality product. On the contrary, after the introduction of its own product the intermediary decreases the fee and accommodates both products in the marketplace. Even though the intermediary is able to perfectly mimic the quality of the

[1]Part of this chapter has been published under the same name in the International Journal of Industrial Organization and has been circulated as a working paper vie2014 of Vienna Economics Papers.

[2]See Berges-Sennou et al. (2004) for theoretical overview and Draganska et al. (2010) and Meza and Sudhir (2010) for empirical studies.

outside seller and monopolize its product space, the intermediary prefers to differentiate its offer and chooses a lower quality for the private label product. The latter is consistent with studies showing that the products of Amazon are, on average, of inferior quality (Marketplace Pulse (2019), Ellis and Hicken (2020)).

Compared to the wholesale model, a private label introduced in the agency model may be less harmful to the third-party seller and improves the consumer surplus less. By choosing the fee, a marketplace operator extracts a larger share of profits from the seller than an intermediary who takes the wholesale price as given and has a stronger incentive to protect these profits. As a result, it chooses a higher price for its lower-quality product and, by decreasing the fee, leaves more space for the seller to operate than in the wholesale model. When a marketplace operator chooses the quality of the private label, it has a stronger incentive to differentiate and offers a lower-quality product than an intermediary in the wholesale model. Although the wholesale model benefits the seller more by allowing it to set the wholesale price, after the optimal choice of the quality of the private label, the seller is left with more profits in the agency model. Finally, while the new product necessarily improves consumer surplus, due to the higher prices and relatively lower quality of the private label it does so less in the agency model.

I now explain these results in more detail. Consider a natural setting where consumers are heterogeneous in their preferences for a particular quality of vertically differentiated products. There is a single seller who needs to use the services of an intermediary to distribute its higher-quality product to consumers. The intermediary operates a marketplace using the agency model and collects a per-unit fee from the sales of the seller. The intermediary can introduce a lower-quality copycat version of the product of the seller. When it does so, it competes in prices with the seller.

In the absence of a private label, the marketplace operator chooses the fee to maximize the revenue from a single higher-quality product. The seller sets the monopoly price for consumers, taking the fee as a given cost and as a monopolist, it passes only a share of it to consumers. The intermediary anticipates the pricing decision of the seller downstream and the fact that the fee is not fully passed to consumers, which gives it a first-mover advantage and the ability to extract a larger share of industry profits. This results in a classic double-marginalization problem where the final price paid by consumers is higher than if the intermediary and seller were vertically integrated.

When a marketplace operator introduces a private label, it competes in prices with the seller. Facing competition, the seller reduces its price to reflect that it can now only extract rents from the market due to the premium that its higher-quality product offers over the lower-quality private label. The reduction in its price is curbed by the fee it has to pay. The marketplace operator does not pay the fee, but incorporates it into its pricing as an opportunity cost - each unit the private label steals away from the seller is fee revenue lost. Due to the opportunity cost and strategic complementarity, the intermediary may choose to set a price higher than if it were to exclude the seller and only offer the private label as a monopolist. This is the case in particular if the fee remains at the pre-entry level.

After the introduction of the private label, the intermediary has an incentive to decrease the fee that it charges to the outside seller. First, competition from the lower-quality substitute limits the rents that the intermediary can extract from the higher-quality product as it did for the seller in the pricing stage. Further, in Bertrand competition prices of both products are strategic complements, meaning that if the seller faces an increasing fee, it increases its price but so does its rival. Due to this strategic complementarity and the fact that the intermediary incorporates the fee in its pricing decision as an opportunity cost, the pass-through of the seller's price with respect to the fee increases after the introduction of the private label. This limits the first-mover advantage of the marketplace operator and reduces even further the margin that can be collected from the seller's channel. Second, by increasing the fee, the intermediary can divert any demand from the seller's product to its own private label. However, because the lower-quality product generates a lower surplus for consumers, its margin can only partially mitigate the decrease in the pure margin collected from the seller's channel, resulting in a lower optimal fee.

To compare to the literature on classic retailers, consider a private label being introduced in the wholesale model. There the seller sets the wholesale price in the upstream contract while the intermediary sets the price for consumers. As pointed out by Johnson (2017), in the absence of a private label, the two vertical models are equivalent from a consumer perspective. The only difference is that the seller, by choosing the wholesale price, extracts a larger share of the industry profits than the intermediary. The two models diverge when a private label is introduced. The intermediary chooses the prices of both products acting as a multi-product monopolist and internalizes the cannibalization effects between them. This results in both products being priced independently from each other at the monopoly price. If the wholesale price remains at the pre-entry level, this implies that the intermediary does not change the price for the seller's product after introducing the private label. However, because some consumers switch to buying the cheaper product, the seller is forced to reduce the wholesale price to recover demand for its product.

Despite Bertrand competition, prices in the agency model are higher than in the wholesale model for a given quality of the private label. Even though the marketplace operator decreases the fee, the price of the private label is still higher than the one chosen by a stand-alone monopolist or the one chosen by the intermediary in the wholesale model. By strategic complementarity, the price of the seller's product is higher in the agency model. Still, because the marketplace operator protects revenues from the seller's product, it leaves the seller with more demand and space to operate in the agency model.

In both vertical models, there is a possibility for the intermediary to exclude the seller from the market if it perfectly mimics the quality of its product. For this reason, I endogenize the quality choice of the intermediary and show that in both vertical models it wants to differentiate its offer from the offer of the seller. Because in the agency model the marketplace operator extracts a larger share of profits from the seller's product, it has a stronger incentive to differentiate and offers a lower quality than in the wholesale model.

Finally, the introduction of a private label improves the consumer surplus and squeezes out the profits of the seller less in the agency model than in the wholesale model. The seller is better off in the wholesale model for a given quality of the private label. However, because the intermediary protects the revenue channel of the seller by differentiating its offer more, this results in higher profits for the seller in the agency model. Nevertheless, because of the higher prices and relatively lower quality, the consumer surplus is lower in the agency model.

Related literature. The literature on intermediaries operating in hybrid mode - running a marketplace for third-party sellers, while also selling its own products through the marketplace - has been growing steadily in the past few years.[3] The closest papers to mine that investigate how entry by the platform affects the market outcome are those by Anderson and Bedre-Defolie (2023), Etro (2023b), Zennyo (2022), Hagiu et al. (2022), and Anderson and Bedre-Defolie (2022). In a framework where there is free entry of sellers and monopolistic competition, Anderson and Bedre-Defolie (2023), Zennyo (2022) and Etro (2023b) show that entry by the platform can lead to some sellers exiting the market and the platform facing a trade-off between increasing fees to further induce exit, limit competition, and steer demand to its own products, or decreasing fees to recover entry and expand demand for all products. In this sense, the distinctive feature of these papers is that the fee (either per-unit or ad-valorem) charged to the sellers affects the variety found on the platform. Moreover, entry in these models is comparable to a form of vertical integration with existing sellers as discussed by Etro (2023b). Differently in this paper, I focus on an exogenous market structure that changes from a monopoly to a duopoly when the platform introduces its own product. Thus, entry by the marketplace operator expands and creates its own demand while also stealing away demand from the seller. This changes the demand conditions and pricing incentives for the hosted seller, and so restricts the margins that can be collected from its higher-quality product. The increased competition downstream drives the decrease of the per-unit fee in my setting.

Finally, Hagiu et al. (2022) and Anderson and Bedre-Defolie (2022) also consider the case where a seller offers a premium product and the platform can offer a lower-quality version. However, because consumers are homogeneous in their valuations of the two products, a private label in Hagiu et al. (2022) is never actively purchased and only serves as a threat to the seller to reduce its price and attract more consumers to the platform. In Anderson and Bedre-Defolie (2022) when the platform offers its own

[3] For an extensive overview see Anderson and Bedre-Defolie (2022) and Etro (2023a).

product it excludes the seller from the platform and competes with the seller through a horizontally differentiated channel.

Further papers focusing on the hybrid model address the question of the opening up of a marketplace by an intermediary to learn about new products (Hervas-Drane and Shelegia (2022)), or choosing the right non-linear contract when the costs of hosted sellers are unknown (Kang and Muir (2022)). Madsen and Vellodi (2022) and Lam and Liu (2021) investigate the role of data when demand is uncertain, with the former addressing innovation incentives of third-party sellers. Padilla et al. (2022) address the question on the use of consumer and seller data for improving the quality of the intermediation service. When consumers have to search for prices, Zennyo (2022) focuses on the possibility of a platform biasing the search algorithm in its own favor, while Kittaka and Sato (2022) explore the ability of a platform to distort the search order. Gautier et al. (2021) and Pouyet and Thomas (2021) consider the importance of network effects. Recently Tremblay (2022) analyzed the implications of fee discrimination, while Bisceglia and Padilla (2023) consider the effects of the presence of a private label when third-party sellers can collude on prices.

This paper borrows from the literature on private labels in the wholesale model. The closest studies to the analysis provided here are found in Mills (1995), Bontems et al. (1999) and Heese (2010) where an intermediary introduces a lower-quality private label to compete with the product of a known manufacturer. I complement these studies by investigating the effects of a private label in the agency model.

The paper is related to the literature on the differences in the vertical organization of markets. The comparison between the wholesale and agency models in the absence of a private label is theoretically founded in Johnson (2017). Gautier et al. (2023) show how, in search markets, a hybrid of the two models can arise. Hagiu and Wright (2015) discuss whether an intermediary wants to act as a reseller or marketplace depending on the cross-product spillovers caused by marketing activities. Complementing these studies, this paper compares the wholesale and agency models in the presence of a lower-quality private label, but does not consider the choice of an optimal vertical model.

The chapter is structured as follows. Section 3.2 presents the model. The analysis starts with a quick review of the vertical models in the absence of a private label as a benchmark in Section 2.3, then proceeds with the analysis of the introduction of a lower-quality private label in Section 2.4. Section 2.5 discusses the optimal choice of quality made by the intermediary. Section 2.6 relaxes some of the main model assumptions. Section 3.4 concludes. Proofs of the propositions that are not displayed in the main text can be found in Appendix 3.A, while additional material is available in Appendix 2.B.

2.2 The Model

On the supply side of the market, there is a single seller (S), who manufacturers product 1 with quality s_1 and needs to use the services of an intermediary to reach consumers. The intermediary can offer a private label product 2 of quality $s_2 \leq s_1$.[4] Both firms, the seller and the intermediary, have the same production technology leading to a marginal cost of producing quality s given by $c(s)$, which is convex and increasing in s. For ease of notation, I use c_i instead of $c(s_i)$ throughout Sections 2.3 and 2.4.

On the demand side, there is a unit mass of consumers who want to buy at most one unit of a product. Consumers differ in their willingness to pay for quality of products which is reflected by their type. In the spirit of Mussa and Rosen (1978), a consumer of type θ purchasing product with quality s_i at price p_i has utility $u(\theta, s_i, p_i) = \theta s_i - p_i$. The consumers' types are distributed over the interval $[0, \bar{\theta}]$ according to a monotone density function $g(\theta)$. Consumers have the option of not buying any product offered in the market, which gives them zero utility ($u(0) = 0$). When the seller offers its product of quality s_1 at price p_1 and the intermediary offers a private label with quality $s_2 \leq s_1$ at price $p_2 \leq p_1$, there are three separate groups of consumers. High type consumers, who have $\theta \geq \frac{p_1 - p_2}{s_1 - s_2} = \hat{\theta}$ prefer to buy the higher-quality product. Intermediate type consumers, who have $\hat{\theta} > \theta \geq \frac{p_2}{s_2} = \tilde{\theta}_2$ prefer to buy the lower-quality product. The rest leave the market without buying anything. The demands for the seller's and intermediary's product are, respectively, given by

$$Q_1(p_2, p_2) = 1 - G(\hat{\theta}) \quad \text{and} \quad Q_2(p_1, p_2) = G(\hat{\theta}) - G(\tilde{\theta}_2)$$

where $G(\theta)$ is the corresponding cumulative distribution function of θ.

When the intermediary operates a marketplace and employs the agency model, the timing of the game is the following. First, as part of the upstream contract stage, the intermediary determines the per-unit fee collected from the sales of the outside seller.[5] Then, on the downstream level, the participants in the marketplace set prices simultaneously for their products - the outside seller chooses the price for its own product and, if offered, the intermediary chooses the price for the private label, taking the fee as given (downstream pricing stage). In the classic wholesale model, the timing and decision variables are different. In the wholesale model, the outside seller first determines the wholesale price at which it sells its product to the intermediary (upstream contract stage), and then

[4]This reflects the fact that intermediaries are only able to replicate, but not improve the product of the seller. In Appendix 2.B.4 I look at the case of a higher-quality private label and show that the main results remain robust.

[5]In practice, marketplaces have various fee structures. For example, Amazon charges a per-unit fee of \$/€0.99 for small sellers, while offering a flat rate of \$/39.99€ for big sellers. On top of that, it charges ad-valorem fees on the final price of each sale, which vary across the categories. Here I assume a per-unit fee as it best represents a direct cost factor for the seller and allows for more tractable analysis. In Appendix 2.B.3 I discuss the case of an ad-valorem fee and show through numerical simulations that the main results remain robust.

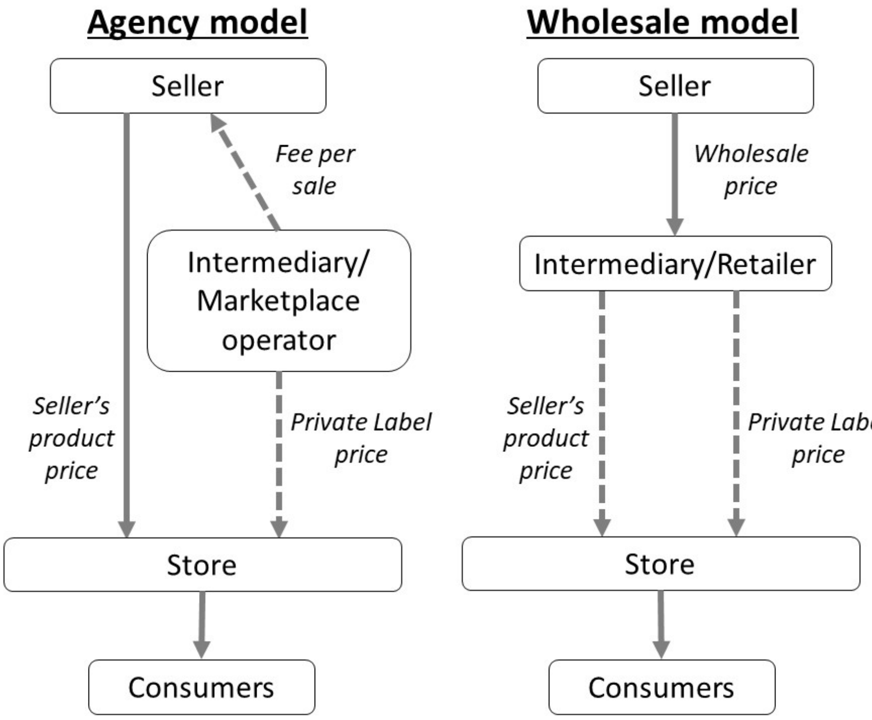

Figure 2.1: Difference between the agency and wholesale models. The timing of the decisions made by intermediary (dashed arrows) and seller (solid arrows) runs from top to bottom.

the intermediary determines the prices of both the seller's product and the private label (if offered), taking the wholesale price as given (downstream pricing stage). As these are dynamic games of complete information, the solution method applied is Subgame Perfect Equilibrium. Figure 2.1 summarizes the differences in the vertical organization and the sequence of decisions taken.

Throughout the analysis of this paper, I assume that the qualities offered in the market by either the seller or the intermediary are viable - they have a positive demand both when offered alone ($\bar{\theta}s_i - c_i > 0$ for $i = 1, 2$) and together ($\bar{\theta}s_i - c_i > \bar{\theta}s_j - c_j$ and $c_is_j > c_js_i$ for $s_i > s_j$), when priced at marginal cost.

Finally, in order to guarantee unique solutions, as well as introduce necessary simplifications at certain stages of the analysis, I make the following assumption over the demand-generating distribution of preferences $g(\theta)$:

Assumption 2.1. *The demand generating probability distribution function of preferences $g_{(\mu,\sigma,\xi)}(\theta)$ belongs to the family of Generalized Pareto Distributions*

$$g_{(\sigma,\xi)}(\theta) = \begin{cases} \frac{1}{\sigma}(1 + \frac{\xi\theta}{\sigma})^{-\frac{1}{\xi}-1} & for\ \xi \neq 0 \\ \frac{1}{\sigma}exp(-\frac{\theta}{\sigma}) & for\ \xi = 0 \end{cases}$$

with scale parameter $\sigma = -\xi\bar{\theta}$, and shape parameter $\xi \in [-1, 0]$.

The family of Generalized Pareto Distributions is a wide family that includes the family of exponential ($\xi = 0$) and the family of uniform ($\xi = -1$) distributions. The parameter ξ measures the curvature (the elasticity of the slope) of inverse demand. The restriction of $\xi \leq 0$ translates to demand being log-concave, which guarantees unique pricing equilibria.[6] The assumption of ξ being constant[7] has an important implication for the inverse of the hazard rate being linear in taste parameters, i.e. $\frac{1-G(\theta)}{g(\theta)} = \gamma(\theta) = \sigma + \xi\theta$. In pricing theory $\gamma(\theta)$ is also known as the Mills' ratio and is a comprehensive measure of demand sensitivity. Thus, a monopolist with quality s_1 collects a margin $p_1 - c_1 = \gamma(\tilde{\theta})s_1$, where $\tilde{\theta}_1 = \frac{p_1}{s_1}$ is the marginal consumer indifferent between buying the product or taking the zero outside option, which changes in price at a constant rate $\gamma' = \xi \leq 0$. This implies a constant pass-through rate of costs $\frac{\partial p_1}{\partial c_1} = \frac{1}{1-\xi}$. In general, Johnson and Myatt (2015) (which expands on Itoh (1983)) demonstrate that common demand distributions (unimodal beta, normal, gamma) often exhibit approximate linearity.

2.3 The Case without a Private Label

In this section, I focus on the benchmark case when there is no private label offered by the intermediary and there is only the product of the outside seller offered in the market. As already pointed out by Johnson (2017) whether the agency with a per-unit fee or the wholesale model is implemented, does not make a difference from the consumer or welfare perspective. In general, one of the firms sets the terms of trade (either per-unit fee or wholesale price), and the other firm sets the consumer price. This results in the same level of double marginalization in both vertical models leading to the same consumer price, demand, and industry profits. The only difference is in the allocation of industry profits - the firm moving on the upstream level (choosing the fee/wholesale price) extracts a larger share of industry profits. Hence, an intermediary is better off when the agency model is implemented which explains why Amazon has been pushing sellers from its retail to its marketplace business (see Berthene (2019)).

To see this in more detail, consider the agency model and the wholesale model with the third-party seller offering product 1 with quality s_1. In the absence of a private label, the share of consumers who purchase the product of the seller for price p_1 is $Q_1(p_1) = (1 - G(\tilde{\theta}_1))$, where $\tilde{\theta}_1 = \frac{p_1}{s_1}$ is the marginal consumer buying the last unit.[8] For ease of notation, I suppress the brackets in $Q(p_1)$ and present demand only as Q_1. The objective functions of the firms in the agency and wholesale models respectively are given by:

[6] See Benassi et al. (2019) for proof of existence for a family of log-concave functions in a vertically differentiated setting when firms have zero costs.

[7] Most recently, it has been used by Rhodes et al. (2021).

[8] The case is equivalent to assuming that the quality of the private label, its cost, and price are all equal to 0, yielding the same utility as the outside option.

Upstream	$\underset{f}{max} \ \ \pi_A = fQ_1$	$\underset{w}{max} \ \ \pi_S = (w - c_1)Q_1$
Downstream	$\underset{p_1}{max} \ \ \pi_S = (p_1 - f - c_1)Q_1$	$\underset{p_1}{max} \ \ \pi_W = (p_1 - w)Q_1$

where π_S denotes the profit function of the third-party seller, while π_A, π_W denote the profit function of the intermediary when, respectively, either the agency or wholesale model is implemented.

The two vertical models are equivalent and yield the same market outcome with $f = w - c_1$. There is a monopoly at each level, optimally choosing the transaction price. On the downstream level, the firm choosing the price for consumers takes f and w as a given cost and sets a margin and price:

$$p_1 - f - c_1 = p_1 - w = -\frac{Q_1}{dQ_1/dp_1} = \frac{1 - G(\tilde{\theta}_1)}{g(\tilde{\theta}_1)} s_1 = \gamma(\tilde{\theta}_1)s_1 \tag{2.1}$$

If demand is log-concave the RHS of the expression above is a monotone non-increasing function and there exists a unique optimal price p_1^*.

Now consider the upstream stage. As pointed out by Johnson (2017), one can alternatively assume that firms moving upstream choose price p_1, while w and f equilibrate respectively according to $w = p_1 - \gamma(\tilde{\theta}_1)s_1$ and $f = p_1 - \gamma(\tilde{\theta}_1)s_1 - c_1$ from optimality condition (2.1). This yields the same optimization problem:

$$\underset{p_1}{max} \ \ \pi = (p_1 - \gamma(\tilde{\theta}_1)s_1 - c_1)(1 - G(\tilde{\theta}_1))$$

with the solution to the F.O.C.:

$$p_1 - c_1 = \gamma(\tilde{\theta}_1)s_1 + [1 - \gamma'(\tilde{\theta}_1)]\gamma(\tilde{\theta}_1)s_1$$

The equation above represents the industry margin on each unit sold. From (2.1) the first term is the margin of the firm that sets the price. Hence, the margin an outside seller in a wholesale model and an intermediary in the agency model collect is $[1 - \gamma'(\tilde{\theta}_1)]\gamma(\tilde{\theta}_1)s_1$. Thus, the firm moving first obtains a larger share of the profits when $\gamma'(\tilde{\theta}_1) < 0$, or when demand is log-concave.

Alternatively, given the optimal price downstream (2.1), the firm moving upstream sets the optimal margin such that:

$$f = -\frac{Q_1}{dQ_1/dp_1}\left(\frac{\partial p_1}{\partial f}\right)^{-1} = w - c_1 = -\frac{Q_1}{dQ_1/dp_1}\left(\frac{\partial p_1}{\partial w}\right)^{-1} = \gamma(\tilde{\theta}_1)s_1[1 - \gamma'(\tilde{\theta}_1)] \tag{2.2}$$

where $\frac{\partial p_1}{\partial f} = \frac{\partial p_1}{\partial w} = \frac{1}{1-\gamma'}$ is the cost pass-through of the downstream monopolist. As long as $\gamma'(\tilde{\theta}_1) < 0$ the downstream monopolist absorbs any increase in its costs and passes less than 100% of it to the consumer side, which allows the upstream monopolist to increase this cost and collect the larger shares of industry profits.

Finally, for completeness, under Assumption 2.1, the prices chosen on the downstream level as a function of f and w respectively are given by

$$p_1(f) = \frac{\sigma s_1 + f + c_1}{1-\xi} \qquad p_1(w) = \frac{\sigma s_1 + w}{1-\xi}$$

while the optimal fee and wholesale margin are $f = w - c_1 = \frac{\sigma s_1 + \xi c_1}{1-\xi}$, leading to the same final price $p_1 = \frac{\sigma s_1(2-\xi)-\xi c_1}{(1-\xi)^2}$.

2.4 The Introduction of a Lower-Quality Private Label

In this section, I solve the game in the presence of a private label. Throughout the following analysis, consistent with the empirical observations, I assume that the intermediary in both vertical models offers a quality of the private label that is at most the quality of the third-party seller, *i.e.*, $s_2 \leq s_1$. Thus, the demand for the high-quality product is given by $Q_1(p_1, p_2) = 1 - G(\hat{\theta})$, and the demand for the lower quality is given by $Q_2(p_1, p_2) = G(\hat{\theta}) - G(\tilde{\theta}_2)$. Again, for ease of notation, I suppress the prices as arguments in Q_1 and Q_2. I first focus on the novel results arising in the agency model and then briefly introduce the well-known to the literature wholesale model.

2.4.1 Agency Model

Price competition downstream. In the presence of the private label, the outside seller and the intermediary choose prices simultaneously in order to maximize profits:

$$\max_{p_1} \ \pi_S = (p_1 - f - c_1)Q_1 \qquad \max_{p_2} \ \pi_A = fQ_1 + (p_2 - c_2)Q_2$$

The optimal pricing rule of the seller changes in the presence of the private label. The seller is no longer a monopolist and its product competes against the lower-quality product of the intermediary. Thus, the seller's optimal margin decreases to reflect that it is only able to extract the added value from its higher-quality product $(s_1 - s_2)$ and that it reacts strategically to price changes made by its rival:

$$p_1^A - f - c_1 = -\frac{Q_1}{dQ_1/dp_1} = \frac{1 - G(\hat{\theta})}{g(\hat{\theta})}(s_1 - s_2) = \gamma(\hat{\theta})(s_1 - s_2) \tag{2.3}$$

After the introduction of the private label, the seller's price becomes more responsive to changes in the fee, thus its fee pass-through increases to reflect that the rival can also pass some of the fee cost to consumers through strategic complementarity, *i.e.*, $\frac{\partial p_1}{\partial f} = \frac{1}{1-\xi} - \frac{\xi}{1-\xi}\frac{\partial p_2}{\partial f}$. More intuitively, the increase in the pass-through of the seller's price is a reflection of the decrease in its market power and the differentiated Bertrand competition with the private label substitute of its product.[9] How much the pass-through increases depends on the strength of competition with the private label and its optimal price.

[9] More generally, as argued by Weyl and Fabinger (2013) and emphasized by Ritz (2018) the rate of pass-through reflects the distribution of welfare in the market between consumers and producers. More

The intermediary chooses the optimal duopoly price for a good with a lower quality, but also incurs an opportunity cost represented by the fee:

$$p_2^A - c_2 = -\frac{Q_2}{dQ_2/dp_2} - \frac{dQ_1/dp_2}{dQ_2/dp_2} f = \underbrace{(1-\omega)\gamma(\tilde{\theta}_2)s_2 - \omega\gamma(\hat{\theta})(s_1 - s_2)}_{\text{duopoly margin}} + \underbrace{\omega f}_{\text{opp. cost}} \tag{2.4}$$

where $-\frac{dQ_1/dp_2}{dQ_2/dp_2} = \frac{g(\hat{\theta})s_2}{g(\tilde{\theta}_2)(s_1-s_2)+g(\hat{\theta})s_2} = \omega \in [0,1]$ is the diversion ratio of the private label product when its price increases. Because the private label is the lower-quality product, when the intermediary increases its price it loses a share ω to the higher-quality product and a share $1-\omega$ to the outside option. Thus, the opportunity cost is weighted down by the share of consumers lost to the higher-quality product.

The price of the private label has a positive pass-through of the fee not only because of strategic complementarity, but also because of the opportunity cost, which increases how much the price of the private label reacts to fee changes. This in turn leads to a higher price pass-through of the seller price.

Exploiting that $\gamma(\theta) = \sigma + \xi\theta$ and recalling that $\hat{\theta} = \frac{p_1-p_2}{s_1-s_2}$ and $\tilde{\theta}_2 = \frac{p_2}{s_2}$, we obtain the expressions for the optimal prices:[10]

Lemma 2.1. *For $f \leq \frac{\sigma s_1+\xi c_1}{1-\xi} + \frac{\sigma(s_1-s_2)+\xi(c_1-c_2)}{-\xi(1-\xi)(1-\omega)}$ the outside seller is not foreclosed from the marketplace and the prices are given by*

$$p_1^A(f) = \frac{\sigma s_1(1-\xi+\xi\omega) + c_1(1-\xi) - (\sigma s_2 + \xi c_2) + f(1-\xi-\xi\omega)}{1-2\xi+\xi^2(1-\omega)} \tag{2.5}$$

$$p_2^A(f) = \frac{\sigma s_2(1-\xi+\xi\omega) + c_2(1-\xi) - \omega(\sigma s_1 + \xi c_1) + \omega f(1-2\xi)}{1-2\xi+\xi^2(1-\omega)} \tag{2.6}$$

Proof. The seller gets foreclosed, when there is no demand for its product, or when $\bar{\theta} - \hat{\theta} = -\frac{\sigma}{\xi} - \frac{p_1^A(f)-p_2^A(f)}{(s_1-s_2)} \leq 0$. Given the optimal prices in (2.5) and (2.6) it is straightforward to derive the foreclosing fee threshold. □

When the quality of the private label converges to 0 the price of the seller converges to the level in (2.1). As the quality s_2 improves the price of the seller decreases under the competitive pressure. In the limit, where $s_2 = s_1$ there is effectively Bertrand competition and prices converge to the effective marginal cost $f + c_1$.

Interestingly, if the seller is not foreclosed, the intermediary may choose a higher price for the private label than if it were to exclude the seller and offer only the private label to consumers. Suppose that the fee is above the level in Lemma 2.1 and the seller is foreclosed. When the seller is excluded from the market, the intermediary is a monopolist offering a product of quality s_2. Thus, analogical to the case in Section 2.3, demand is given by $Q_2 = 1 - G(\frac{p_2}{s_2})$ and the price is chosen to maximize profits $(p_2 - c_2)Q_2$ which

competition decreases market power and pushes the prices closer to marginal costs, leaving more surplus for consumers.

[10] Note that ω is still a function of prices for $\xi \neq -1$.

yields an optimal price of $p_2 = \frac{\sigma s_2 + c_2}{1-\xi}$. However, if the fee is not too high and the seller is active in the market, by strategic complementarity and due to the opportunity cost, the intermediary may prefer to set a higher price for its product than the stand-alone monopoly one. This happens for a particular level of the fee as summarized by the following Lemma:

Lemma 2.2. *If $f > \frac{\sigma s_1 + \xi c_1}{1-\xi} + \frac{\xi}{1-\xi}\gamma(\hat{\theta})(s_1 - s_2)$ then the price of the private label exceeds the price chosen by a stand-alone monopolist.*

Proof. The result follows on from comparing Equation (2.6) to the stand-alone monopoly price $p_2 = \frac{\sigma s_2 + c_2}{1-\xi}$. □

Recall that pre-entry the fee that the intermediary collects from the seller's higher-quality product is given by $f = \frac{\sigma s_1 + \xi c_1}{1-\xi}$. If the fee is not adjusted after the introduction of the private label, according to Lemma 2.1, the seller is not excluded from the market. Furthermore, this fee level is above the one in Lemma 2.2. Thus, if the intermediary does not adjust the fee, the price it sets in competition with the seller is higher than what it would choose as a stand-alone monopolist. In the following section, I show that the intermediary does not want to foreclose the seller by increasing the fee; it prefers to decrease it instead.

Optimal per-unit fee. Anticipating the outcome of the competition downstream, the intermediary chooses the fee in order to maximize the revenues from both products:

$$\max_{f} \ \ \pi_A = fQ_1(p_1(f), p_2(f)) + (p_2(f) - c_2)Q_2(p_1(f), p_2(f))$$

Using the envelope theorem, the first-order condition of the intermediary's problem can be expressed as:

$$\frac{\partial \pi_A}{\partial f} = \frac{d\pi_A}{dp_1}\frac{\partial p_1}{\partial f} + \frac{d\pi_A}{df} = \left(f\frac{dQ_1}{dp_1} + (p_2(f) - c_2)\frac{dQ_2}{dp_1}\right)\frac{\partial p_1}{\partial f} + Q_1 \overset{!}{=} 0$$

Like any standard first-order condition, the fee is chosen to balance the lost revenues from the decrease in seller demand and the increase in revenues from charging a higher fee. However, in the presence of a lower-quality private label, any demand lost by the higher-quality product following its price increase is fully covered by the lower-quality product, *i.e.*, $-\frac{dQ_1}{dp_1} = \frac{dQ_2}{dp_1}$. Thus, revenues from the private label mitigate the loss from increasing the fee and create an incentive for the intermediary to charge a higher fee. Whether this fee will be higher than the original fee depends on the way that the introduced competition downstream affects the demand of the seller and pricing of the two products. To see how these interact in the final fee choice, rearranging is done to obtain the following:

$$f = \underbrace{-\frac{Q_1}{dQ_1/dp_1}\left(\frac{\partial p_1}{\partial f}\right)^{-1}}_{\text{pure marginalization}} \underbrace{-\frac{dQ_2/dp_1}{dQ_1/dp_1}(p_2(f) - c_2)}_{\text{revenue from private label}} = \gamma(\hat{\theta})(s_1 - s_2)\left(\frac{\partial p_1}{\partial f}\right)^{-1} + (p_2^A - c_2) \tag{2.7}$$

There are three main differences between the optimal fee chosen in the absence of a private label in (2.2) and the one chosen in the presence of a private label. The first two differences arise in the pure marginalization element which is the common element between the new and the original fee setting rule $-\frac{Q_1}{dQ_1/dp_1}\left(\frac{\partial p_1}{\partial f}\right)^{-1}$. As we have already seen in the pricing stage, the margin $-\frac{Q_1}{dQ_1/dp_1}$ decreases from the monopoly to the duopoly level, reflecting that firms can now extract from the higher-quality product only the premium that it offers over the lower quality. Then, recall that the pass-through $\frac{\partial p_1}{\partial f}$ increases to reflect the increase in competition and the opportunity cost, which the intermediary incorporates in the price of the private label. This further decreases the pure margin that can be collected from the seller's product. The second term in (2.7) reflects that the fee should extract at least as much revenue from the seller as it can recover through the private label. The private label generates a lower surplus for consumers due to its lower quality and lower margin. Therefore it can only partially compensate for the decrease in pure marginalization on the seller's higher-quality product and the fee ends up being lower than it would be in the absence of a private label.[11] The result is summarized in the proposition below:

Proposition 2.1. *An intermediary operating a marketplace and employing the agency model reduces the per-unit fee charged to the third-party seller after the introduction of a lower-quality private label.*

Alternatively, a more heuristic explanation for why the fee decreases after the introduction of a lower-quality private label is that if the intermediary maintains the pre-entry fee, the prices downstream end up being too high, especially the price of the lower-quality product. Recall from the pricing stage that if the intermediary were to offer only the private label to consumers, it would choose a lower price than it chooses in competition with the seller given the original fee level. As the intermediary has only one instrument to adjust the margins collected from the two vertically differentiated products, it decreases the fee as a way to correct for the high prices. Still, the fee is above the threshold defined in Lemma 2.2 for which the price of the private label is higher than the one chosen by a stand-alone monopolist.

The results in this and the previous section highlight how a marketplace operator does not want to distort competition in favor of its own product when it offers a lower-quality version of an outside seller's product. Because it has a revenue interest in the higher-

[11]When the private label has a higher quality than the seller's product it generates a higher margin. However, because it only diverts a share of seller's demand when the fee increases, it still cannot compensate for the decrease in pure marginalization.

quality product, it competes less fiercely against the seller, by choosing a price higher than if it were to monopolize the market and only offer the private label to consumers. Furthermore, the intermediary decreases the fee after introducing its private label. The price of the seller is more responsive to changes in the fee than the price of the private label, *i.e.*, $\frac{\partial p_1}{\partial f} > \frac{\partial p_2}{\partial f}$, which reflects both the advantage of the private label product only facing a weighted fee as an opportunity cost and not the whole fee as a direct cost, but also the more general property of lower-quality products having lower cost pass-through due to the lower diversion ratio.[12] Thus the fee reduction shifts back some of the demand to the seller's product from the private label. Because the price of the private label decreases as well, more consumers prefer to buy the private label to taking the outside option. Thus, when the intermediary decreases the fee it improves the consumer surplus further, besides the new variety that it offers and the competition it imposes on the seller's price.

Still, the introduction of the private label makes the outside seller worse off as it squeezes its profits. This leads us to the question of whether the outcome is any different from the well-known results from the literature on classic retailers. For this comparison, I next present the results for the wholesale model.

2.4.2 Wholesale Model

To put the results of the agency model into perspective, in this section I present the outcome of the introduction of a private label into the classic wholesale model.

Multi-product monopoly pricing downstream. In the pricing stage of the wholesale model, the intermediary acts as a multi-product monopolist and sets both prices p_1 and p_2. It takes the wholesale price w as given which enters its optimization problem as a marginal cost:

$$\max_{p_1,p_2} \ \pi_W = (p_1 - w)Q_1 + (p_2 - c_2)Q_2$$

If $\frac{s_1}{s_2}c_2 \leq w \leq \bar{\theta}(s_1 - s_2) + c_2$ there is an interior solution to the optimization problem where both products have positive demand:[13]

$$p_1^W - w = -\frac{Q_1}{dQ_1/dp_1} - \frac{dQ_2/dp_1}{dQ_1/dp_1}(p_2 - c_2) = \gamma(\hat{\theta})(s_1 - s_2) + (p_2 - c_2) = \frac{\sigma s_1 + w}{1 - \xi} \quad (2.8)$$

$$p_2^W - c_2 = -\frac{Q_2}{dQ_2/dp_2} - \frac{dQ_1/dp_2}{dQ_2/dp_2}(p_1 - w) = \gamma(\tilde{\theta}_2)s_2 = \frac{\sigma s_2 + c_2}{1 - \xi} \quad (2.9)$$

[12]If the marketplace operator enters with a higher-quality product, the pass-through of the seller's price will increase for the same reason as in the lower-quality private label case. The difference is that the private label's fee pass-through will now be higher than that of the seller's product. For more details see the Appendix 2.B.4.

[13]The higher-quality product has zero demand if $\bar{\theta} \leq \hat{\theta}$. This implies that $1 - G(\hat{\theta}) = 0$, thus $\gamma(\hat{\theta})(s_1 - s_2) = 0$. Then $\bar{\theta} \leq \hat{\theta} \Leftrightarrow \bar{\theta}(s_1 - s_2) \leq p_1 - p_2 = \gamma(\hat{\theta})(s_1 - s_2) + p_2 - c_2 + w - p_2 = w - c_2 \Leftrightarrow w \geq \bar{\theta}(s_1 - s_2) + c_2$. Similarly, the lower-quality product has zero demand if $\hat{\theta} \leq \tilde{\theta}_2 \Leftrightarrow p_1/s_1 \leq p_2/s_2 \Leftrightarrow w \leq \frac{s_1}{s_2}c_2 + s_1\gamma(\tilde{\theta}_2) - s_2\gamma(\tilde{\theta}_2) - \gamma(\hat{\theta})(s_1 - s_2) \leq \frac{s_1}{s_2}c_2$ with the last inequality following for $\bar{\theta} = \tilde{\theta}_2$.

The intermediary always internalizes the cannibalization effect for the lower-quality product and prices it as a stand-alone monopolist. This result has been pointed out by Itoh (1983) and generalized by Johnson and Myatt (2015). It is specific to vertically differentiated products, where the diversion ratio of the higher-quality product to the lower-quality product is equal to one, *i.e.*, $-\frac{dQ_2/dp_1}{dQ_1/dp_1} = 1$, as any demand lost by the higher-quality product is completely covered by the lower-quality product. Intuitively, if one regards the lower-quality as a base quality product and the higher-quality product as a composition of the base quality product plus an upgrade, a multi-product monopolist can always extract the base quality margin from all consumers who buy either product while charging an additional premium for the upgrade that the higher quality offers. As the base quality is independent of the premium, so is its price.[14]

The markup set for the third-party seller's product is at least the margin collected on the private label plus a premium for the higher quality it offers. The following Proposition summarizes the outcome for the price of the seller's product depending on the curvature of the Mills ratio $\gamma(\theta)$ for a general $g(\theta)$:

Proposition (Johnson and Myatt (2015))**.** *Depending on whether $\gamma''(\theta)$ is greater, equal, or less than zero, the price of the higher-quality product set by a multi-product monopolist is going to be higher, equal, or lower than the stand-alone monopolist price set by a stand-alone monopolist with the same cost.*

In other words, if $\gamma(\theta)$ is convex (concave) in θ then the upgrade of the outside seller's quality $(s_1 - s_2)$ acts as complementary (substitute) to the base quality s_2 and results in a higher (lower) price than what a stand-alone monopolist would charge. In the case of linear $\gamma(\theta)$ (constant curvature/pass-through demand) the base quality s_2 and the upgrade are independent of each other and their margin increments add up to the stand-alone monopoly margin, i.e., $p_1^W - w = \gamma(\tilde{\theta}_1)s_1$. Under Assumption (2.1) $\gamma' = \xi$ is constant. Hence, if the seller does not adjust the wholesale price post entry, the price p_1^W remains the same. However, demand is going to shrink, as the consumer who prefers the seller's product to the outside option is now buying the lower-quality product, i.e. $\hat{\theta} > \tilde{\theta}_1$. This demand-stealing effect is the main drive for the reduction of the wholesale price, which is discussed next.

Optimal wholesale price. After the introduction of the private label, the seller adjusts the wholesale price according to the changed demand and competition conditions:

$$\max_{w} \;\; \pi_S = (w - c_1)(1 - G(\hat{\theta}))$$

This means that the optimal margin it sets on top of its cost is given by:

[14] More technically, take the general formulas in (2.8) and (2.9) and solve for $p_2 - c_2$. If the total demand $Q = Q_1 + Q_2$ satisfies discrete choice and Slutsky symmetry ($\frac{dQ_1}{dp_2} = \frac{dQ_2}{dp_1}$), one can rewrite and obtain $p_2 - c_2 = -\frac{Q_1+Q_2}{dQ_2/dp_2 + dQ_1/dp_2} = -\frac{Q}{dQ/dp_2}$. Thus, the lowest-quality product offered by a multi-product monopolist is priced as if it were the only product in the market.

$$w - c_1 = \gamma(\hat{\theta})(s_1 - s_2)\left(\frac{\partial p_1}{\partial w}\right)^{-1} = \frac{\sigma(s_1 - s_2) + \xi(c_1 - c_2)}{1 - \xi}$$

where $\frac{\partial p_1}{\partial w}$ is the seller's price's pass-through with respect to the wholesale price after the introduction of the private label. Because the prices of both products downstream are independent for γ' constant, the pass-through remains the same as before, *i.e.*, $\frac{\partial p_1}{\partial w} = \frac{1}{1-\gamma'}$. Therefore the only drive for change in the wholesale price is the demand-stealing effect of the lower quality in the market. As a result, the seller reduces the wholesale price with the decrease being larger the better the quality of the private label. This directly reduces the price p_1^W. Finally, so long as the qualities offered are viable, the wholesale price is in the interval $[\frac{s_1}{s_2}c_2, \bar{\theta}(s_1 - s_2) + c_2]$.[15]

2.4.3 Comparison between the Agency and Wholesale models

While the intermediary in the wholesale model chooses the prices of both products as a multi-product monopolist, a marketplace operator competes directly in prices with the outside seller. As a result, in the wholesale model the intermediary can internalize the downstream margins and the pricing of both products, while in the agency model it only partially internalizes the lost fee revenues in its pricing, resulting in a relatively high price for the private label. Even though the intermediary decreases the fee, this one is still above the level in Lemma 2.2 for which the price of the private label in the agency model exceeds the price set by a stand-alone monopolist or by an intermediary in the wholesale model. By strategic complementarity, the price of the outside seller's product in the agency model is higher than in the wholesale model. This result is summarized in the following proposition:

Proposition 2.2. *Prices are higher in the agency model compared to the wholesale model after the introduction of a lower-quality private label.*

Overall, when the quality of the private label is $s_2 = 0$, buying the private label is equivalent to taking the zero outside option. In this case, $p_2 = 0$ and p_1 is the same as in the absence of a private label as in Equation (2.1). Once the quality of the private label improves, in both vertical models, the intermediary can charge a higher p_2. Due to the competitive pressure in both vertical models the consumer price p_1 decreases with s_2. In the limit where $s_2 = s_1$, the products are homogeneous and are both priced at the stand-alone monopoly price for quality s_1.

Corollary 2.1. *The seller has higher sales in the agency model than in the wholesale model, while the intermediary has higher sales in the wholesale model than in the agency model.*

[15] Recall that $\bar{\theta} = -\frac{\sigma}{\xi}$.

As the marketplace operator protects the revenue channel of the seller by increasing its price and decreasing the fee, the seller ends up having higher sales in the agency model compared to the wholesale model. To see this, consider the indifferent between the two products consumer $\hat{\theta} = \frac{p_1 - p_2}{s_1 - s_2}$. In the wholesale model this one is higher than in the agency model. The reason is that the relative difference in the lower-quality product prices $p_2^A - p_2^W$ is higher than the difference in the higher-quality product prices $p_1^A - p_1^W$, as seen in the proof of Proposition 2.2. In a similar way, the intermediary has higher sales of the private label in the wholesale model, as $\hat{\theta}^W$ is higher and $\tilde{\theta}_2^W = \frac{p_2^W}{s_2} < \frac{p_2^A}{s_2} = \tilde{\theta}_2^A$ is lower due to the lower price of the private label in the wholesale model. In total, more consumers buy a product in the wholesale model than in the agency model.

The introduction of a private label in an uncovered market always improves the consumer surplus. Consumers who bought the higher-quality product before and are still buying it after the introduction of the private label are better off due to its reduced price. Consumers, who switch to buying the cheaper product are better off as it generates higher utility for them. Finally, consumers who did not buy any product before and buy the private label after are better off as well. The only difference is that due to the higher prices and lower overall demand, the consumer surplus is improved less by the introduction of a private label in the agency model than in the wholesale model.

In both vertical models, the profit of the outside seller decreases as the quality of the private label s_2 improves. Inherent from the case without a private label, for a given s_2, in the agency model the seller extracts the smaller share of revenues generated through its channel, while the intermediary extracts the larger share of revenues, making the seller better off in the wholesale model.

Finally, there is a possibility for the intermediary to exclude the seller from the market in both vertical models. This happens when the private label is a "perfect copy" ($s_2 = s_1$) of the outside seller's product. In the agency model, the intermediary can increase the fee slightly by $\varepsilon > 0$, while in the wholesale model, it can increase slightly the price p_1^W by $\varepsilon > 0$, and capture the whole demand for itself. Whether such an outcome is profitable for the intermediary in both vertical models, is subject to the following Section 2.5.

2.5 Optimal Quality of the Private Label

Given that the intermediary can foreclose the outside seller when the two products are perfect substitutes ($s_1 = s_2$), in this section I endogenize the quality choice of the private label. As a longer-term decision, this choice is taken by the intermediary before the upstream contract stage, thus before the per-unit fee or wholesale price is determined. Because the main focus falls on the private label strategy of the intermediary, I do not model a quality choice by the outside seller. So far, I have assumed that the quality the seller offered to consumers was restricted only by the viability assumption. To keep the results tractable here I make three additional assumptions. First, I focus on the case where θ is uniformly distributed. Second, I assume that the seller is offering the optimal quality s_1^* conditional on being a monopolist in the market. It is easy to see that this quality is independent of which (or whether at all a) vertical model is implemented. For instance, the profits of the seller in the agency model, the wholesale model, or as a vertically integrated monopolist are respectively given by:

$$\pi_S^A = \frac{(\bar{\theta}s_1 - c_1(s_1))^2}{16\bar{\theta}s_1} \qquad \pi_S^W = \frac{(\bar{\theta}s_1 - c_1(s_1))^2}{8\bar{\theta}s_1} \qquad \pi_S^{Int} = \frac{(\bar{\theta}s_1 - c_1(s_1))^2}{4\bar{\theta}s_1}$$

In all three cases the optimal quality s_1^* is determined by the same optimality condition $c_1'(s_1^*)s_1^* = \frac{\bar{\theta}s_1^* + c_1(s_1^*)}{2}$. Thus, making my third assumption about the explicit form for the cost function $c(s) = \alpha s^2$, yields $s_1^* = \frac{\bar{\theta}}{3\alpha}$.

Conditional on the seller offering the optimal quality as a monopolist before the intermediary offers its private label, I find that:

Proposition 2.3. *The equilibrium qualities are such that* $s_2^{A*} < s_2^{W*} < s_1$.

In neither vertical model the intermediary wants to mimic the quality of the seller but rather prefers to differentiate its offer. In general, the choice of private label quality in both vertical models is influenced by two opposing effects. One is the reduced demand for the seller's product due to double marginalization, leading to competition for higher value consumers, pushing for a quality closer to s_1. The other is the incentive for maximum differentiation among heterogeneous consumers. The latter clearly dominates, leading to the intermediary offering a lower-quality product in both models.

In the agency model, the marketplace operator optimally chooses a lower quality for the private label than in the wholesale model. Inherent from the case without private label, the intermediary earns higher revenues from the seller's channel in the agency model compared to the wholesale model, resulting in a stronger incentive to differentiate

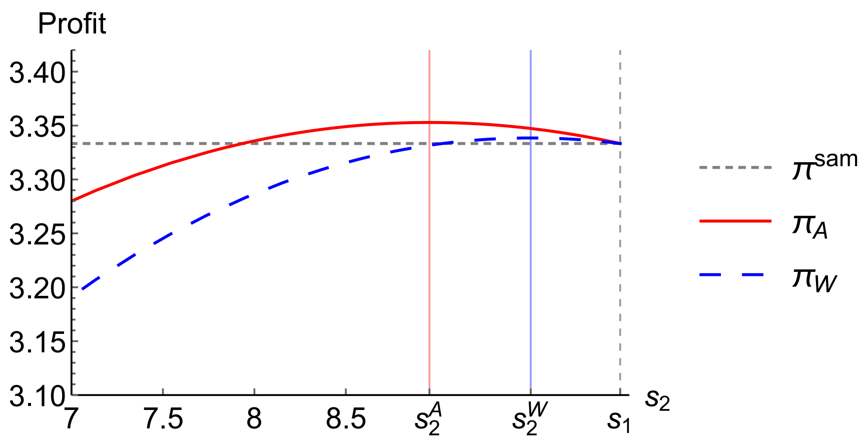

Figure 2.2: Profits of the intermediary in the agency (red, solid) and wholesale (blue, large dashed) model, and as a stand-alone monopolist with quality s_1 (grey, small dashed), as a function of quality s_2 with their respective maximum. Calibrated for $\bar{\theta} = 3$, $\alpha = 0.1$, $s_1 = \frac{\bar{\theta}}{3\alpha} = 10$.

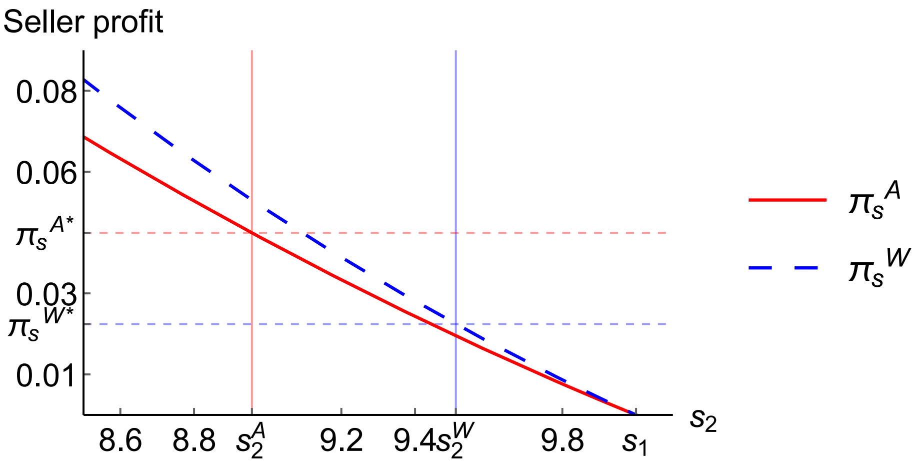

Figure 2.3: Profits of the seller in the agency (red, solid) and wholesale (blue, dashed) models and the optimal quality of the private label in each model. Calibrated for $\bar{\theta} = 3$, $\alpha = 0.1$, $s_1 = \frac{\bar{\theta}}{3\alpha} = 10$.

and offer a lower-quality product. This result is best illustrated by Figure 2.2 which plots the profits of the intermediary in both the agency and wholesale models and as a stand-alone monopolist, carrying the outside seller's product for a given level of the private label's quality.

As discussed in Section 2.4.3, the seller generates lower profits in the agency model than in the wholesale model for a given s_2. However, as in the agency model, the intermediary differentiates the private label more than in the wholesale model, and this outcome gets reversed. As illustrated in Figure 2.3, the seller actually ends up with more profits after the introduction of the private label in the agency than in the wholesale model. The difference in profits $\pi_S^A(s_2^{A*}) - \pi_S^W(s_2^{W*})$ increases in the heterogeneity of consumers (as $\bar{\theta}$ increases) and decreases in the cost factor α. In other words, in the markets in which

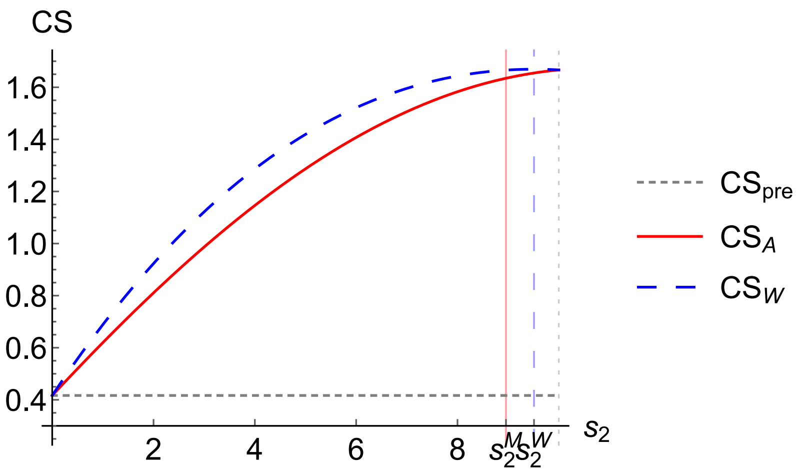

Figure 2.4: Consumer surplus pre- (gray, small dashed) and post-entry in the agency (red) and wholesale (blue, large dashed) models and the optimal quality of the private label. Calibrated for $\bar{\theta} = 3$, $\alpha = 0.1$, $s_1 = \frac{\bar{\theta}}{3\alpha} = 10$.

consumers have a greater heterogeneity in preference for quality, this difference is more prominent.

Due to the higher prices, the consumer surplus is lower in the agency model than in the wholesale model. After the quality choice, this result is further strengthened, as consumers enjoy a lower quality in the agency model (see Figure 2.4). The same holds true for total welfare. The differences increase in $\bar{\theta}$ and decrease in α.

Finally, a marketplace operator optimally chooses a lower quality for its product. This is consistent with the observations made by Marketplace Pulse (2019) about Amazon's products having, on average, lower ratings and reviews revealing the inferior quality of the products (Ellis and Hicken (2020)).

2.6 Extensions

In the following sections, I relax or alter some of the assumptions and discuss how they may affect the outcome in the agency and wholesale models. In Section 2.6.1, I allow for perfect contracts and outline their effect on the market outcome in both vertical models. In Section 2.6.2, I look at the case where the market power of the intermediary is limited while in 2.6.3, I look at a possible policy remedy that caps the ability of the marketplace operator to extract revenues through the fee.

2.6.1 Two-part Tariff

In the classic wholesale model two-part tariffs have been used to avoid the double-marginalization problem. In a scenario, where the seller sets a per-unit wholesale price w and a fixed tariff T_W, it is optimal to set $w = c_1$ and $T_W = \pi_W = (p_1 - c_1)Q_1$. Thus, the price p_1 is equal to the stand-alone monopoly price and the whole surplus goes to the seller and the intermediary obtains zero profits. Similarly, in the agency model, when there is no private label present, the marketplace operator can set the per-unit fee $f = 0$ and extract the whole surplus through a fixed tariff $T_A = (p_1 - c_1)Q_1$.

When the intermediary introduces a private label of lower quality in the wholesale model, the seller changes the optimal tariff to extract the duopoly surplus of the higher quality, while $w = c_1$. In such a case, the only channel through which the intermediary obtains positive profits is the one of the private label. Therefore, the best strategy is to perfectly copy the seller's product and foreclose the seller.

If a marketplace operator introduces a lower-quality private label, it has three instruments through which to maximize profits: the price of the private label p_2, the per-unit fee f, and the tariff T_A. As it has an interest in maximizing the joint producer surplus, its objective function changes in both the pricing and fee stages:

$$\max_{p_2} \pi_A = (p_1 - c_1)Q_1 + (p_2 - c_2)Q_2$$
$$\max_{f} \pi_A = (p_1(f) - c_1)Q_1(f) + (p_2(f) - c_2)Q_2(f)$$

The seller uses the same pricing rule as in Equation (2.3). The best response of the intermediary when choosing p_2 is also similar to the one in (2.4) with the opportunity cost of f replaced by the whole margin $p_1 - c_1$. Also in the optimal fee choice the fee is only represented indirectly by the cost it imposes on the seller. In particular, the first order condition of the intermediary's profit with respect to the fee is given by:

$$\frac{\partial \pi_A}{\partial f} = \frac{\partial p_1}{\partial f} Q_1 + (p_1 - c_1)\frac{dQ_1}{dp_1}\frac{\partial p_1}{\partial f} + (p_2 - c_2)\frac{dQ_2}{dp_1}\frac{\partial p_1}{\partial f} \stackrel{!}{=} 0$$

One can notice that the pass-through of the seller's price is not driving the choice of fee anymore as the intermediary can perfectly internalize the competition between the two products and choose $f = (p_2 - c_2)$. This, in turn, leads to the optimal prices in the agency model being determined by the same first-order conditions and final prices as in the wholesale model in Section 2.4.2 with $w = c_1$ in Equations (2.8) and (2.9).

The non-linear contract allows the marketplace operator to perfectly internalize competition between the products and act as a multi-product monopolist. The per-unit fee f increases with the quality of the private label in order to push the price of the seller's product up and guarantee that it extracts the monopoly margin of a higher-quality product from the market. The tariff $T_A = (p_1 - f - c_1)Q_1 = \frac{\sigma(s_1 - s_2) + \xi(c_1 - c_2)}{1 - \xi} Q_1$ decreases as the quality of the private label improves, to reflect the reduced revenues from the higher-quality product.

While consumer prices in both vertical models are the same, the incentives to choose the quality for the private label are different in the agency model. Because the marketplace operator obtains the revenues of a multi-product monopolist, it will choose a socially sub-optimal lower quality for the private label as demonstrated by Mussa and Rosen (1978), Itoh (1983), Maskin and Riley (1984), Besanko et al. (1987) and Tirole (1988). The seller will not be foreclosed and the consumer surplus will be higher.

Non-linear contracts avoid double marginalization which is why they are often referred to as "perfect" contracts as they allow the one dictating the contract to perfectly extract all rents from the market. In reality, markets are not perfect and double marginalization cannot be fully avoided as it is the case in Kang and Muir (2022). Still, the results in this section outline the drastic differences in the roles adopted by the intermediary in the different vertical models. Even though an intermediary prices like a multi-product monopolist in the wholesale model, it has an incentive to act as a direct competitor to the outside seller and fully forecloses the seller. Despite competing in prices with the seller in the agency model, by dictating the contract, a marketplace operator acts as a multi-product monopolist. It accommodates both products and protects the outside product more by differentiating and setting a sub-optimal quality for the private label.

2.6.2 Reduced Market Power of the Intermediary

In this section, I relax the assumption that the intermediary has complete power over the access to the consumer side of the market. Suppose that a share $\lambda \in [0,1]$ of consumers is loyal (or captive) to the intermediary, and only considers and buys products offered on its storefront. In order to reach them, the seller must use the services of the intermediary. The rest $(1-\lambda)$ of the consumers are shoppers and can also be reached directly by the seller, e.g., through its own shop/website. This implies that the seller has an additional pricing instrument - the price in the direct channel p_1^d. The timing of the game remains the same with the prices for consumers in both the direct and indirect/intermediary channels set after the wholesale price or the per-unit fee is determined.

In the absence of a private label, the results from Section 2.3 remain the same. As loyal consumers always buy from the intermediary, while others buy from the direct channel, in each channel there is monopoly pricing. In the intermediary's channel, there is double marginalization. The only difference is that as $\lambda \to 0$, the profits of the intermediary converge to zero, while the profits of the seller converge to the stand-alone monopoly profits.

When the intermediary introduces a lower-quality private label in either the agency or wholesale model, it decreases the price of the private label to be able to compete for shoppers. Thus, the price of the private label is a weighted average between the duopoly margin for loyal consumers in the indirect channel and the duopoly margin set for shoppers. Due to the strategic complementarity in the indirect channel and the opportunity cost the per-unit fee represents, the price of the private label in the agency model is higher than in the wholesale model. This leads to a higher price for the seller's product in both the direct and indirect channels.

Further, in the agency model, the per-unit fee decreases after the introduction of the private label. It does even more so given that the per-unit fee is independent of λ in the absence of the private label. Besides the decrease in pure marginalization, the margin that can be collected from the private label is lower than before, compensating even less for the pure marginalization decrease. Thus, the fee decreasing as $\lambda \to 0$ reflects the decreasing "gatekeeper" power of the intermediary.

Interestingly in the wholesale model, for a given quality of the private label, the wholesale price remains independent of the private label's price and is also independent of λ. Thus, the wholesale price decreases due to the same reason as before - decreased demand in the intermediary's channel. Finally, as λ decreases and the prices and profits converge to the ones in a market with two independent firms and no vertical structure in the market, the optimal quality chosen by the intermediary is going to decrease and converge to the one chosen by an independent competitor.

2.6.3 Regulating the Fee

In this section, I consider the case in which the market power of the intermediary is exogenously restricted by a regulator that puts a cap on the fees in order to leave more profits to third-party sellers. The timing changes to 1) the intermediary choosing the quality of the private label that it introduces and 2) the outside seller and the intermediary competing in prices.

First, any fee cap above the optimal fee chosen in the presence of the private label as in Equation (2.7) is non-binding. Once the cap falls below that level, it becomes binding and the intermediary will always choose to set the fee equal to the cap. If the fee is still above the threshold defined in Lemma 2.2, the price of the private label is still going to be higher than the stand-alone monopoly price or the price in the wholesale model. For any fee cap below the threshold defined in Lemma 2.2, the prices of both products are going to decrease with a stricter cap on the fee. If the regulator decides to completely forbid the intermediary from collecting any fees, the prices of the two products are going to be equal to the optimal prices set by the two independent firms.

For the case where the cap on the fee is not binding, the intermediary chooses the optimal quality of the private label as in Proposition 2.3. If the cap on the fee is binding, the intermediary will adjust the quality of the private label downwards. Choosing the same quality as the one of the seller's product is not optimal, as it results in lower profits than what a stand-alone monopolist can generate due to the lower fee. Therefore, the intermediary would choose to differentiate its offer from the offer of the seller. The lower the cap on the fee, the stronger the incentive to differentiate.

A policy that puts a cap on the fee possibly improves the profits of the outside seller and improves the consumer surplus as it decreases the double marginalization and intensifies competition between the two firms, leading to more differentiation. The profits of the intermediary are, however, harmed and lower than what it can generate by mimicking the seller's quality and excluding the rival from the market. In the model discussed in this paper, the cap restricts the amount of instruments that the intermediary can use in order to control competition on its platform. That's why it prefers to differentiate its offer. However, in reality, the intermediary may resort to other non-pricing instruments like, for example, decreasing the visibility of the seller's product in order to guarantee its product a larger share of the market. Thus, a cap on fees could be pro-competitive so long as it comes with a guarantee that other instruments cannot be used by platforms to exclude sellers.

2.7 Conclusion

In this paper, I have addressed the concern of the regulatory authorities about the conflict of interest arising from the introduction of a private label product by a marketplace operator. I show that a marketplace operator does not have an incentive to distort competition in favor of its own product. An intermediary in the agency model competes head-to-head in prices with an outside seller which drives the intermediary to internalize the fee revenue lost from competing with the seller in its own price. This leads to too much cost being passed to the consumer side by both the seller and the private label's pricing. To correct for this, the intermediary reduces the fee compared to what it would charge absent of the private label product. To further protect the revenue from the seller's channel, the intermediary differentiates its offer and offers a lower-quality product.

When it comes to an intermediary like Amazon targeting sellers in niche categories and copying their innovative products, I show that the protection of sellers is a greater concern in the wholesale model than it is in the agency model. The introduction of the new product always squeezes out the profits of the seller. Nevertheless, because in the agency model the intermediary sets a higher price for the private label and decreases the fee, this leaves the seller with more space to operate in the agency model than in the wholesale model. Furthermore, by differentiating its offer more, the intermediary in the agency model leaves the seller with more profits than in the wholesale model.

Finally, a private label version of a third-party product aimed at consumers with a lower willingness to pay always improves consumer surplus. Nevertheless, in the agency model, due to the higher prices and lower quality product compared to the wholesale model, the consumers' surplus is improved less.

The current analysis points to multiple possibilities for future research. The model at present is stylized and assumes that the market structure is exogenously given. The entry with a product by the platform always improves consumer surplus as the market is uncovered and no seller gets foreclosed. This contrasts with the results in my paper in relation to the papers that allow for free entry (Anderson and Bedre-Defolie (2023), Etro (2023b), Zennyo (2022)). Free entry in the current framework with vertically differentiated products would transform the current analysis into a vertical merger problem, rather than entry with new variety, potentially having different implications for the outcome of pricing and the fees charged to sellers (see Akgün et al. (2020)).

As I focus on the private label strategy of the intermediary, I keep the quality of the seller's product exogenous. Thus, I do not address the discussion about how entry by a platform affects the innovation incentives of third-party sellers. If the seller could adjust its quality, it would differentiate upwards, similar to independent sellers in markets with vertical differentiation. However, endogenizing the seller's quality in the current framework adds an additional layer of complexity without gaining significant insights into the private label strategy of the intermediary. Future research can address this issue.

2.A Omitted Proofs

Proof of Proposition 2.1. The closed-form solution for the optimal fee is given by

$$f = \frac{(\sigma s_1 + \xi c_1)(\left(\frac{\partial p_1}{\partial f}\right)^{-1}(1-\xi(1-\omega)) - \omega) + (\sigma s_2 + \xi c_2)(1-\xi(1-\omega) - \left(\frac{\partial p_1}{\partial f}\right)^{-1})}{(1-\xi)(1-\xi-\xi\left(\frac{\partial p_1}{\partial f}\right)^{-1})(1-\omega)} \tag{2.A.1}$$

Subtracting the original fee $f = \frac{\sigma+\xi c_1}{1-\xi}$ from the expression above yields

$$-\frac{\left(\sigma(s_1-s_2)+\xi(c_1-c_2)\right)\left(1-\xi(1-\omega)-\left(\frac{\partial p_1}{\partial f}\right)^{-1}\right)}{(1-\xi)\left(1-\xi-\xi\left(\frac{\partial p_1}{\partial f}\right)^{-1}\right)(1-\omega)}$$

The denominator is strictly positive for $\xi \leq 0$. The expression $\sigma(s_1-s_2)+\xi(c_1-c_2) > 0$ is equivalent to a basic assumption that the higher-quality product is viable and has positive demand when products are priced at marginal cost, *i.e.*, $\bar{\theta}s_1 - c_1 > \bar{\theta}s_2 - c_2$, where $\bar{\theta} = -\frac{\sigma}{\xi}$. Therefore, given $\frac{\partial p_1}{\partial f} = \frac{1-\xi\frac{\partial p_2}{\partial f}}{1-\xi}$ the fee decreases post entry if:

$$1-\xi(1-\omega)-\left(\frac{\partial p_1}{\partial f}\right)^{-1} > 0 \Leftrightarrow \frac{\partial p_2}{\partial f} \geq \frac{\omega}{1-\xi+\xi\omega} \tag{2.A.2}$$

The proof proceeds with demonstrating that this holds for $\xi \in [-1, 0]$.

Let us first explore the limits of the interval. For $\xi \to 0$ the distribution of preferences $g(\theta)$ converges to the exponential distribution. This family of distributions has the property that the Mills ratio is a constant, *i.e.*, $\gamma(\theta) = \sigma$. The expression of the optimal price from equation (2.4) reduces to:

$$p_2 - c_2 = \sigma s_2 - \omega\sigma s_1 + \omega f$$

Taking the first derivative w.r.t. the fee we obtain:

$$\frac{\partial p_2}{\partial f} = -\frac{\partial \omega}{\partial f}\sigma s_1 + \frac{\partial \omega}{\partial f} f + \omega$$

where $\omega = \frac{g(\hat{\theta})s_2}{g(\tilde{\theta}_2)(s_1-s_2)+g(\hat{\theta})s_2} \in [0,1]$. The price that the outside seller sets for the high-quality product is given by $p_1 - f - c_1 = \gamma(\hat{\theta})(s_1-s_2) = \sigma(s_1-s_2)$ implying that its pass-through is exactly equal to 1. The optimal fee is then given by $f = \sigma(s_1-s_2)+p_2-c_2 = \sigma s_1$. Hence, we obtain that $\frac{\partial p_2}{\partial f} = \omega \in [0,1]$. This also satisfies $\frac{\partial p_2}{\partial f} \geq \frac{\omega}{1-\xi+\xi\omega} = \omega$.

For the limit $\xi = -1$ the distribution of preferences $g(\theta)$ is given by the uniform distribution, where $\omega = \frac{s_2}{s_1}$ and $0 < \frac{\omega}{1-\xi+\xi\omega} = \frac{s_2}{3s_1-s_2} < \frac{\partial p_2}{\partial f} = \frac{3s_2}{4s_1-s_2} < 1$ for $s_2 < s_1$. This follows from the expression of the optimal price from equation (2.4) for uniform distribution:

$$p_2^*(f) = \frac{1}{4s_1 - s_2}\left(s_2\bar{\theta}(s_1 - s_2) + s_2c_1 + 2s_1c_2 + 3s_2f\right)$$

For $\xi \in (-1, 0)$ the proof proceeds as follows. Start with the first order condition of the marketplace operator w.r.t. the price of the private label and take its derivative w.r.t. the fee to obtain an expression containing $\frac{\partial p_2}{\partial f}$:

$$\frac{\partial \pi_A}{\partial p_2} \stackrel{!}{=} 0 \Leftrightarrow f\frac{g(\hat{\theta})}{(s_1 - s_2)} + G(\hat{\theta}) - G(\tilde{\theta}_2) = (p_2 - c_2)\left[\frac{g(\hat{\theta})}{(s_1 - s_2)} + \frac{g(\tilde{\theta}_2)}{s_2}\right] \quad \Big/\partial f \Rightarrow$$

$$\frac{g(\hat{\theta})}{(s_1 - s_2)} + f\frac{g(\hat{\theta})}{(s_1 - s_2)}\frac{-\xi - 1}{\gamma(\hat{\theta})(s_1 - s_2)}\left(\frac{\partial p_1}{\partial f} - \frac{\partial p_2}{\partial f}\right) +$$
$$+\left[\frac{g(\hat{\theta})}{(s_1 - s_2)}\left(\frac{\partial p_1}{\partial f} - \frac{\partial p_2}{\partial f}\right) - \frac{g(\tilde{\theta}_2)}{s_2}\frac{\partial p_2}{\partial f}\right] =$$
$$= \frac{\partial p_2}{\partial f}\left[\frac{g(\hat{\theta})}{(s_1 - s_2)} + \frac{g(\tilde{\theta}_2)}{s_2}\right] +$$
$$+(p_2 - c_2)\left[\frac{g(\hat{\theta})}{(s_1 - s_2)}\frac{-\xi - 1}{\gamma(\hat{\theta})(s_1 - s_2)}\left(\frac{\partial p_1}{\partial f} - \frac{\partial p_2}{\partial f}\right) + \frac{g(\tilde{\theta}_2)}{s_2}\frac{-\xi - 1}{\gamma(\tilde{\theta}_2)s_2}\frac{\partial p_2}{\partial f}\right]$$

Substitute $\frac{\partial p_1}{\partial f} - \frac{\partial p_2}{\partial f} = \frac{1 - \frac{\partial p_2}{\partial f}}{1 - \xi}$ and multiply by $(1 - \xi)$, then the expression can be simplifies to

$$\left[(2 - \xi) + (f - (p_2 - c_2))\frac{-\xi - 1}{\gamma(\hat{\theta})(s_1 - s_2)}\right]\frac{g(\hat{\theta})}{(s_1 - s_2)}\left(1 - \frac{\partial p_2}{\partial f}\right)$$
$$= \frac{\partial p_2}{\partial f}(1 - \xi)\frac{g(\tilde{\theta}_2)}{s_2}\left[2 + (p_2 - c_2)\frac{-\xi - 1}{\gamma(\tilde{\theta}_2)s_2}\right]$$

Take the expression for the optimal f from equations (2.7) and simplify equation (2.4) to $p_2 - c_2 = \frac{\omega}{1-\omega}\gamma(\hat{\theta})(s_1 - s_2)\left[\left(\frac{\partial p_1}{\partial f}\right)^{-1} - 1\right] + \gamma(\tilde{\theta}_2)s_2$. Then, after substituting f and $p_2 - c_2$ into the above expressions, this one simplifies to

$$\left[(2 - \xi) + \frac{(-\xi - 1)(1 - \xi)}{1 - \xi\frac{\partial p_2}{\partial f}}\right]\frac{g(\hat{\theta})}{(s_1 - s_2)}\left(1 - \frac{\partial p_2}{\partial f}\right) =$$
$$= \frac{\partial p_2}{\partial f}(1 - \xi)\frac{g(\tilde{\theta}_2)}{s_2}\left[2 + \frac{-\xi - 1}{\gamma(\tilde{\theta}_2)s_2}\left(\frac{\omega}{1 - \omega}\gamma(\hat{\theta})(s_1 - s_2)(-\xi)\left[\frac{1 - \frac{\partial p_2}{\partial f}}{1 - \xi\frac{\partial p_2}{\partial f}}\right] + \gamma(\tilde{\theta}_2)s_2\right)\right]$$

Notice that $\frac{\frac{g(\tilde{\theta}_2)}{s_2}}{\frac{g(\hat{\theta})}{(s_1 - s_2)}} = \frac{1-\omega}{\omega}$. Then multiply by $1 - \xi\frac{\partial p_2}{\partial f}$ and obtain a quadratic equation for $\frac{\partial p_2}{\partial f}$:

$$\left(\frac{\partial p_2}{\partial f}\right)^2 \underbrace{\xi\left[(2-\xi)+(1-\xi)\left((1-\xi)\frac{1-\omega}{\omega}+(1+\xi)\frac{\gamma(\hat{\theta})(s_1-s_2)}{\gamma(\tilde{\theta}_2)s_2}\right)\right]}_{a}+$$

$$+\frac{\partial p_2}{\partial f}\underbrace{\left[(-1-\xi)-(1-\xi)\left((1-\xi)\frac{1-\omega}{\omega}+\xi(1+\xi)\frac{\gamma(\hat{\theta})(s_1-s_2)}{\gamma(\tilde{\theta}_2)s_2}\right)\right]}_{b}]+ \tag{2.A.3}$$

$$+\underbrace{(1-\xi+\xi^2)}_{c}\stackrel{!}{=}0$$

The solution to the quadratic equation is given by $\frac{-b\pm\sqrt{b^2-4ac}}{2a}$. The denominator $2a$ is always negative for $\xi\in(-1,0)$. $-b$ is always positive, although not apparent straight from the expression. While $(1-\xi)\geq 0$ for $\xi\leq 0$ and $(-1-\xi)\leq 0$ for $\xi\geq -1$, the expression $\left((1-\xi)\frac{1-\omega}{\omega}+\xi(1+\xi)\frac{\gamma(\hat{\theta})(s_1-s_2)}{\gamma(\tilde{\theta}_2)s_2}\right)$ is not that obvious. Whether the latter is greater than zero is equivalent to:

$$\begin{aligned}(1-\xi)(1-\omega)\gamma(\tilde{\theta}_2)s_2 &> -\xi(1+\xi)\omega\gamma(\hat{\theta})(s_1-s_2)\\ \Leftrightarrow (1-\xi)(1-G(\tilde{\theta}_2)) &> -\xi(1+\xi)(1-G(\hat{\theta}))\end{aligned}$$

$(1-\xi)>-\xi(1+\xi)$ for all ξ and the total demand $(1-G(\tilde{\theta}_2))$ is greater than the demand for the higher-quality product $(1-G(\hat{\theta}))$. Thus, $-b$ is always positive.[16]

To show that $\frac{\partial p_2}{\partial f}\geq\frac{\omega}{1-\xi+\xi\omega}$, we can plug in $\left(\frac{\partial p_2}{\partial f}+\frac{\omega}{1-\xi+\xi\omega}\right)$ instead of $\frac{\partial p_2}{\partial f}$ in (2.A.3) and show that the function still has a positive root. Denote $x=\frac{\omega}{1-\xi+\xi\omega}$ for simplicity. Then we obtain a new quadratic equation:

$$\left(\frac{\partial p_2}{\partial f}\right)^2 a+(b+2ax)\frac{\partial p_2}{\partial f}+c+ax^2+bx=\left(\frac{\partial p_2}{\partial f}\right)^2 a_1+b_1\frac{\partial p_2}{\partial f}+c_1$$

We are interested in the root $\frac{-b_1-\sqrt{b_1^2-4a_1c_1}}{2a_1}>0$. As $a_1=a<0$ and $-b_1=-b-2ax>0$, we need $-b_1-\sqrt{b_1^2-4a_1c_1}<0\Leftrightarrow 0<-4a_1c_1$. This one is positive if $c_1>0$, or:

$$-\frac{\xi(1-\xi)^2(1-\omega)\left[\frac{\gamma(\hat{\theta})(s_1-s_2)}{\gamma(\tilde{\theta}_2)s_2}(1+\xi)\omega+\omega-\xi(1-\omega)\right]}{(1-\xi(1-\omega))^2}>0$$

For $\xi\in[-1,0]$ the above expression is positive, thus $\frac{\partial p_2}{\partial f}\geq\frac{\omega}{1-\xi+\xi\omega}$ and the fee decreases post entry.

Proof of Proposition 2.2. The price p_1 for the outside seller's product in the agency model is higher than the one set in the wholesale model as long as the difference

[16] One can check that $\frac{\partial p_2}{\partial f}\leq 1$ is satisfied as long as the root $\frac{-b\pm\sqrt{b^2-4ac}}{2a}\leq 1\Leftrightarrow a+b+c\leq 0$.

$$p_1^A - p_1^W = \frac{((1-\xi\omega)\left(\frac{\partial p_1^A}{\partial f}\right)^{-1} - (1-\xi))(\sigma(s_1-s_2)+\xi(c_1-c_2))}{(1-\xi)^2(1-\omega)(1-\xi-\xi\left(\frac{\partial p_1^A}{\partial f}\right)^{-1})}$$

is positive. This happens when $\frac{\partial p_2}{\partial f} \leq \omega$. From equation (2.A.3) $\frac{-b-\sqrt{b^2-4ac}}{2a} \leq \omega \Leftrightarrow -c > (a+b)\omega$. This reduces to $-c(1-\omega) > (a+b-c)\omega \Leftrightarrow 1-\xi+\xi^2 < (1-\xi)^3$. As $\xi \leq 0$ then $1-\xi+\xi^2 < (1-\xi)^2$, thus $\frac{\partial p_2}{\partial f} \leq \omega$ and the price of the higher-quality product is higher in the agency model than in the wholesale model.

From Lemma 2.2 it follows that the price of the private label is higher than the one set by a stand-alone monopolist if $f > \frac{\sigma s_1+\xi c_1}{1-\xi} + \frac{\xi}{1-\xi}\gamma(\hat{\theta})(s_1-s_2)$. Thus, one needs to show that the optimal fee in (2.7) is higher than the threshold in Lemma 2.2. Subtracting the latter from the former yields:

$$\frac{\left(\left(\frac{\partial p_1}{\partial f}\right)^{-1} - 1\right)(\sigma(s_1-s_2)+\xi(c_1-c_2))(1-2\xi+\xi^2(1-\omega))}{(1-\xi)(1-2\xi)(1-\omega)\left(1-\xi-\xi\left(\frac{\partial p_1}{\partial f}\right)^{-1}\right)} \geq 0$$

For efficient qualities and $\xi \leq 0$ one only needs to show that $\frac{\partial p_1}{\partial f} \leq 1 \Leftrightarrow \frac{\partial p_2}{\partial f} \leq 1$. As $\frac{\partial p_2}{\partial f} \leq \omega$, it follows that the private label is priced above monopoly level in the agency model.

Alternatively, the difference between the prices for the private label in the two vertical models is:

$$p_2^A - p_2^W = \frac{\omega(\left(\frac{\partial p_1^A}{\partial f}\right)^{-1} - 1)(\sigma(s_1-s_2)+\xi(c_1-c_2))}{(1-\xi)(1-\omega)(1-\xi-\xi\left(\frac{\partial p_1^A}{\partial f}\right)^{-1})}$$

This difference is higher than $p_1^A - p_1^W$ as long as $\frac{\partial p_2}{\partial f} \geq 0$, which is satisfied as per proof of Proposition 2.1.

Proof of Proposition 2.3. The profit of the intermediary in both vertical models is equal to the integrated monopolist profit for $s_2 = s_1$. To show that $s_2^{A*} < s_2^{W*} < s_1$ holds in equilibrium, it must be that the profit of the intermediary in the agency model is higher (I), concave in s_2 (II), and has a negative slope at both s_2^{W*} (III) and s_1 (IV). For $s_2^{W*} < s_1$ we need to show that the profit of the intermediary in the wholesale model is concave (V) and has a negative slope at s_1 (VI). Statement (I) holds for all $s_2 < s_1$ inherent from the fact that in the agency model the intermediary extracts more profit from the outside seller. For statements (IV) and (VI), take the first-derivatives of the intermediary's profits in the agency and the wholesale model with respect to s_2 and evaluate at $s_1 = \frac{\bar{\theta}}{3\alpha}$:

$$\left.\frac{\partial \pi_A}{\partial s_2}\right|_{s_2=s_1=\frac{\bar{\theta}}{3\alpha}} = -\frac{\bar{\theta}}{81} < \left.\frac{\partial \pi_W}{\partial s_2}\right|_{s_2=s_1=\frac{\bar{\theta}}{3\alpha}} = -\frac{\bar{\theta}}{144}$$

Thus, statements (IV) and (VI) are fulfilled.

The optimal quality choice of the intermediary in the wholesale model s_2^{R*} is given by:

$$s_2(s_1) = \frac{6\bar{\theta} + \alpha s_1 - \sqrt{9\bar{\theta}^2 + 12\bar{\theta}\alpha s_1 - 8\alpha^2 s_1^2}}{9\alpha} = \frac{\bar{\theta}(19 - \sqrt{109})}{27\alpha}$$

Substituting into the first derivative of the intermediary's profit in the agency model proves statement (III):

$$\left.\frac{\partial \pi_A}{\partial s_2}\right|_{s_2 = s_2^{R*}} = -\frac{\bar{\theta}(56303 - 3383\sqrt{109})}{486(91 - \sqrt{109})^2} \approx -\frac{\bar{\theta}}{150.3} < 0$$

Statement (V) requires that the second derivative of the intermediary's profit in the wholesale model is negative:

$$\frac{\partial^2 \pi_W}{\partial s_2^2} = \frac{\alpha(27\alpha s_2 - 10\bar{\theta})}{24\bar{\theta}} \leq 0 \Leftrightarrow s_2 < \frac{19\bar{\theta}}{27\alpha}$$

which is true for all $s_2 < s_1 = \frac{\bar{\theta}}{3\alpha}$. The second derivative of the profit function in the agency model is given by:

$$\frac{\partial^2 \pi_A}{\partial s_2^2} = \frac{\alpha(243\alpha^4 s_2^4 + 1674\alpha^3 s2^3\bar{\theta} + 3024\alpha^2 s_2^2\bar{\theta}^2 - 1152\alpha s_2\bar{\theta}^3 - 1520\bar{\theta}^4)}{6\bar{\theta}(8\bar{\theta} + 3\alpha s_2)^3}$$

Let $k = \frac{\alpha s_2}{\bar{\theta}}$. Then the expression above is negative if:

$$243k^4 + 1674k^3 + 3024k^2 - 1152k - 1520 < 0$$

For $s_2 \leq \frac{\bar{\theta}}{3\alpha}$ it means that $k \leq \frac{1}{3}$ and the expression above can be, at most, equal to $-1503 < 0$. This proves statement (II).

2.B Extensions and Supplementary Materials

2.B.1 Reduced Market Power of the Intermediary

Agency vs. wholesale model in the absence of a private label. The objective functions in the two vertical models in the absence of a private label can be described by the following equations for the agency model:

$$\text{US} \ \max_{f} \ \pi_A = \lambda f(1 - G(\tilde{\theta}_1^i))$$

$$\text{DS} \ \max_{p_1^i, p_1^d} \ \pi_S = \lambda(p_1^i - f - c_1)(1 - G(\tilde{\theta}_1^i)) + (1 - \lambda)(p_1^d - c_1)(1 - G(\tilde{\theta}_1^d))$$

and for wholesale model:

$$\text{US} \ \max_{w} \ \pi_S = \lambda(w - c_1)(1 - G(\tilde{\theta}_1^i)) + (1 - \lambda)(p_1^d - c_1)(1 - G(\tilde{\theta}_1^d))$$

$$\text{DS} \ \max_{p_1^i} \ \pi_R = \lambda(p_1^i - w)(1 - G(\tilde{\theta}_1^i)) \ \max_{p_1^d} \ \pi_S = (1 - \lambda)(p_1^d - c_1)(1 - G(\tilde{\theta}_1^d))$$

where the superscript i stands for the intermediary's channel, and d for the direct channel.

The prices in the intermediary's channel p_1^i, the optimal fee f and the the wholesale price w remain the same as before. The intuition is simple. The consumer prices in the direct and indirect channel are independent from each other - they are the monopoly price given the cost of $f + c_1/w$ in the indirect, and c_1 in the direct channel. As shoppers will always go to the cheaper channel and loyal consumers will buy through the intermediary, the firm setting the upstream contract (f/w) focuses only on the demand of the captives in this channel. The profits generated through the intermediary channel are directly proportionate to the profits in Section 2.3. The profits of the outside seller are higher, as they converge to stand-alone monopolist profits as $\lambda \to 0$.

Private label in the agency model. When the intermediary introduces a private label, consumers from both channels are going to consider it as an option. Like in the case without a private label discussed above, shoppers will compare the offer of the private label with the offer of the seller in the direct channel, as this one will be always cheaper due to lack of double marginalization. Thus, in the direct channel there is an indifferent consumer $\hat{\theta}^d = \frac{p_1^d - p_2}{s_1 - s_2}$. Similarly, the captive consumers are only going to compare offers in the intermediary channel, i.e., $\hat{\theta}^i = \frac{p_1^i - p_2}{s_1 - s_2}$. In the pricing stage we have an additional price, which the seller chooses - the price in the direct channel:

$$\max_{p_1^i, p_1^d} \pi_S = \lambda(p_1^i - f - c_1)(1 - G(\hat{\theta}^i)) + (1 - \lambda)(p_1^d - c_1)(1 - G(\hat{\theta}^d))$$

$$\max_{p_2} \pi_A = \lambda f(1 - G(\hat{\theta}^i)) + (p_2 - c_2)[\lambda(G(\hat{\theta}^i) - G(\tilde{\theta}_2)) + (1 - \lambda)(G(\hat{\theta}^d) - G(\tilde{\theta}_2))]$$

In both channels the seller sets a duopoly markup, and the prices differ only through the presence of double marginalization in the intermediary channel: $p_1^d - c_1 = \gamma(\hat{\theta}^d)(s_1 - s_2)$ and $p_1^i - c_1 = \gamma(\hat{\theta}^i)(s_1 - s_2) + f$.

The price of the private label is then the duopoly markup of the weighted demand plus the opportunity cost from the fee in the intermediary channel:

$$p_2 - c_2 = \frac{\lambda G(\hat{\theta}^i) + (1 - \lambda)G(\hat{\theta}^d) - G(\tilde{\theta}_2)}{\lambda \frac{g(\hat{\theta}^i)}{(s_1 - s_2)} + (1 - \lambda)\frac{g(\hat{\theta}^d)}{(s_1 - s_2)} + \frac{g(\tilde{\theta}_2)}{s_2}} + \frac{\frac{g(\hat{\theta}^i)}{(s_1 - s_2)}}{\lambda \frac{g(\hat{\theta}^i)}{(s_1 - s_2)} + (1 - \lambda)\frac{g(\hat{\theta}^d)}{(s_1 - s_2)} + \frac{g(\tilde{\theta}_2)}{s_2}} f$$

Hence, as λ decreases, so will the price of the private label, and as $\lambda \to 0$, the price will be equal to the price set by an independent duopolist carrying the lower quality in a market.

The optimal fee is the solution to

$$\lambda f = \lambda \left(\gamma(\hat{\theta}^i)(s_1 - s_2) \left(\frac{\partial p_1^i}{\partial f} \right)^{-1} + (p_2 - c_2) \right) + (1 - \lambda)(p_2 - c_2) \frac{g(\hat{\theta}^d)}{g(\hat{\theta}^i)} \frac{\partial p_1^d}{\partial f} \left(\frac{\partial p_1^i}{\partial f} \right)^{-1}$$

As before, the optimal fee consists of the pure marginalization on top of the seller‘s product plus the margin lost on the private label in the own channel, plus an additional term. This additional term represents the relative margin of the private label lost from shoppers.

The fee decreases as $\lambda \to 0$. We can see this for the uniform case where

$$f(\lambda) = \frac{[8\bar{\theta}(s_1^2 + s_2^2) - s_2 c_2 - 8c_1 s_1](s_1 - s_2)}{16s_1^2 - 2(4+3\lambda)s_1 s_2 + (1-3\lambda)s_2^2} < f(\lambda = 1) = \frac{8\bar{\theta}(s_1^2 + s_2^2) - s_2 c_2 - 8c_1 s_1}{2((s_1 + s_2)}$$

Private label in the wholesale model. In the pricing stage we have the following objective functions:

$$\begin{aligned} &\underset{p_1^d}{max}\ \pi_S = \lambda(w - c_1)(1 - G(\hat{\theta}^i)) + (1-\lambda)(p_1^d - c_1)(1 - G(\hat{\theta}^d)) \\ &\underset{p_1^i, p_2}{max}\ \pi_R = \lambda(p_1^i - w)(1 - G(\hat{\theta}^i)) + (p_2 - c_2)[\lambda(G(\hat{\theta}^i) - G(\tilde{\theta}_2)) + (1-\lambda)(G(\hat{\theta}^d) - G(\tilde{\theta}_2))] \end{aligned}$$

In the direct channel, the seller sets a duopolistic markup $p_1^d - c_1 = \gamma(\hat{\theta}^d)(s_1 - s_2)$.

In the intermediary channel the price of the outside seller's product is set as before $p_1^i - w = \gamma(\hat{\theta}^i)(s_1 - s_2) + (p_2 - c_2)$. The price of the private label, however, is pushed down to account for the competition for shoppers, and is a weighted average between the monopoly and duopoly prices:

$$\begin{aligned} p_2 - c_2 = {} & \lambda \frac{g(\tilde{\theta}_2)/s_2}{g(\tilde{\theta}_2)/s_2 + (1-\lambda)g(\hat{\theta})/(s_1 - s_2)} \gamma(\tilde{\theta}_2)s_2 \\ & + (1-\lambda)\frac{g(\tilde{\theta}_2)/s_2 + g(\hat{\theta})/(s_1 - s_2)}{g(\tilde{\theta}_2)/s_2 + (1-\lambda)g(\hat{\theta})/(s_1 - s_2)}((1-\omega^d)\gamma(\tilde{\theta}_2)s_2 - \omega^d \gamma(\hat{\theta}^d)(s_1 - s_2)) \end{aligned}$$

where $\omega^d = \frac{g(\hat{\theta^d})s_2}{g(\hat{\theta^d})s_2 + g(\hat{\theta}_2)(s_1 - s_2)}$. As $\lambda \to 1$ the prices in the wholesale model are going to converge to the monopoly prices derived in Section 2.4.2, while as $\lambda \to 0$ the prices in the market are going to converge to the pure duopoly level.

If the seller does not adjust w, for $\lambda < 1$ the price p_1^i decreases after the introduction of the private label, in contrast to the case when $\lambda = 1$. This is a direct result of the lower price of the private label. Nevertheless, the demand in the intermediary channel for the seller's product is going to be constant for any $\lambda \in [0, 1]$.[17]

When choosing the optimal wholesale price the seller's objective is

$$\underset{w}{max}\ \pi_S = \lambda(w - c_1)(1 - G(\hat{\theta}^i)) + (1-\lambda)(p_1^d - c_1)(1 - G(\hat{\theta}^d))$$

which gives the F.O.C.:

$$\frac{\partial \pi_s}{\partial w} = \frac{d\pi_s}{dp_1^i}\frac{\partial p_1^i}{\partial w} + \frac{d\pi_s}{dp_2}\frac{\partial p_2}{\partial w} + \frac{d\pi_s}{dp_1^d}\frac{\partial p_1^d}{\partial w} + \frac{d\pi_s}{dw} \Leftrightarrow w - c_1 = \gamma(\hat{\theta}^i)(s_1 - s_2)(1 - \gamma'(\hat{\theta}^i))$$

[17] The reason is that $\hat{\theta} = \frac{p_1 - p_2}{(s_1 - s_2)} = \frac{\gamma(\hat{\theta})(s_1 - s_2) + p_2 - c_2 + c_1 + w - p_2}{(s_1 - s_2)} = \gamma(\hat{\theta}) - c_2 + c_1 + w$. As the costs and the Mills ratio $\gamma(\cdot)$ are exogenous and independent of λ, for given w so will be demand.

As the intermediary internalizes competition effects between the private label and the seller's product in its channel, the price of the private label is independent of the price p_1^i and the wholesale price w, i.e., $\frac{\partial p_2}{\partial w} = 0$. The wholesale price is the same and decreases for the same reason and at the same rate for all $\lambda \in [0, 1]$[18] - demand-stealing effect from the private label.

Comparison of the two models. The price of the private label in the agency model is still higher than in the wholesale model:

$$\begin{aligned}
p_2^A = & \frac{1}{16s_1^2 - 2(4+3\lambda)s_1 s_2 + (1-3\lambda)s_2^2}(\bar{\theta}s_2((4+6\lambda)s_1 - s_2)(s_1 - s_2)+ \\
& + c_2(8s_1^2 - (2+3\lambda)s_1 s_2 - 3\lambda s_2^2) + c_1 s_2(s_1(4-6\lambda) - s_2(1-3\lambda)) \\
& > p_2^W = \frac{\bar{\theta}s_2(s_1 - s_2)(1+\lambda) + 2c_2(s_1 - \lambda s_2) + (1-\lambda)c_1 s_2}{4s_1 - s_2(1+3\lambda)}
\end{aligned}$$

Showing that $p_2^A - p_2^W > 0$ reduces to showing that:

$$\bar{\theta}(8s_1^2 - 4s_1 s_2(3\lambda+1) + (3\lambda+5)s_2^2) - 4c_1(2s_1 + s_2(1-3\lambda)) + c_2(8s_1 - (3\lambda+5)s_2) > 0$$

The expression is linear and decreasing in λ, and is strictly positive at $\lambda = 1$, thus $p_2^A > p_2^W$. As a direct result of this, it follows that the price $p_1^{Ad} > p_1^{Wd}$ as both have the same best-response function to the price of the private label, namely $p_1^j = \frac{\bar{\theta}(s_1 - s_2) + p_2^j + c_1}{2}$ for $j = Ad, Wd$.

For the seller's price in the intermediary's channel we can show that $p_1^{Ai} > p_1^{Wi}$:

$$\begin{aligned}
p_1^{Ai} = & \frac{1}{32s_1^2 - 4(4+3\lambda)s_1 s_2 + 2(1-3\lambda)s_2^2}(\bar{\theta}(s_1 - s_2)(24s_1^2 - 4s_1 s_2 + (1-3\lambda)s_2^2)+ \\
& + 4c_1 s_1(2s_1 + s_2(1-3\lambda)) + c_2(8s_1^2 - 3(1+\lambda)s_1 s_2 + (1-3\lambda)s_2^2)) > \\
p_1^{Wi} = & \frac{1}{4(4s_1 - s_2(1+3\lambda))}(\bar{\theta}(s_1 - s_2)(12s_1 + s_2(1-5\lambda))+ \\
& + c_1(4s_1 + (3-7\lambda)s_2) + c_2(4s_1 + s_2(1-5\lambda))
\end{aligned}$$

This simplifies to showing that

$$\begin{aligned}
& \bar{\theta}(s_1 - s_2)(8\lambda s_1^2 + 2(6+\lambda(1-15\lambda))s_1 s_2 - (3+\lambda)(1-3\lambda)s_2^2) - \\
& - c_1(8\lambda s_1^2 - 2(6-\lambda(13-15\lambda))s_1 s_2 + (1-3\lambda)(3-7\lambda)s_2^2) - \\
& - c_2(8(3-4\lambda)s_1^2 - 2(9-\lambda(11+6\lambda))s_1 s_2 + (3+\lambda)(1-3\lambda)s_2^2) > 0
\end{aligned}$$

If one expresses the condition above as $A\bar{\theta}(s_1 - s_2) - Bc_1 + Cc_2$ it is straightforward to show that $A > 0$, $A > B$, $A > C$ and $s_1(A - B) = s_2(A - C)$ so that the expression is then positive if the high quality is efficient, i.e. $\bar{\theta}s_1 - c_1 > \bar{\theta}s_2 - c_2$.

[18] The problem from the previous footnote only changes to $\hat{\theta} = \gamma(\hat{\theta})(2 - \gamma'(\hat{\theta})) - c_2 + c_1$.

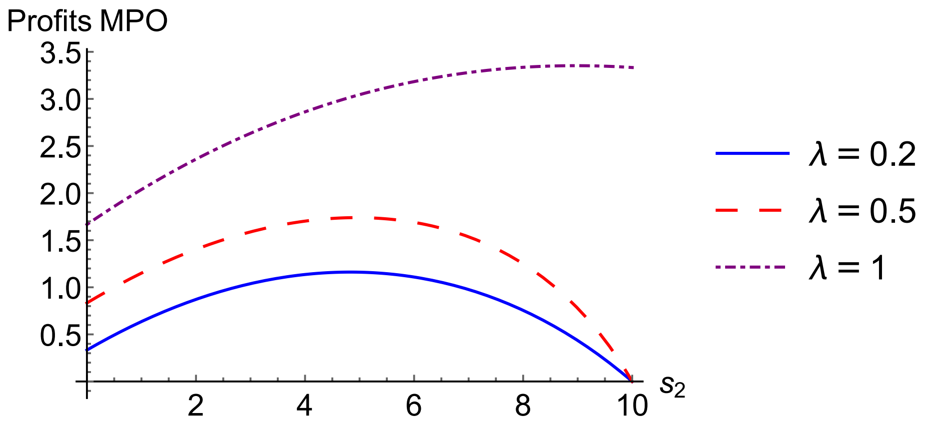

Figure 2.B.1: Profit function of the intermediary in the agency model as a function of the quality of the private label, for different values of λ. Calibrated for uniformly distributed preferences, $\bar{\theta} = 3$, $\alpha = 0.1$, $s_1 = \frac{\bar{\theta}}{3\alpha} = 10$.

Finally, given that in the baseline model ($\lambda = 1$) deriving the optimal quality of the private label in the agency model required numerical methods, for $\lambda \neq 1$ the problem becomes even more challenging. Still, comparing the profit curves of the intermediary in the agency model in Figure 2.B.1, it is not surprising that for lower λ the maximum of the profit curve is reached for lower s_2, thus converging to the s_2 chosen by an independent competitor in a market with no vertical structures.

2.B.2 Regulating the Fee

Here I demonstrate that the optimal quality chosen by the marketplace operator decreases when the fee cap falls below the optimal fee in (2.7) and decreases further. From the Proof of Proposition 2.3, we know that the profit of the intermediary in the agency model is concave. Therefore we only need to show that the derivative of the intermediary's profit evaluated at the seller' quality is negative and decreases as f^c decreases.

$$\left.\frac{\partial \pi_A}{\partial s_2}\right|_{s2=s_1=\frac{\bar{\theta}}{3\alpha}} = -\frac{\bar{\theta}}{9} + \frac{5\alpha}{3\bar{\theta}} f - \frac{7\alpha^2}{\bar{\theta}^2} f^2$$

At $s_2 = s_1$ the optimal fee without a cap from equation (2.7) is equal to $f = \frac{t^2}{9\alpha}$ and is lower for any $s_2 < s_1$ as per Proposition 2.1. Thus, when there is a cap on the fee, it is going to be in the interval $[0, \frac{t^2}{9\alpha}]$. At the lower bound it is apparent that the $\left.\frac{\partial \pi_A}{\partial s_2}\right|_{s2=s_1=\frac{\bar{\theta}}{3\alpha}} = -\frac{\bar{\theta}}{9} < 0$. At the upper bound $\left.\frac{\partial \pi_A}{\partial s_2}\right|_{s2=s_1=\frac{\bar{\theta}}{3\alpha}} = -\frac{\bar{\theta}}{81} < 0$, which is the same as in the proof of Proposition 2.3. Thus, as the cap on the fee is set lower, the derivative evaluated at the seller's quality becomes "more negative", emphasizing that the optimal s_2 moves away from s_1. Finally, for $f^c \in [0, \frac{t^2}{9\alpha}]$ $\left.\frac{\partial \pi_A}{\partial s_2}\right|_{s2=s_1}$ is monotone, as the turning point occurs at $f = \frac{5t^2}{42\alpha} > \frac{t^2}{9\alpha}$.

2.B.3 Ad-valorem Fee

In reality, marketplaces like Amazon have a complex menu of fees and charge an ad-valorem fee on the revenues from each sale. The differences between per-unit and ad-valorem fees in the absence of a private label are discussed by Johnson (2017). It is a well-known result in the literature on taxation that ad-valorem taxes are more efficient than specific taxes (Bishop (1968)). Thus, in the absence of a private label, the price is lower with an ad-valorem fee than with per-unit fee $p_1^*(r) < p_1^*(f)$. This leads to more demand and higher industry profits. The seller receives a lower share of these profits, and as a result less profits under ad-valorem than under per-unit fee contract. The marketplace operator and consumers, however, are better off with an ad-valorem fee contract. In this section I assume a uniform distribution of preferences and a cost function given by $c(s) = \alpha s^2$.

Let us first briefly go over the case without private label. The profit of the seller is given by

$$\pi_S = ((1-r)p_1 - c_1)(1 - G(\tilde{\theta}_1))$$

and yields an optimal monopolist margin of $p_1 - \frac{c_1}{1-r} = \gamma(\tilde{\theta}_1)s_1$, where the ad-valorem fee r collected on the revenues enters the pricing decision as a cost-augmenting factor. For the uniformly distributed θ the price is then $p_1 = \frac{1}{2}(\bar{\theta}s_1 + \frac{c_1}{1-r})$.

The marketplace operator then sets r so that it maximizes $\pi_A = rp_1Q_1 = rp_1(1 - G(\tilde{\theta}_1))$. The optimal ad-valorem fee r is the solution of the following equation:

$$\begin{aligned}\frac{r}{1-r} &= -\frac{Q_1}{dQ_1/dp_1}\frac{p_1}{c_1}\left(\frac{\partial p_1}{\partial r}\right)^{-1} = \frac{p_1\gamma(\tilde{\theta}_1)s_1}{c_1}\left(\frac{\partial p_1}{\partial r}\right)^{-1} \\ &= \frac{p_1(\bar{\theta}s_1 - p_1)}{(2p_1 - \bar{\theta}s_1)^2} = \frac{(\bar{\theta}s_1)^2 - (\frac{c_1}{1-r})^2}{2(\frac{c_1}{1-r})^2} \Leftrightarrow \frac{1+r}{(1-r)^3} = \left(\frac{\bar{\theta}s_1}{c_1}\right)^2\end{aligned}$$

In general, like the per-unit fee, the ad-valorem fee decreases in the cost c_1. In the uniform case the optimal ad-valorem fee is the solution to $\frac{1+r}{(1-r)^3} = \left(\frac{\bar{\theta}s_1}{c_1}\right)^2$.

Consistent with the result, derived by Johnson (2017), the marketplace operator has higher revenues with ad-valorem fee than with per-unit fee. One interesting result is that the revenue per unit with ad-valorem fee $rp_1 > f$ is higher than with the per-unit fee. Under Assumption (2.1) the pre-entry per-unit fee f is equal to the stand-alone monopolist margin. This implies that with an ad-valorem fee, the marketplace operator extracts an even higher revenue per unit than a stand-alone monopolist.

After the introduction of a lower-quality private label the profit functions of the seller and the marketplace operator change to:

$$\begin{aligned}\pi_S &= ((1-r)p_1 - c_1)Q_1(p_1, p_2) = ((1-r)p_1 - c_1)(1 - G(\hat{\theta})) \\ \pi_A &= r\,p_1Q_1(p_1, p_2) + (p_2 - c_2)Q_2(p_1, p_2) = r\,p_1(1 - G(\hat{\theta})) + (p_2 - c_2)(G(\hat{\theta}) - G(\tilde{\theta}_2))\end{aligned}$$

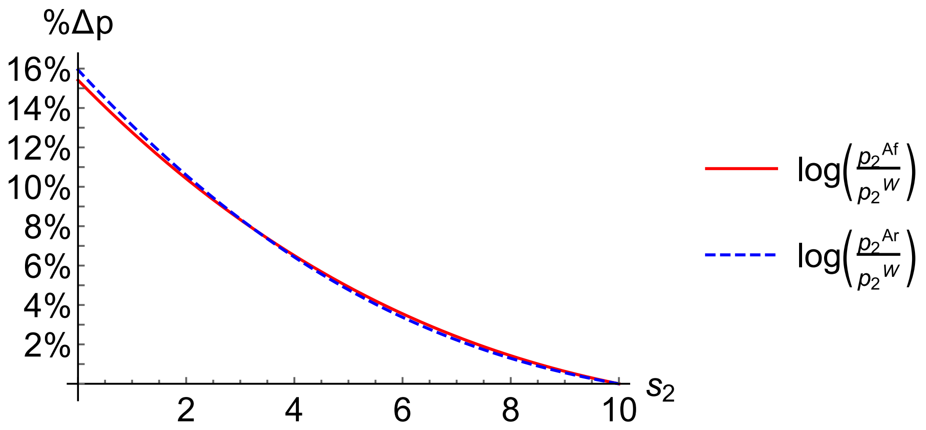

Figure 2.B.2: Differences in marketplace private label prices under per-unit (red) and ad-valorem (blue, dashed) fee to stand-alone monopolist or intermediary's private label price in the wholesale model for a given quality of the private label s_2. Calibrated for uniformly distributed preferences, $\bar{\theta} = 3$, $\alpha = 0.1$, $s_1 = \frac{\bar{\theta}}{3\alpha} = 10$.

where the term in the brackets (p_1, p_2) is omitted from now on for convenience of notation. The seller switches to setting the duopoly margin $p_1 - \frac{c_1}{1-r} = \gamma(\hat{\theta})(s_1 - s_2) = \bar{\theta}s_1 - p_2 - (\bar{\theta}s_2 - p_2)$. Its price decreases with the improvement of the private label quality due to the competitive pressure. Because it is lower than under the per-unit fee contract, and also lower than in the wholesale model, the seller has higher sales with ad-valorem fee than with per-unit fee, or in the wholesale model.

The result that the marketplace operator sets a higher than monopoly price for the private label compared to a retailer or independent monopolist carries on with an ad-valorem fee contract. The share of the seller's price collected by the marketplace operator rp_1 serves the same purpose in the optimal expression for the private label price

$$p_2 - c_2 = -\frac{Q_2}{dQ_2/dp_2} - \frac{dQ_1/dp_2}{dQ_2/dp_2}rp_1 = (1-\omega)\gamma(\tilde{\theta}_1)s_2 - \omega\gamma(\hat{\theta})(s_1 - s_2) + \omega r p_1$$

and represents an opportunity cost. The relative difference between the price of a private label in the agency model and its price in the wholesale structure after the adjustment of the ad-valorem fee is similar to the case with a per-unit fee (see Figure 2.B.2).

The prices with an ad-valorem fee are then given by:

$$p_1 = \frac{(2\bar{\theta}(s_1 - s_2) + c_2 + 2\frac{c_1}{1-r})s_1}{4s_1 - (1+r)s_2}$$
$$p_2 = \frac{(1+r)s_2(\bar{\theta}(s_1 - s_2) + \frac{c_1}{1-r}) + 2c_2 s_1}{4s_1 - (1+r)s_2}$$

The price p_1 decreases the better the quality of the private label.

Solving for the optimal ad-valorem commission in the absence of a private label is already mathematically involved. In the presence of a private label the optimal commission is the solution to:

$$r = \frac{1-r}{c_1}\Bigg[\frac{s_1(2\bar{\theta}(s_1-s_2)+c_2+2\frac{c_1}{1-r})(2\bar{\theta}(s_1-s_2)+c_2-\frac{c_1}{1-r}(2s_1-(1+r)s_2)}{s_2(2\bar{\theta}(s_1-s_2)+c_2)+\frac{c_1}{(1-r)^2}(8s_1-4rs_2)} + \frac{(1+r)s_2(\bar{\theta}(s_1-s_2)+\frac{c_1}{1-r})-c_2(2s_1-(1+r)s_2)}{4s_1-(1+r)s_2}\Bigg] \Leftrightarrow$$

$$\frac{rc_1}{1-r} = \Bigg[\frac{s_1(2\bar{\theta}(s_1-s_2)+c_2+2\frac{c_1}{1-r})(2\bar{\theta}(s_1-s_2)+c_2-\frac{c_1}{1-r}(2s_1-(1+r)s_2)}{s_2(2\bar{\theta}(s_1-s_2)+c_2)+\frac{c_1}{(1-r)^2}(8s_1-4rs_2)} + \frac{(1+r)s_2(\bar{\theta}(s_1-s_2)+c_2+\frac{c_1}{1-r})-c_2 2s_1}{4s_1-(1+r)s_2}\Bigg]$$

Numerical simulations show how this one changes in the presence of a private label.

After the introduction of the private label the optimal ad-valorem fee r increases. The mechanism, which drives the change of the ad-valorem fee post private label introduction is more complex than, and different from, the one of the per-unit fee. The fact that the price of the seller's product decreases after the introduction of the private label has an additional direct implication for the revenues collected by the intermediary. While with a per-unit fee the revenues of the intermediary decrease due to the demand-stealing effect and increased pass-through of the seller's price, with an ad-valorem fee they additionally decrease as a direct consequence of decreasing p_1. Thus, the intermediary has an incentive to increase the ad-valorem fee in order to recover some, but not all of the revenues from the seller's sales.

The increase in the ad-valorem fee cannot be directly interpreted as an incentive for the marketplace operator to increase the cost for the seller, even though it has this effect in its optimality condition. As pointed out by Bishop (1968) (who compares per-unit and proportionate taxes imposed by the government on a monopolist), one should be cautious when comparing a per-unit fee with the same dimension as the marginal cost, to an ad-valorem fee, which is a dimensionless ratio. For the right comparison one needs to look at how rp_1 behaves after the introduction of the private label. As depicted by Figure 2.B.3,

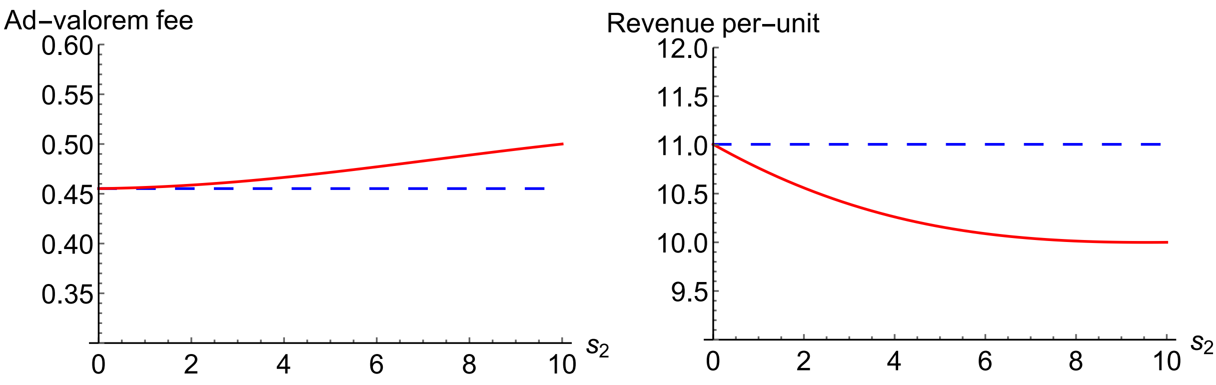

Figure 2.B.3: Ad-valorem fee and revenue per unit extracted by the intermediary before (blue, dashed) and after (red, solid) the introduction of the private label for different s_2. Calibrated for uniformly distributed preferences, $\bar{\theta} = 3$, $\alpha = 0.1$, $s_1 = \frac{\bar{\theta}}{3\alpha} = 10$.

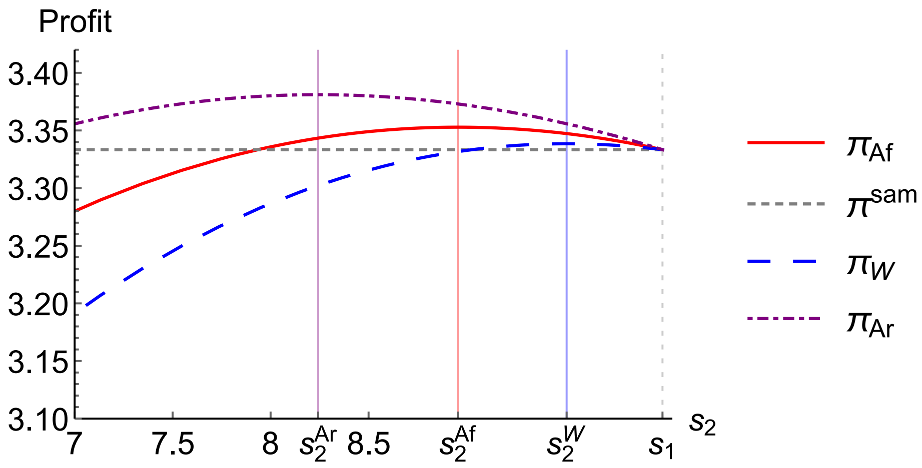

Figure 2.B.4: Profits of the intermediary in the agency model with per-unit (red) and ad-valorem (purple, dot-dashed) fee, in the wholesale model (blue, large dashed), and as a stand-alone monopolist with quality s_1 (grey, small dashed) as a function of quality s_2 with respective optimum. Calibrated for uniformly distributed preferences, $\bar{\theta} = 3$, $\alpha = 0.1$, $s_1 = \frac{\bar{\theta}}{3\alpha} = 10$.

even though the ad-valorem fee increases, it is to only partially recover revenues, and the overall result is that the extracted revenue per unit decreases.

Finally, with an ad-valorem fee the result of the marketplace operator choosing optimally a lower quality than the seller continues to hold. Given that the vertical structure generates larger industry revenues and a larger share goes to the marketplace operator, there is an even stronger incentive to differentiate the product, compared to the per-unit fee case (see Figure 2.B.4). This also reverses the result that a seller generates lower profits with an ad-valorem fee and leaves the seller with even more profits ex-post compared to a per-unit fee contract or in the wholesale model (see Figure 2.B.5). Interestingly, numerical simulations also show that, even though for a given quality of the private label

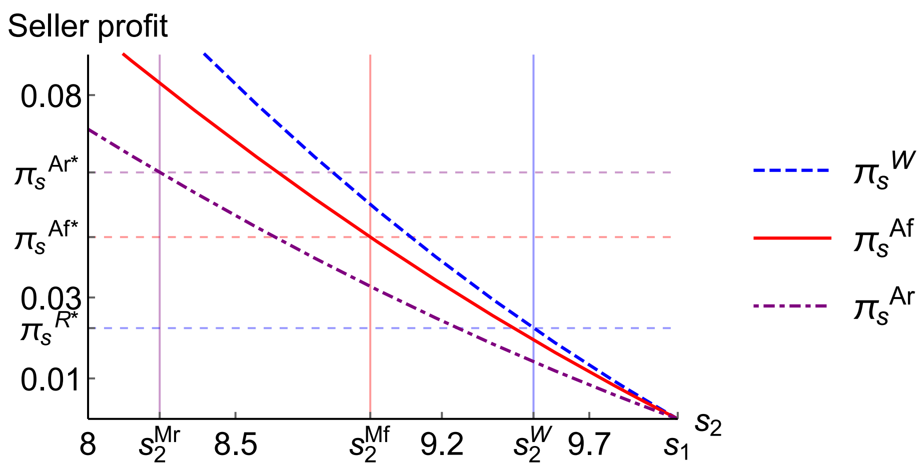

Figure 2.B.5: Profits of the seller in the agency model with per-unit (red, solid) and ad-valorem (purple, dot-dashed) fee, and in the wholesale model (blue, dashed) with optimal quality of the private label and ex-post profits in each mode. Calibrated for uniformly distributed preferences, $s_1 = \frac{\bar{\theta}}{3\alpha} = 10$, $\alpha = 0.1$.

consumer surplus is higher under ad-valorem fee, this result gets reversed by the quality choice of the marketplace operator, and the consumer surplus is the lowest ex-post among the three vertical models.

2.B.4 The Case of a Higher-Quality Private Label

In this section, I discuss the outcome if an intermediary introduces a higher quality than the seller, i.e. $s_2 \geq s_1$. This implies that the demand for the intermediary's product is now given by $Q_2(p_1, p_2) = 1 - G(\hat{\theta})$ and for the seller by $Q_1(p_1, p_2) = G(\hat{\theta}) - G(\tilde{\theta}_1)$, where $\hat{\theta} = \frac{p_2 - p_1}{s_2 - s_1}$. I demonstrate that the results presented here hold for uniformly distributed θ, but as the logic is similar to the case with a lower-quality private label, these should hold more generally. For ease of notation the terms (p_1, p_2) are omitted from $Q_1,\ Q_2$.

In the agency model when offering the higher-quality private label the results are similar under similar conditions. In the Proof of Proposition 2.1 the condition for the fee to decrease and the private label to be priced above monopoly level is that the quality of the outside seller product is efficient, or in other words, when both products are priced at marginal cost there is demand for the higher-quality product. The same holds in the case where the private label has the higher quality, as long as the lower-quality product (the seller's product) has positive demand when both products are priced at marginal cost.

In particular, the consumer prices in the agency model are given by:

$$\begin{aligned}
p_1 - f - c_1 &= (1 - \omega_1)\gamma(\tilde{\theta}_1)s_1 - \omega_1 \gamma(\hat{\theta})(s_2 - s_1) \\
p_2 - c_2 &= \gamma(\hat{\theta})(s_2 - s_1) + f \Leftrightarrow \\
p_1 &= \frac{\bar{\theta}(s_1 s_2 - s_1^2) + f(2s_2 + s_1) + c_1 2s_2 + c_2 s_1}{4s_2 - s_1} \\
p_2 &= \frac{s_2(2\bar{\theta}(s_2 - s_1) + 3f + c_1 + 2c_2)}{4s_2 - s_1}
\end{aligned}$$

where $\omega_1 = -\frac{dQ_2/dp_1}{dQ_1/dp_1} = \frac{g(\hat{\theta})s_1}{g(\hat{\theta})s_1 + g(\tilde{\theta}_1)(s_2 - s_1)}$ is the diversion ratio of the lower-quality product. Notice that with a higher-quality private label the opportunity cost that the per-unit fee represents is not weighted down as in the lower-quality case. This is because all of the demand lost for the high-quality product by a price increase goes directly to the low-quality product. Like in Lemma 2.2 for the lower-quality private label case, there exists a threshold such that if $f \geq \frac{\bar{\theta} 3 s_1 s_2 - 2s_2 c1 - s_1 c_2}{6s_2}$ the price of the private label is higher than the one set by a stand-alone monopolist. Given that before the private label the optimal price of the seller product as a function of the per-unit fee was given by $p_1 = \frac{\bar{\theta}s_1 + c_1 + f}{2}$, it is apparent from the expression above that its pass-through has increased after the introduction of the private label.

When choosing the optimal fee after the introduction of the private label, a similar optimality condition arises:

$$f = \Big((1-\omega_1)\gamma(\tilde{\theta}_1)s_1 - \omega_1\gamma(\hat{\theta})(s_2-s_1)\Big) \Big/ \frac{\partial p_1}{\partial f} + \omega_1(p_2 - c_2) \Leftrightarrow$$
$$f = \frac{s_2(\bar{\theta}(8s_2+s1)s_1 - 8s_2c_1) - s_1^2c_2}{2s_2(8s_2+s_1)}$$

Because demand for the outside seller's product is more responsive to price changes, revenues from the private label can only compensate a share ω_1 of demand lost from the seller's product. Thus, the reason the fee is lower after the introduction of a higher-quality private label is similar to a lower-quality private label. First, there is a demand-stealing effect, that drives pure marginalization down. Then, the pass-through of the seller's price increases, further limiting the ability of the marketplace operator to collect revenues from the seller's product. As the private label can only partially compensate for the decrease in pure marginalization, the fee decreases after the introduction of the private label.

Compared to pre entry $f = \frac{\bar{\theta}s_1 - c_1}{2}$, the fee clearly decreases after the introduction of the private label as long as $\frac{c_2}{s_2} \geq \frac{c_1}{s_1}$, or the lower-quality product is efficient and has positive demand when both products are priced at marginal cost.

In the wholesale model, the intermediary chooses prices independently under the assumption of a constant curvature demand:

$$p_1 - w = \gamma(\tilde{\theta}_1)s_1 \qquad \Leftrightarrow p_1 = \frac{\bar{\theta}s_1 + w}{2}$$
$$p_2 - c_2 = \gamma(\hat{\theta})(s_2 - s_1) + p_1 - w = \gamma(\tilde{\theta}_2)s_2 \qquad \Leftrightarrow p_2 = \frac{\bar{\theta}s_2 + c_2}{2}$$

The outside seller decreases the wholesale price due to a demand-stealing effect by the newly introduced product:

$$w - c_1 = \Big((1-\omega_1)\gamma(\tilde{\theta}_1)s_1 - \omega_1\gamma(\hat{\theta})(s_2-s_1)\Big) \Big/ \frac{\partial p_1}{\partial w} \Leftrightarrow w = \frac{c_2s_1 + c_1s_2}{2s_2}$$

Much like the lower-quality private label, the wholesale price decreases due to the demand-stealing effect. The closer quality s_2 is to s_1, the lower the wholesale price, reaching $w = c_1$ when $s_2 = s_1$.

Comparing prices p_1, p_2 in both vertical models reveals higher prices in the agency model, particularly evident in the private label price difference. This contributes to diminished private label sales in the agency model versus the wholesale model.

Similar to the situation with the lower-quality private label, choosing the same quality for the private label as the seller's product ($s_2 = s_1$) enables the intermediary to foreclose the outside seller. However, the intermediary's preference for differentiation remains intact in both vertical models, being more pronounced in the agency model. Assuming the seller offers stand-alone monopoly quality, the convex cost function implies that higher-quality entry yields lower profits than lower-quality entry, favoring the intermediary's choice of a lower-quality private label when entering in both vertical models.

Bibliography

Akgün, U., Caffarra, C., Etro, F., & Stillman, R. (2020). On the welfare impact of mergers of complements: Raising rivals' costs versus elimination of double marginalization. *Economics Letters*, *195*, 109429.

Anderson, S., & Bedre-Defolie, Ö. (2022). Online trade platforms: Hosting, selling, or both? *International Journal of Industrial Organization*, *84*, 102861.

Anderson, S. P., & Bedre-Defolie, Ö. (2023). *Hybrid platform model: Monopolistic competition and a dominant firm* (tech. rep.). The RAND Journal of Economics.

Benassi, C., Chirco, A., & Colombo, C. (2019). Vertical differentiation beyond the uniform distribution. *Journal of Economics*, *126*(3), 221–248.

Berges-Sennou, F., Bontems, P., & Réquillart, V. (2004). Economics of private labels: A survey of literature. *Journal of Agricultural & Food Industrial Organization*, *2*(1).

Berthene, A. (2019). Amazon pushes first-party suppliers toward brand registry and seller central.

Besanko, D., Donnenfeld, S., & White, L. J. (1987). Monopoly and quality distortion: Effects and remedies. *The Quarterly Journal of Economics*, *102*(4), 743–767.

Bisceglia, M., & Padilla, J. (2023). On sellers' cooperation in hybrid marketplaces. *Journal of Economics & Management Strategy*, *32*(1), 207–222.

Bishop, R. L. (1968). The effects of specific and ad valorem taxes. *The Quarterly Journal of Economics*, *82*(2), 198–218.

Bontems, P., Monier-Dilhan, S., & Requillart, V. (1999). Strategic effects of private labels. *European Review of Agricultural Economics*, *26*(2), 147–165.

Draganska, M., Klapper, D., & Villas-Boas, S. B. (2010). A larger slice or a larger pie? an empirical investigation of bargaining power in the distribution channel. *Marketing Science*, *29*(1), 57–74.

Ellis, B., & Hicken, M. (2020). Dozens of amazon's own products have been reported as dangerous – melting, exploding or even bursting into flames. many are still on the market.

Etro, F. (2023a). The economics of amazon. *Available at SSRN 4307213*.

Etro, F. (2023b). Hybrid marketplaces with free entry of sellers. *Review of Industrial Organization*, 1–30.

Gautier, A., Madio, L., & Shekhar, S. (2021). *Platform duality and network externalities* (Working paper).

Gautier, P., Hu, B., & Watanabe, M. (2023). Marketmaking middlemen. *The RAND Journal of Economics*, *54*(1), 83–103.

Hagiu, A., Teh, T.-H., & Wright, J. (2022). Should platforms be allowed to sell on their own marketplaces? *The RAND Journal of Economics*, *53*(2), 297–327.

Hagiu, A., & Wright, J. (2015). Marketplace or reseller? *Management Science*, *61*(1), 184–203.

Heese, H. S. (2010). Competing with channel partners: Supply chain conflict when retailers introduce store brands. *Naval Research Logistics (NRL)*, *57*(5), 441–459.

Hervas-Drane, A., & Shelegia, S. (2022). *Retailer-led marketplaces* (tech. rep.). CEPR Discussion Paper No. DP17351.

Itoh, M. (1983). Monopoly, product differentiation and economic welfare. *Journal of Economic Theory*, *31*(1), 88–104.

Johnson, J. P. (2017). The agency model and mfn clauses. *The Review of Economic Studies*, *84*(3), 1151–1185.

Johnson, J. P., & Myatt, D. P. (2015). The properties of product line prices. *International Journal of Industrial Organization*, *43*, 182–188.

Kang, Z. Y., & Muir, E. V. (2022). Contracting and vertical control by a dominant platform. *Unpublished manuscript, Stanford University*.

Kittaka, Y., & Sato, S. (2022). Dual role platforms and search order distortion. *Available at SSRN 3736574*.

Lam, W. M. W., & Liu, X. (2021). *Data usage and strategic pricing: Does platform entry benefit independent traders?* (Working paper).

Madsen, E., & Vellodi, N. (2022). Insider imitation. *Available at SSRN 3832712*.

Marketplace Pulse. (2019). *Amazon private label brands* (Survey).

Maskin, E., & Riley, J. (1984). Monopoly with incomplete information. *The RAND Journal of Economics*, *15*(2), 171–196.

Meza, S., & Sudhir, K. (2010). Do private labels increase retailer bargaining power? *Quantitative Marketing and Economics*, *8*(3), 333–363.

Mills, D. E. (1995). Why retailers sell private labels. *Journal of Economics & Management Strategy*, *4*(3), 509–528.

Mussa, M., & Rosen, S. (1978). Monopoly and product quality. *Journal of Economic theory*, *18*(2), 301–317.

Padilla, J., Piccolo, S., & Vasconcelos, H. (2022). Business models, consumer data and privacy in platform markets. *Journal of Industrial and Business Economics*, 1–36.

Pouyet, J., & Thomas, T. (2021). The competitive effects of vertical integration in platform markets.

Rhodes, A., Watanabe, M., & Zhou, J. (2021). Multiproduct intermediaries. *Journal of Political Economy*, *129*(2), 421–464.

Ritz, R. A. (2018). Oligopolistic competition and welfare. In *Handbook of game theory and industrial organization, volume i* (pp. 181–200). Edward Elgar Publishing.

Tirole, J. (1988). *The theory of industrial organization*. MIT press.

Tremblay, M. J. (2022). Fee discrimination by hybrid platforms. *Available at SSRN 4195017*.

Weyl, E. G., & Fabinger, M. (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy*, *121*(3), 528–583.

Zennyo, Y. (2022). Platform encroachment and own-content bias. *The Journal of Industrial Economics*, *70*(3), 684–710.

Chapter 3

Hybrid Platforms: Managing Price Competition with Strategic Sellers

3.1 Introduction

Online trade platforms often operate in a so-called hybrid mode - they run a marketplace for third-party sellers and collect fees from their sales, and offer their own products in direct competition with sellers. This has drawn the attention of regulatory authorities both in the EU and the US, as a hybrid platform constitutes a conflict of interest. A platform operator that introduces its own product may have an incentive to distort competition in its favor, which can be detrimental to outside sellers, consumers, and total welfare.

This paper addresses the question of whether a platform operator has an incentive to increase the costs for outside sellers to compete on the platform when it introduces its own product. I show that when the platform operates in hybrid mode, it sets a lower fee than when it operates as a pure platform. This result is in line with the predictions of Shopova (2023), but expands it to a horizontal differentiation framework and allows for multiple sellers, which highlights the role of sellers' market power in the mechanism behind the fee decrease. Hybrid mode is always more profitable than only operating a pure platform, as it strengthens competition and expands demand not only by increasing the number of products but also by avoiding double marginalization. Finally, a platform has no incentive to foreclose third-party sellers as long as there are variety effects. Consumers are always better off when the platform introduces its own product, whereas sellers' profits get squeezed out.

To understand these results, consider a platform that runs a marketplace for third-party sellers. The sellers are symmetric and offer horizontally differentiated products, and compete in prices to capture consumer demand. The platform sets a per-unit fee it collects from their sales. When the platform offers its own product, it has the same quality and cost and is symmetrically differentiated from all other products. Demand for the sellers' products stems from a representative consumer with quasi-linear quadratic utility, resulting in a linear demand system for differentiated goods.

When the platform only operates a pure marketplace with horizontally differentiated sellers (pure mode), it chooses the optimal per-unit fee to be equal to the stand-alone monopoly margin that can be extracted for each product. Under the assumption of symmetric sellers with constant marginal cost and linear demand, a standard result is that the optimal per-unit fee is independent of the number of sellers or the level of product differentiation. Nevertheless, a larger number of sellers implies more competition and lower prices, which leads to greater demand and profits for the platform but lower profits for the sellers.

When the platform introduces its own product, it competes in prices with the sellers. Even though the platform's product enjoys the cost advantage of not paying the fee, it faces an opportunity cost in the form of the fee: any demand that the platform product diverts from third-party sellers results in lost fee revenues. As the opportunity cost for the platform is lower than the direct cost of the sellers, it can set a lower price. However, the more sellers there are in the market, the higher the opportunity cost, the lower the cost advantage, and the closer the price of the platform product is to the prices of the sellers.

The platform sets a lower fee for third-party sellers when it operates in hybrid mode than when it operates as a pure platform without its product. The reason is that the platform's product enjoys an effective cost advantage, allowing it to charge a lower price and divert demand from sellers. Since sellers have market power, they are reluctant to sufficiently decrease their prices in response to the cost advantage, resulting in a stronger decrease in demand for their products. To compensate for this effect, the platform reduces the fee. Similarly to the price, the presence of more firms in the market reduces the market power of sellers and the cost advantage of the platform, leading to a weaker fee decrease when there are more sellers on the platform.

The platform always has an incentive to introduce its own product. As more products increase competition, lower prices, and increase demand, even in pure mode, the platform gains from accommodating more sellers. Additionally, since the platform's product avoids double marginalization, the platform prefers to introduce its own product rather than accommodating another seller.

The platform prefers to accommodate all sellers and operate in hybrid mode rather than foreclosing sellers and monopolizing the product space. The profits of the platform in both pure and hybrid modes increase when more sellers join the platform due to two distinct effects: the competition effect, which brings prices down and increases demand; and the variety effect, which expands the market when new products are introduced, even if prices are kept fixed and do not adjust. In the extreme case where the products are homogeneous, the entry of a new product does not grow the market, and the platform is indifferent between foreclosing or operating in hybrid mode.

Finally, the decision of the platform to introduce its own product always benefits consumers as it avoids double marginalization and intensifies competition. Due to variety effects it improves total welfare through market expansion and decrease in double marginalization. However, this comes at the detriment of sellers, whose profits are squeezed out.

Related literature. This paper contributes to the literature by expanding the results in Shopova (2023) to a multi-seller horizontal differentiation framework. It focuses on the strategic interaction arising from differentiated Bertrand competition between a hybrid platform and an exogenous number of third-party sellers it hosts on its marketplace. I show that sellers lose more sales than is optimal from the platform's perspective due to their insufficient price response to the platform product's fee-cost advantage. To correct for their too-high prices the platform lowers the per-unit fee it charges for access to the marketplace. The literature on intermediaries operating in hybrid mode has grown steadily in recent years.[1] The closest papers that investigate how platform entry affects market outcomes are Anderson and Bedre-Defolie (2023), Etro (2023b), Zennyo (2022), and Anderson and Bedre-Defolie (2022). Anderson and Bedre-Defolie (2023), Zennyo (2022), and Etro (2023b) focus on a free entry framework that leads to an endogenous number of sellers. They show that when the platform introduces its own product, it causes some sellers to leave the marketplace. The platform then faces a trade-off between increasing the fee to further induce exit, relax competition, and shift demand to its own product, or decreasing the fee to compensate for the exit of third-party sellers. Etro (2023b) highlights how the decision to raise or lower the fee is tied to the relative (constant) elasticity of demand for platform products versus third-party products. However, these papers shut down the strategic price response of sellers to platform entry either because they assume monopolistic competition (Anderson and Bedre-Defolie (2023), Etro (2023b)), or a fixed sample search framework with logit demand (Zennyo (2022)). Thus, the current paper proposes a complementary mechanism through which the optimal fee choice is affected by the entry of the platform. Finally, in Anderson and Bedre-Defolie (2022), the platform excludes the third-party seller after introducing its product and competes with the seller through an alternative channel, not in a hybrid platform setting.

Currently, there is limited but growing empirical literature on platforms that operate in hybrid mode. In their study, Zhu and Liu (2018) were the first to investigate the entry of a trading platform into the product space of hosted sellers. They discovered that when Amazon offers its own products, it discourages rival sellers from expanding their business on the platform. In a recent study, Crawford et al. (2022) investigated Amazon's entry with its products on its Marketplace in Germany. They found that Amazon tends to enter categories that have low competition but high growth. After Amazon enters a category, prices slightly decrease, with Amazon's product prices being slightly lower

[1]For a comprehensive overview, see Anderson and Bedre-Defolie (2022) and Etro (2023a).

than the prices set by hosted third-party sellers. The authors conclude that Amazon's entry aims to expand the market, not to steal business. These observations align with the findings presented in this paper. Another study by Lee and Musolff (2021) analyzes Amazon's self-preferencing in the Buy Box and suggests that a prohibition on it could harm consumers. In conclusion, Gutierrez (2021) proposes that prohibiting the hybrid mode could increase fees and reduce the selection of products on Amazon.

The paper relates to the literature on vertical mergers.[2] In the current model framework, entry with the own product has the same ex-post market structure and effect on prices and fees as a two-step process, where first the platform attracts an additional seller and then vertically integrates with this seller. However, by varying the number of sellers and the substitutability of products, I demonstrate that the reason for the decrease in the input price, or fee in the current setting, is related to the market power of sellers and how strongly they react and pass on a shock to their demand in their pricing. In particular, the shock in this paper is that the platform's product avoids double marginalization and has a lower effective cost, to which outside sellers react by decreasing their prices insufficiently from the platform's perspective. Thus, the paper supports and expands upon results found in recent papers on vertical mergers, which focus on downstream duopoly (Arya et al. (2008), Domnenko and Sibley (2020), and Das Varma and De Stefano (2020)), to downstream oligopoly.

The paper is also related to the retail management literature, as the problem presented here is akin to a manufacturer deciding between distributing its product through different retailers or selling directly to consumers. However, most of the existing literature in this domain focuses on efficiency discussions (Tsay and Agrawal (2004)), cost disadvantages (Arya et al. (2007)), demand advantage (Cai (2010)), and channel decisions based on the type of contract (Ha et al. (2022)). In contrast, the analysis in this paper centers on efficient sellers and instead focuses on the strength of competition in terms of differentiation, variety effects, and competition level, and how these factors influence pricing decisions.

The paper is structured as follows: Section 3.2 introduces the model. Section 3.3 discusses the effects of the introduction of a product by a platform. Section 3.4 concludes. Omitted proofs and additional material can be found in the Appendix.

3.2 Model

There is a representative consumer with a quasi-linear quadratic utility (QQU) of the form

$$U(q_1, ..., q_n, \widetilde{q}) = \sum_i a q_i - s \sum_i \sum_{j>i} q_i q_j - \frac{b}{2} \sum_i q_i^2 + \widetilde{q} \tag{3.1}$$

[2]For recent advances and discussions, see Salinger (2021).

where q_i is the amount consumed of product $i = 1, .., n$ and $\widetilde{q}$ is a a numéraire good with a price normalized to 1. All products are symmetric with the marginal quality index of each product equal to a, $b > 0$, and $s \in [0, 1]$. When $s = 0$, products are independent, while when $s = b$, products are homogeneous. Thus, any $s \in (0, b)$ corresponds to products being imperfect substitutes. Consumers maximize their utility in (3.1) by choosing their optimal consumption level $(q_1, ..., q_n)$ taking the budget constraint $\widetilde{q} + \sum_i p_i q_i \leq m$ as given, where p_i is the price of variety i and m is the income level. Assuming that m is large enough and eliminating $\widetilde{q}$ simplifies to the following optimization problem:[3]

$$\max_{q_1,..,q_n} \sum_i (a - p_i) q_i - s \sum_i \sum_{j>i} q_i q_j - \frac{b}{2} \sum_i q_i^2$$

This formulation is due to Spence (1976), who used it to derive a linear demand system for differentiated goods.[4] Thus, the optimization problem yields a system of indirect demand functions $p_i = a - bq_i + s\sum_{j\neq i} q_j$. Solving for the individual direct demand functions yields:

$$q_i(p_1, ..., p_n) = \alpha_n - \beta_n p_i + \gamma_n \sum_{j\neq i} p_j \tag{3.2}$$

for $i = 1, .., n$ where the parameters $\beta_n = \frac{b+(n-2)s}{(b-s)(b+(n-1)s)}$, $\gamma_n = \frac{s}{(b-s)(b+(n-1)s)}$, and $\alpha_n = (\beta_n - (n-1)\gamma_n)a = \frac{a}{b+(n-1)s}$, will be used from now on for ease of notation and exposition.

A few remarks are in order at this point. The maximum market share achievable by a single firm α_n decreases in n to reflect the fact that each new substitutable product in the market necessarily steals some demand from existing firms. In addition, the aggregate diversion ratio $-\frac{\sum_{j\neq i} dq_j/dp_i}{dq_i/dp_i} = \delta_n = \frac{(n-1)\gamma_n}{\beta_n} = \frac{(n-1)s}{b+(n-2)s} \leq 1$ increases in n to reflect that when a firm raises its price, less demand is lost to the outside option relative to rivals when there are more competitors in the market. However, due to variety effects, the total market size $n\alpha_n$ increases in n as long as $s < b$.

Each of the products in the market is offered by a separate firm, and all firms have the same marginal cost of production $c < a$. Firms must use the services of a monopoly platform to reach the consumer side. The platform provider charges them a per-unit fee f to sell on the platform. The platform provider may also offer its own product, which is symmetric with the products of other outside sellers and faces the same demand function as in (3.2). I assume that the platform has no advantage in producing its product and has the same marginal cost c.[5] The timing of the game is as follows: first, the platform sets the fee f; then, all firms offering a product on the platform compete in prices for the demand of the representative consumer. The solution concept is Subgame Perfect Equilibrium.

[3]Note that this assumption is mild only for the case of products being substitutes, as emphasized by Amir et al. (2017).

[4]See Choné and Linnemer (2020) for the history of QQU models and the various formulations used to derive linear demand systems.

[5]Besides the advantage that the platform product does not have to pay the fee.

3.3 The Effects of a Platform Introducing its own Product

To analyze the effects of platform owner entry with its own product, I first analyze the case where the platform operates in pure mode and only runs a marketplace for n third-party sellers. Then, I allow the platform to introduce its own product and operate in hybrid mode, still hosting the n outside sellers. Next, I discuss the incentives for the platform to operate in hybrid mode to running a pure platform or foreclosing sellers and monopolizing their product space by stressing the importance of variety effects and sellers' market power. By backward induction, I analyze each of the two platform modes by starting at the pricing stage and then moving to the optimal fee stage.

3.3.1 Pure Platform

Pricing stage. Firms compete in prices with their objective function to maximize profits taking the number of firms active in the market n, the marginal cost c and the fee they have to pay f as given, i.e.,

$$\max_{p_i} \pi_i = (p_i - f - c)q_i(p_1, ..., p_n)$$

for $i = 1, .., n$ and q_i is as in (3.2). Due to symmetry, the optimal price is the same for all products $i = 1, .., n$ and is equal to

$$p_i^*(f) = \frac{\alpha_n + \beta_n(f + c)}{2\beta_n - (n-1)\gamma_n} = \frac{a(1-\delta_n) + f + c}{2 - \delta_n} \tag{3.3}$$

This price decreases with the number of sellers as the diversion ratio δ_n increases in n. Despite the decreasing prices of firms, their individual demand still shrinks when new rivals enter the market, i.e., $q_i^*(f) = (a - p_i^*(f))(\beta_n - (n-1)\gamma_n) = \frac{(a-f-c)(\beta_n-(n-1)\gamma_n)}{2-\delta_n} = \frac{(a-f-c)(b+(n-2)s))}{(2b+(n-3)s)(b+(n-1)s)}$ decreases in n.

Optimal fee. The platform maximizes the revenues from sellers' sales by setting the per-unit fee in anticipation of the pricing decisions of sellers in (3.3) and their pricing reactions to the fee, i.e.,

$$\max_f \Pi_P = f \sum_i q_i(p_1^*(f), ..., p_n^*(f)).$$

The optimal fee balances the loss in revenues from the contraction in demand due to the higher fee and higher prices with the gain from collecting this higher fee from the existing demand:

$$\underbrace{-f\sum_{i=1}^{n}\left[\frac{dq_i}{dp_i}+\sum_{j\neq i}\frac{dq_j}{dp_i}\right]\frac{\partial p_i}{\partial f}}_{\text{Marginal loss}} \stackrel{!}{=} \underbrace{\sum_i q_i}_{\text{Marginal gain}} \Leftrightarrow -fn\left[\frac{dq_1}{dp_1}+\sum_{i=2}^{n}\frac{dq_i}{dp_1}\right]\frac{\partial p_1}{\partial f} \stackrel{!}{=} nq_1$$

$$\Leftrightarrow fn(\beta_n-(n-1)\gamma_n)\frac{1}{2-\delta_n} \stackrel{!}{=} n(\beta_n-(n-1)\gamma_n)\frac{a-f-c}{2-\delta_n} \tag{3.4}$$

where the first equivalence follows from the fact that all third-party firms are symmetric. The optimal fee is then such that

$$f = a-f-c = \frac{a-c}{2} \tag{3.5}$$

The fee is completely independent of b and s and extracts the monopoly surplus from each product in the market.

The optimal fee in a pure platform is also independent of the number of firms. The marginal loss from raising the fee increases as more firms join the platform - because there is more demand in the market and thus more consumers would prefer to take the outside option when the price rises. In addition, the pass-through of sellers' prices of the fee $\frac{\partial p_i}{\partial f} = \frac{1}{2-\delta_n}$ increases, as long as δ_n increases with entry.[6] However, as more competition and products in the market imply lower prices and thus more demand, the marginal gain also increases and offsets the increase in the cost, resulting in the same optimal fee for any n.

The fact that the per-unit fee in a pure platform is independent of the substitution parameter s and especially of the number of firms n is a result of the assumption of perfectly symmetrical firms facing linear demand and constant marginal costs. Nevertheless, it is a useful result that allows one to emphasize the difference for a platform between introducing its own product and accommodating another seller and how this changes the incentive to set the fee in hybrid mode.

Finally, due to the vertical structure, there is a double marginalization problem such that the final price of each product in the market necessarily exceeds the price chosen by a vertically integrated monopolist. Technically, for f as in (3.5) the optimal price in (3.3) reduces to $p_i^* = \frac{a(3-2\delta_n)+c}{2(2-\delta_n)} > \frac{a+c}{2} = p_m$ for $a > c$. Since linear demand is log-concave, this implies that the platform extracts a larger share of the total profit per product in the market.

3.3.2 Hybrid Platform

When the platform introduces a new substitute product in the market, the number of products to choose from for consumers increases from n to $n+1$. Optimising consumption

[6] More generally, as argued by Weyl and Fabinger (2013) and emphasized by Ritz (2018) the rate of pass-through rate reflects the distribution of welfare in the market between consumers and producers. More competition decreases market power and pushes the prices closer to marginal costs, leaving more surplus for consumers.

then implies that demand for the n third-party sellers' and the platform's products are respectively given by:

$$q_i(p_0, p_1, ..., p_n) = \alpha_{n+1} - \beta_{n+1} p_i + \gamma_{n+1} \sum_{j \neq i,0} p_j + \gamma_{n+1} p_0$$
$$q_0(p_0, p_1, .., p_n) = \alpha_{n+1} - \beta_{n+1} p_0 + \gamma_{n+1} \sum_{i} p_i$$

for $i = 1, .., n$ where p_0 denotes the price of the platform product. The total market share of an individual firm decreases ($\alpha_{n+1} \leq \alpha_n$), and the share of demand that a firm loses to its rivals when it raises its price also increases ($\delta_{n+1} \geq \delta_n$) to account for the new substitute product. In other words, demand for existing products shrinks as consumers can substitute them with the new product. How firms change their pricing decision in response to the platform's entry will be discussed in the following section.

Pricing stage. When the platform introduces its own product, it competes with n other firms in the market. Its product faces the same demand function and is produced at the same marginal cost c. The only difference is that the platform, as a seller, does not pay the fee and avoids double marginalization, giving it an advantage over the other firms. Thus, the objective functions of the third-party sellers and the platform are given by:

$$\max_{p_i} \pi_i = (p_i - f - c) q_i(p_0, p_1, ..., p_n)$$
$$\max_{p_0} \Pi_H = f \sum_i q_i(p_0, p_1, .., p_n) + (p_0 - c) q_0(p_0, p_1, .., p_n)$$

The best response functions of firms $i = 1, ..., n$ and of the platform are:

$$p_i = \frac{\alpha_{n+1} + \beta_{n+1}(f + c) + \gamma_{n+1} \sum_{j \neq i,0} p_j + \gamma_{n+1} p_0}{2\beta_{n+1}}$$
$$p_0 = \frac{\alpha_{n+1} + \beta_{n+1} c + \gamma_{n+1} \left(f + \sum_i p_i\right)}{2\beta_{n+1}}$$

for $j \neq i$. Since the platform does not have to pay the fee, it has a cost advantage. However, because when it lowers its price it diverts a share $\delta_{n+1} = -\frac{\sum_i dq_i/dp_0}{dq_0/dp_0}$ of the demand of third-party sellers to its own product, it loses $\delta_{n+1} f$ revenue from sales from sellers. This opportunity cost drives the platform to increase its price and relax competition with sellers. Nonetheless, as long as $\delta_{n+1} < 1$, the platform charges a lower price than third-party sellers, which is reflected in the expressions for optimal prices below:

$$p_i^*(f) = \frac{a(1 - \delta_{n+1}) + c}{2 - \delta_{n+1}} + f \frac{(2n + \delta_{n+1}^2)}{(2 - \delta_{n+1})(2n + \delta_{n+1})} \tag{3.6}$$
$$p_0^*(f) = \frac{a(1 - \delta_{n+1}) + c}{2 - \delta_{n+1}} + f \frac{\delta_{n+1}(3n - (n-1)\delta_{n+1})}{(2 - \delta_{n+1})(2n + \delta_{n+1})} \tag{3.7}$$

Entry by the platform leads to lower prices for consumers. First, the entry of the platform product reduces the prices of all the sellers to the optimal margin for $n + 1$

products, reflected in the first fraction in both expressions (3.6) and (3.7) above. Second, the cost advantage of the platform pushes the prices of outside sellers down even further, reflected in the second fraction in (3.6). This cost advantage also has an additional effect on the cost pass-through of outside sellers. Even though the entry increases the pass-through, it does so by less from platform entry than from outside entry due to the cost advantage of the platform product, i.e., $\frac{1}{2-\delta_n} < \frac{(2n+\delta_{n+1}^2)}{(2-\delta_{n+1})(2n+\delta_{n+1})} < \frac{1}{2-\delta_{n+1}}$. Furthermore, for the same reason, the pass-through of the platform product's price is lower than that of any third-party seller, meaning that a fee increase will shift demand from sellers to the platform product.

Lemma 3.1. *For* $f \geq \frac{a-c}{2} - \frac{(a-c)n(1-\delta_{n+1})}{2(3n-(n-1)\delta_{n+1})}$ *the price of the platform product* p_0 *is higher than the price chosen by a stand-alone monopolist* $p_m = \frac{a+c}{2}$.

Proof. The result for the threshold fee level follows from comparing the expression in (3.7) with the stand-alone monopoly price $p_m = \frac{a+c}{2}$. □

Due to strategic complementarity and the opportunity cost, the platform sets a price that is above the competitive level for a product in a market with n other substitutes. According to Lemma 3.1, if the platform does not adjust the fee after the entry, it sets a price for its own product that is even higher than the stand-alone monopoly level. In particular, this implies that the revenue per unit from its own product $p_0 - c > \frac{a-c}{2}$ is higher than the revenue per unit from third-party sellers. Thus, one might expect that the platform can improve its profits by raising the fee and shifting demand to its own product. Whether this is the case is subject of to the following section.

Optimal fee The platform operates in hybrid mode when it introduces its own product and accommodates sellers on its platform. The platform has a cost advantage because it does not have to pay the per-unit fee. This allows it to set lower prices compared to third-party sellers. According to Lemma 3.1, if the platform keeps the fee at the optimal level in pure mode in (3.5), it earns a higher revenue per unit from its own products than from the sales of third-party sellers. Although the platform's own product is more profitable, it chooses not to raise the fee and shift demand to it but rather lowers the fee to restore demand for third-party products.

Proposition 3.1. *When operating in hybrid mode, the platform sets a lower fee than when operating in pure mode.*

The platform chooses the optimal fee to maximize profits from all products and anticipates how downstream firms' prices and demand respond to the fee, i.e.,

$$\max_f \Pi_H = f \sum_i q_i \left(p_0(f), p_0(f), p_1(f), ..., p_n(f)\right) + (p_0(f) - c) q_0 \left(p_0(f), p_1(f), ..., p_n(f)\right)$$

The optimal fee, as in the pure platform case, balances the cost of raising the fee further and the benefit of collecting a higher fee from sellers' demand. Applying the envelope theorem, the marginal loss and gain of raising the fee can be expressed in general terms as follows

$$\underbrace{-\sum_i \left(f \left[\frac{dq_i}{dp_i} + \sum_{j \neq i} \frac{dq_j}{dp_i} \right] + (p_0 - c) \frac{dq_0}{dp_i} \right) \frac{\partial p_i}{\partial f}}_{\text{Marginal loss}} \stackrel{!}{=} \underbrace{\sum_i q_i}_{\text{Marginal gain}} \Leftrightarrow$$

$$\underbrace{-n \left(f \left[\frac{dq_1}{dp_1} + \sum_{i=2}^{n} \frac{dq_i}{dp_1} \right] + (p_0 - c) \frac{dq_0}{dp_1} \right) \frac{\partial p_1}{\partial f}}_{\text{Marginal loss}} \stackrel{!}{=} \underbrace{nq_1}_{\text{Marginal gain}}$$

where the equivalence between the two formulations follows from the assumption that all outside sellers are symmetric to each other and to the platform product. Thus, when the platform enters the market, it affects all existing firms symmetrically. This, allows us to simplify the argument for ease of exposition by looking at the limit case where $n = 1$ (whereas the general n case is the focus of the Appendix). From the pure platform case, we know that if another third-party firm enters the market, the fee remains unchanged - the marginal loss changes by the same amount as the marginal gain. This is not the case when the platform operates in hybrid mode.

Thus, let us look at the optimal fee in the presence of the platform product and one outside seller $n = 1$:

$$-\left(\frac{dq_1}{dp_1} f + \frac{dq_0}{dp_1} (p_0 - c) \right) \frac{\partial p_1}{\partial f} \stackrel{!}{=} q_1$$

The marginal loss from increasing the fee diminishes after the introduction of the platform product. As mentioned at the end of the previous section, the platform can shift demand to its own product, which generates higher revenues at the old fee level (3.5). Even though entry increases the cost pass-through of the outside seller, it does so to a lesser extent when the platform enters with its own product due to its fee-cost advantage. Nevertheless, the increase in the pass-through mitigates the decrease in the marginal loss.

The marginal gain from increasing the fee decreases as well. Entry by the platform leads to a stronger decrease in seller's demand q_1 due to its cost advantage compared to entry by another third-party seller. Especially at the original fee level in (3.5), this decrease is stronger than the decrease in the marginal loss. To illustrate this, reformulate the optimality condition for the fee in the presence of the platform product by using the best response function of the sellers' price w.r.t. p_0:

$$(\beta_2 f - \gamma_2(p_0 - c)) \left(\frac{1}{2} + \frac{\gamma_2}{2\beta_2} \frac{\partial p_0}{\partial f} \right) \stackrel{!}{=} \frac{1}{2} \left(\beta_2(a - f - c) - \gamma_2[(a - c) - (p_0 - c)] \right)$$

From equation (3.5) at the original fee level we know that $f = a - f - c$. Because of the fee-cost advantage, p_0 is lower than a third-party seller's price leading to a stronger decrease in demand for product 1. Even though the platform has a lower cost of raising the fee as it can divert demand to its more profitable product, the strategic response in the sellers' pass-through rate mitigates this effect.

Alternatively, one can view the platform as a multi-product monopolist aiming to maximize revenues from $n+1$ products in the market while having to set a single per-unit fee for third-party sellers.[7] Under the assumption of symmetric sellers and linear demand, this constraint is not limiting in the pure platform case, as the platform can set the fee at the stand-alone monopoly margin. Consequently, when the platform introduces its own symmetric product, it would prefer to collect the same margin from it. Particularly, if the platform can commit to pricing its own product at $p_0 = \frac{a+c}{2}$ in the future, it would maintain the fee at the pre-entry level, regardless of whether sellers possess market power or not.[8] However, as demonstrated by Lemma 3.1, in the pricing subgame, the platform prefers to deviate to a higher price due to strategic complementarity and sellers' insufficient decrease in prices in response to the platform's cost advantage.

3.3.3 Discussion

The analysis thus far reveals that when a platform operator, running a marketplace for third-party sellers, introduces its own product, it charges a lower fee to outside sellers. This adjustment aims to compensate for the insufficient reduction in sellers' prices in response to the cost advantage of the platform product. In the following section, I go into a more comprehensive discussion beyond the qualitative result of Proposition 3.1, exploring the quantitative aspect of sellers' market power, which is influenced by the number of competitors n and the strength of competition s. Additionally, I explore the platform's incentives to introduce its own product and potentially exclude sellers to monopolize their product space, examining the implications for consumers and sellers in the market. Lastly, I underscore the importance of variety effects when new products enter the market.

3.3.3.1 Competition and Market Power of Sellers

The prices of products in both the pure and the hybrid modes are higher than the monopoly price $p_m = \frac{a+c}{2}$ chosen by a vertically integrated monopolist in an independent market. Since both the platform and the sellers have market power, this results in a classic case of double marginalization. When there are more firms in the market and/or

[7]For a more comprehensive discussion on the limits of marketplace fee discrimination, refer to Tremblay (2021).

[8]Etro (2021) observes similar results and relates them to the market power of sellers that causes double marginalization. If sellers have no market power the strategic interaction is eliminated and the fee does not depend on the presence of a platform product. However, double marginalization stemming from sellers' market power can still occur if the platform can commit to setting the monopoly price for its own product.

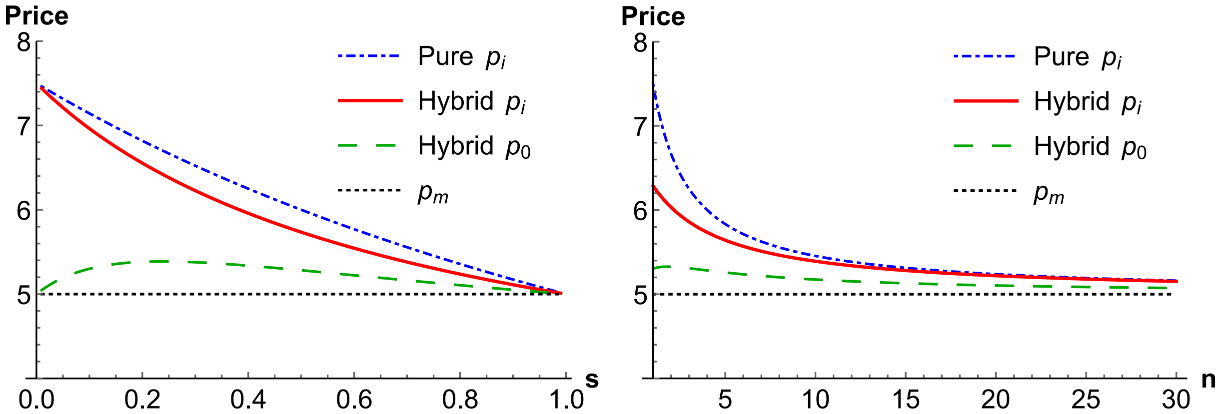

(a) Calibrated for $a = 10$, $b = 1$, $c = 0$, $n = 4$. (b) Calibrated for $a = 10$, $b = 1$, $c = 0$, $s = 0.5$.

Figure 3.1: Prices of sellers in a pure (blue, dot-dashed) and hybrid (red, solid) platform, as well as the price of the platform product in hybrid (green, dashed) and monopoly mode (black, dotted) as a function of s (a) and n (b).

the products are closer substitutes, prices for all products are lower because more competition means sellers have less market power and their prices move closer to the marginal cost $f + c$. As a result, sellers' prices are also more sensitive to cost and demand shocks, which they pass on to consumers to a greater extent. The platform's product does not have to pay the fee, but due to strategic complementarity and its opportunity cost it has, its price is also higher than the stand-alone monopoly price (as per Lemma 3.1), but is still lower than the sellers' prices. As n increases the opportunity cost of the platform increases as well so that its cost advantage vanishes as $n \to \infty$. In the limit where sellers' market power is negligible (either $n \to \infty$ or $s \to b$), all prices converge to the stand-alone monopoly price p_m. This is illustrated in Figure 3.1.

The fee reduction is less pronounced after the platform introduces its product when there is a higher number of third-party sellers already present in the market. First, more firms imply lower cost advantage, as there is more demand from which the fee can be collected. Second, as more firms imply less market power and stronger pass-through, firms adjust their prices more and their demand shrinks relatively less in response to the fee cost advantage of the platform product. As a result, the platform needs to decrease the fee by less to compensate for losses in revenues from third-party sellers' products. The result is illustrated in Figure 3.2b.

When it comes to the strength of substitution s, the amount by which the fee decreases is non-monotonic. If the products are rather weak substitutes (s close to 0), the decrease is low as the platform product does not divert much demand from sellers despite its fee-cost advantage. If the products are strong substitutes (s close to b), the diversion to the own product is strong and the sellers' market power is low. This leads to a smaller decline in the demand for third-party seller products, and thus a smaller reduction in the fee. For the extreme cases where $s = \{0, b\}$ the fee does not adjust, as the prices of third-party sellers do not react to entry. The result is illustrated in Figure 3.2a.

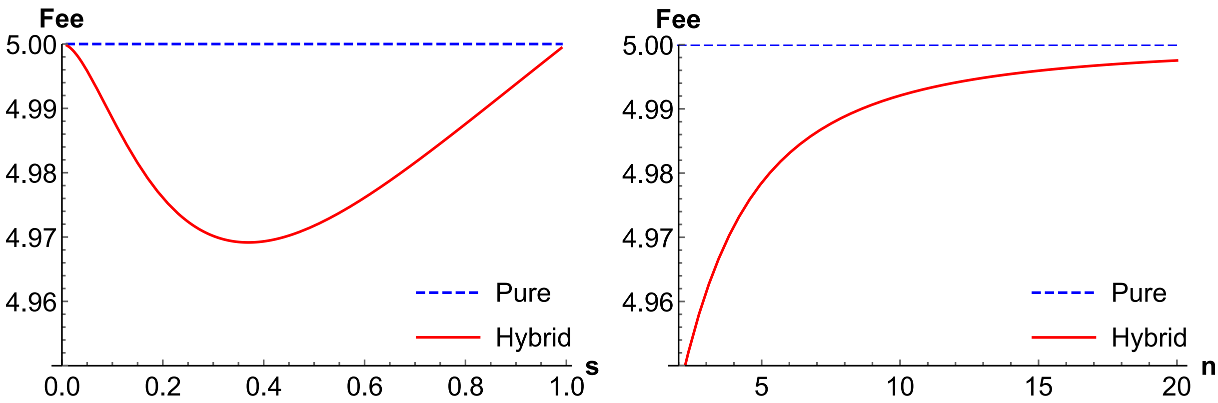

(a) Calibrated for $a = 10$, $b = 1$, $c = 0$, $n = 4$. (b) Calibrated for $a = 10$, $b = 1$, $c = 0$, $s = 0.5$.

Figure 3.2: Optimal fee in pure (blue, dashed) and hybrid (red, solid) mode for different s (a) and n (b).

3.3.3.2 The incentives of a platform to enter with its own product

Online platforms are not completely different from classical retailers - they serve as distribution intermediaries, allowing manufacturers' products to reach consumers. In the vast literature on retailing (see for example Bontems et al. (1999)) there are various different incentives as to why a retailer would want to offer its own product, also known as a private label. For example, a private label allows a retailer to increase its bargaining power vis-a-vis manufacturers and drive wholesale prices down resulting in higher profits for the retailer. However, an online platform with considerable market power has the bargaining power to dictate the terms of trade. Still, the own product is more profitable, as it avoids double marginalization and intensifies competition leading to higher demand. Thus, in the current setting:

Proposition 3.2. *The platform always has an incentive to enter with its own product.*

Proof. To see this compare the profits of operating a pure marketplace with n sellers, and operating in hybrid mode with n sellers and one own product:

$$\Pi_P = \frac{(a-c)^2 n(\beta_n - (n-1)\gamma_n)}{4(2-\delta_n)} \tag{3.8}$$

$$\Pi_H = \frac{(a-c)^2(\beta_{n+1} - n\gamma_{n+1})(2n+\delta_{n+1})(4n+2n^2-n\delta_{n+1}+\delta_{n+1}^2)}{4(8n^2-4n(n-1)\delta_{n+1}+n\delta_{n+1}^2-(n-1)\delta_{n+1}^3)} \tag{3.9}$$

First, the profit of the platform (3.8) always increases when more firms join the marketplace, as this increases the total demand in the market. Then, the profit of the platform from having another third-party seller is lower than introducing its own product. Thus, rewriting the expression in (3.8) for $n+1$ firms and comparing it with (3.9) and omitting the common multipliers reduces to showing that:

$$\frac{n+1}{2-\delta_{n+1}} \leq \frac{(2n+\delta_{n+1})(4n+2n^2-n\delta_{n+1}+\delta_{n+1}^2)}{8n^2-4n(n-1)\delta_{n+1}+n\delta_{n+1}^2-(n-1)\delta_{n+1}^3} \Leftrightarrow$$
$$0 \leq (1-\delta_{n+1})(8n^2-(n-1)\delta_{n+1}^2+4n\delta_{n+1}+\delta_{n+1}^3)$$

which is satisfied for $n \geq 1$ and $\delta_{n+1} \leq 1$. □

The intuition for this Proposition is simple. First, having more firms on the marketplace benefits the platform. In general, when a new firm joins the market, this necessarily leads to two effects, a variety effect and a competition effect. Suppose there are n firms in the market charging the price p_i^* in (3.3). If an additional firm with marginal quality index a enters and charges the same price p_i^*, this implies that total market demand $(n+1)q_i(p_i^*) = (a-p_i^*)(n+1)(\beta_{n+1}-n\gamma_{n+1})$ increases due to variety effects. However, because existing firms would react by decreasing their prices, this would further increase demand in the market - a competition effect. These two together result in more demand from which the platform can extract its per-unit fee, and thus, higher profits. Second, the platform prefers to offer its own product over accommodating an additional seller because its own product has a fee-cost advantage, which puts more pressure on sellers' prices. The platform sets an even lower price than third-party sellers, but one higher than a stand-alone monopolist and the fee. This results in more overall demand for the products offered by a hybrid platform and allows the platform to collect higher margins on the demand for its own product. Thus, the platform is always better off offering its own product than not.

The platform has a stronger incentive to enter a market where there are fewer third-party sellers. Suppose there is a fixed cost $K \sim U[0, \overline{K}]$, where $\overline{K}$ is high, that the platform must overcome to enter with a new product. A higher cost K requires a greater improvement in profits $\Pi_H - \Pi_P$. When the platform introduces its product, it both expands the market and takes market share from existing firms. Because the platform avoids paying the fee and can charge a lower price, its product is in greater demand than the products of other sellers. As seen in Section 3.3.2 and illustrated in Figure 3.1, the platform earns more revenue per unit from its own product than from sellers' products, and this revenue is higher than the stand-alone monopoly margin. The fewer sellers there are in the market, the higher the price and the higher the revenue from its own product. Thus, the lower the number of competitors n, the greater the improvement in profits $\Pi_H - \Pi_p$ for a given s. Therefore, the probability that a platform introduces its own product $\kappa = \frac{\Pi_H - \Pi_P}{\overline{K}}$ decreases with n. In addition, κ decreases with the substitutability of products s. Thus, the current framework supports the observations of Crawford et al. (2022) that Amazon is more likely to enter product spaces with fewer competitors.

3.3.3.3 The importance of variety effects and the incentives to foreclose

Due to the conflict of interest a hybrid platform constitutes, a natural question arises of whether a platform operator has the incentive to keep its marketplace open to outside sellers. In the case of Shopova (2023), the marketplace operator offers a lower quality private label product aimed at consumers with a lower willingness to pay and serves demand that was not served by the third-party seller and its higher quality product. Thus, the decision not to foreclose the seller and to differentiate the offer is linked to the provision of additional variety, which expands the market.

In the current framework, when a platform decides to foreclose sellers and monopolize the market space by becoming a vertically integrated monopolist, it enjoys profits:

$$\Pi_{mon} = \frac{(a-c)^2\beta_1}{4} = \frac{(a-c)^2}{4b} \tag{3.10}$$

This profit equals the hybrid platform profit in (3.9) when all products in the market are perfect substitutes ($s = b$). In this limit case, the prices of the sellers and the price of the platform product are equal to the marginal cost $f + c$, as explained in Section 3.3.3.1. For $f = \frac{a-c}{2}$, all prices converge to the monopoly price, and the market size becomes $(n+1)\alpha_{n+1} = a$.

As long as the products in the market are differentiated, the platform has no incentive to foreclose sellers and prefers to operate a hybrid platform. When $s < b$ each new product offered in the market, whether by an outside seller or the platform, necessarily expands the total market size and total demand as discussed in the previous Section 3.3.3.2. This also results in higher profits for the hybrid platform. This is illustrated in Figure 3.3.

Markets with fixed size and no variety effects. In the current framework, the platform always has the incentive to introduce its own product as it expands the market due to variety effects. However, if the opposite is true - each additional firm in the market does not expand the market but steals business away from existing firms - then the incentives of the platform to foreclose may change.

Shubik and Levitan (1980) suggest a linear demand system where they assume a fixed market size. They achieve this by setting the parameter $b = n - (n-1)s$ in the

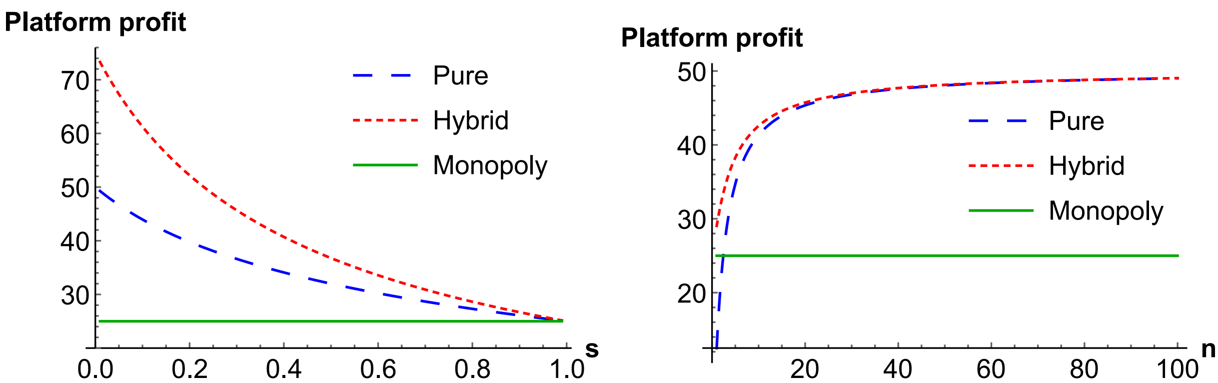

(a) Calibrated for $a = 10$, $b = 1$, $c = 0$, $n = 4$.

(b) Calibrated for $a = 10$, $b = 1$, $c = 0$, $s = 0.5$.

Figure 3.3: Platform profits in pure (blue, big dashed), hybrid (red, small dashed), and as an integrated monopolist (green, solid) for different levels of substitution s and number of third-party firms n.

utility function (3.1), thereby constraining the degree of substitutability parameter s to the interval $[0, 1]$. The total market size is $\frac{a}{b}$, and each firm holds a market share of $\alpha_n = \frac{a}{n}$ when all prices are set to zero. Under Shubik and Levitan's formulation total market demand in equilibrium is equal to $Q = \sum_i q_i = (a - p_i^*)$. As stated by Choné and Linnemer (2020), if another firm enters the market and sets the same price p_i^* total demand and consumer surplus remain unchanged. However, with more competing firms the equilibrium price p_i^* should decrease. Therefore, in the model proposed by Shubik and Levitan (1980), total demand and consumer surplus increase only due to competition effects, while variety effects are shut down.

Despite the assumption of a fixed market size, the competitive mechanisms remain the same. Having more firms in the market lowers prices and increases demand. A pure platform sets the same optimal fee as in (3.5) in pure mode and a lower fee in hybrid mode when it has to accommodate and compete with n outside sellers as a result of the same mechanism described in Proposition 3.1. Offering its own product improves the platform's profits so there is a strict incentive to do so as in Proposition (3.2). Due to double marginalization and the opportunity cost of the platform, prices are higher than the stand-alone monopoly price, and behave as in the Spence (1976) framework depicted in Figure 3.1 in Section 3.3.3.1.

In the Shubik and Levitan (1980) framework the platform prefers to be a monopolist over accommodating sellers and running a hybrid marketplace. Following Chapter 9 of Shubik and Levitan (1980), if the platform monopolizes the market and offers only its own product, it faces demand $q_1 = a - p_1$. Charging the monopoly price then yields a profit of $\frac{(a-c)^2}{4}$, which is equal to the monopoly profit in the Spence (1976) formulation in equation (3.10) for $b = 1$. Operating in hybrid mode with n outside sellers yields a profit as in (3.9) with $\delta_{n+1} = \frac{ns}{n+1-s}$ and $\beta_{n+1} - n\gamma_{n+1} = \frac{1}{n+1}$, that is lower than under monopolized market, or:

$$\begin{aligned}\Pi_H^{LS} &= \frac{(a-c)^2}{4} \frac{(2(n+1)-s)(2n^2+n(6-5s)+4(1-s)^2)}{8(n+1)^3+4(n+1)^2(5+n)s-(n+1)(16+9n)s^2+(2+n)^2s^3} \\ &< \frac{(a-c)^2}{4}\end{aligned}$$

The platform's higher profits as a vertically integrated monopolist compared to operating a hybrid or pure platform can be explained as follows. Under the assumptions made by Shubik and Levitan (1980), the market is divided among all present firms. With a vertical structure, double marginalization occurs, leading to prices higher than the stand-alone monopoly price. In contrast, a vertically integrated monopolist sets a lower price and does not share the market with other firms. Given that the market does not grow, monopolizing it is more advantageous than sharing it with other firms.

Even though Choné and Linnemer (2020) recommend the Shubik and Levitan (1980) formulation to researchers, who want to vary the number of firms, the framework proves to be problematic exactly for such purposes. Höffler (2008) highlighted that using the Shubik and Levitan (1980) approach for merger analysis may result in inaccurate estimates of consumer surplus as changing n in the utility function alters preferences. In particular, in the current model where a platform offers its own product and chooses whether to include sellers on the platform or exclude them and monopolize the market, there is a stark difference in results depending on the assumed number of products n in the utility function. In a scenario where the market size is constant and the platform's product stands alone, demand should be derived from a utility function with $b = 1$. As a result, the platform can capture the entire market as a monopolist. In another scenario, the platform could accommodate the n third-party sellers, but charge an excessive fee to discourage buyers from choosing competitors. A market of fixed size is divided equally among the $n + 1$ symmetric products. If the platform excludes sellers by raising their prices, it does not gain control over the entire market, but only a monopoly over its own $\frac{1}{n+1}$ share of the market. The latter scenario results in much lower monopoly profits for the platform than the former.[9] A possible solution to this problem is the assumption that there can be at most N varieties, of which $n < N$ are available. However, this restores the variety effects, since the market with n active firms is a fraction of the total market, and each new firm expands the market and brings it closer to its full size. In such a case, the analysis is equivalent to that using the Spence (1976) framework.

3.3.3.4 Consumer, third-party sellers, and total welfare

Consumers are always better off when the platform introduces its own product. First, because of variety effects, there is demand that was not covered before - the market expansion increases consumer welfare. Second, the entry of the platform with its own product puts competitive pressure on sellers and leads to lower prices, which further expands total demand and in this way, increases consumer surplus.

While the platform lacks incentives to foreclose sellers, consumers might prefer a monopolized marketplace to a hybrid platform in the case of one outside seller ($n = 1$) and the products being close substitutes (s close to b). The reason is that both the prices of the third-party seller product and the platform product are above the monopoly level, as discussed in Section 3.3.3.1. When the products are close substitutes, the variety effects are not strong enough to compensate for the excessively high prices. This leaves consumers worse off with two highly-priced products compared to having one product priced at the monopoly level. For $n \geq 2$, variety effects are always stronger, and consumers are always better off when the platform chooses to operate in hybrid mode rather than monopolizing the market. This is illustrated by Figure 3.4.

Outside sellers are always worse off when the platform offers its product. In the pure platform case, rival product entry reduces their prices and demand, leading to decreased profits. The platform's fee cost advantage further lowers their prices and demand, exacerbating the decrease in profits.

[9] For an explicit discussion, see Appendix 3.B.

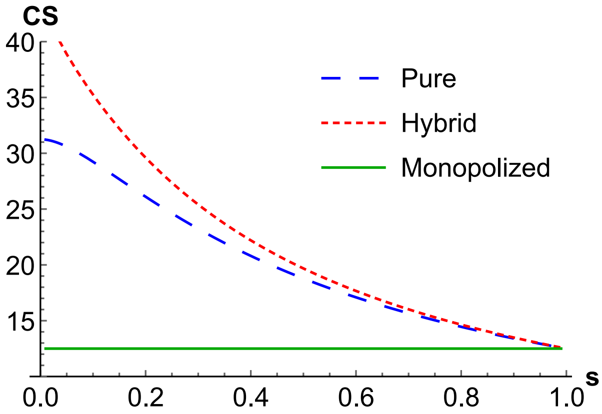

(a) Calibrated for $a = 10$, $b = 1$, $c = 0$, $n = 10$.

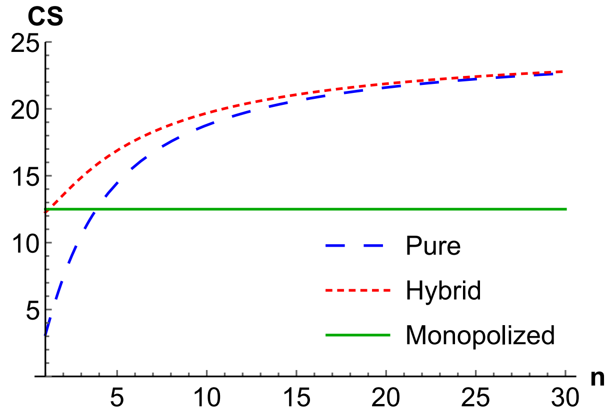

(b) Calibrated for $a = 10$, $b = 1$, $c = 0$, $s = 0.5$.

Figure 3.4: Consumer surplus in pure (blue, big dashed), hybrid (red, small dashed), and under monopoly (green, solid) for different levels of substitution s and number of third-party firms n.

Finally, total welfare increases due to both variety and competition effects. Besides increasing consumer surplus and capturing new demand through market expansion, the introduction of the platform product also leads to a stronger decrease in double marginalization, resulting in an overall increase in total welfare.

3.4 Conclusion

This paper expands the framework of Shopova (2023) to include horizontal differentiation with multiple sellers and shows that the fee in hybrid mode is lower than in pure mode. The main driver behind this decrease is the insufficient reaction of the sellers with market power to the cost advantage of the platform. Thus, the platform decreases the fee to shift back demand to sellers' products. Sellers with less market power and in more competitive markets tend to adjust their prices more, resulting in a weaker need for the platform to decrease the fee. In addition, the platform has no incentive to exclude sellers and monopolize their product space, as each variety, including the product of the platform, has a market expansion effect. These results are in line with findings made by Crawford et al. (2022) on Amazon's entry strategy and effect on sellers' prices.

As the main objective of the paper is to analyze the competitive effects of a platform introducing its own product on fees and prices, it abstracts from other relevant issues related to the policy discussion on hybrid platforms. In particular, this paper does not address the role of asymmetric information with respect to costs (Kang and Muir (2022)) or demand (Madsen and Vellodi (2022), Lam and Liu (2021)), and it does not explore the role of data (e.g. Madsen and Vellodi (2022), Padilla et al. (2022)), nor network effects (e.g. Gautier et al. (2021), Pouyet and Thomas (2021)). Additional research is needed to assess the interplay between these distinct aspects and the findings presented in this paper.

3.A Omitted Proofs

Proof of Proposition 3.1. Applying the Envelope Theorem to the optimization problem of the platform, one can express the first order condition w.r.t. the fee as:

$$\underbrace{-\sum_i \left(f\left[\frac{dq_i}{dp_i} + \sum_{j\neq i}\frac{dq_j}{dp_i}\right] + (p_0 - c)\frac{dq_0}{dp_i}\right)\frac{\partial p_i}{\partial f}}_{\text{Marginal loss}} \stackrel{!}{=} \underbrace{\sum_i q_i}_{\text{Marginal gain}}$$

The marginal loss of increasing the fee decreases after the introduction of the own product. To see this write out the expression for the marginal loss:

$$\begin{aligned} n\left[f(\beta_{n+1} - (n-1)\gamma_{n+1}) - \gamma_{n+1}(p_0 - c)\right]\frac{2n + \delta_{n+1}^2}{(2-\delta_{n+1})(2n+\delta_{n+1})} &= \\ = n\left[f(\beta_{n+1} - n\gamma_{n+1}) - \gamma_{n+1}(p_0 - c - f)\right]\frac{2n + \delta_{n+1}^2}{(2-\delta_{n+1})(2n+\delta_{n+1})} & \end{aligned} \tag{3.A.1}$$

Let us compare the marginal loss in (3.4) and in (3.A.1). First, as the demand for each seller shrinks in the presence of the platform product, the marginal loss of increasing the fee decreases, i.e., $\beta_{n+1} - n\gamma_{n+1} \leq \beta_n - (n-1)\gamma_n$. Second, as Lemma 3.1 shows, at the original fee level in pure mode, the platform product generates a higher revenue per unit than the fee. Thus, since the platform can divert demand to its own more profitable product, this further reduces the cost of increasing the fee. Some of this reduction is offset by the increased pass-through of seller prices. However, because pass-through increases less in the presence of a platform product than in the presence of another third-party seller, the overall effect is to reduce the marginal loss from raising the fee.[10]

The marginal gain of increasing the fee shrinks more than the marginal loss. Consider the expression for the marginal gain:

$$nq_i = n(\beta_{n+1} - n\gamma_{n+1})\left(\frac{a-f-c}{2-\delta_{n+1}} - f\frac{\delta_{n+1}}{(2-\delta_{n+1})(2n+\delta_{n+1})}\right) \tag{3.A.2}$$

Comparing the marginal gains in (3.4) and in (3.A.2), it is straightforward to see that total demand for sellers' products has decreased - first because the entry of a new rival product shrinks individual demand, and second, because the platform product has a fee cost advantage which allows it to set a lower price. Due to the market power of sellers and the incomplete pass-through of their prices, this cost advantage is not fully passed on to consumers, resulting in an even stronger decrease in demand for their products.

Finally, to see that the marginal gain of increasing the fee beyond the level in the pure platform case in (3.5) has decreased more than the marginal loss, compare (3.A.1) and (3.A.2) evaluated at $f = \frac{a-c}{2}$:

[10] Technically, even if pass-through had increased to $\frac{1}{2-\delta_{n+1}}$, the marginal loss would still have been lower.

$$\frac{(a-c)n(1-\delta_{n+1})\left(2n+\delta_{n+1}^2\right)\left(2n-(n-1)\delta_{n+1}+\delta_{n+1}^2\right)}{(2n+\delta_{n+1})^2(2-\delta_{n+1})^2} > \frac{(a-c)n^2\beta^2(1-\delta_{n+1})}{(2n+\delta_{n+1})(2-\delta_{n+1})}$$
$$\Leftrightarrow \frac{(a-c)n\beta_{n+1}(n+\delta_{n+1})(1-\delta_{n+1})^2}{(2n+\delta_{n+1})^2(2-\delta_{n+1})^2} \geq 0$$

The optimal fee in hybrid mode in reduced form is thus given by

$$f = \frac{a-c}{2}\frac{(2\beta_{n+1}+\gamma_{n+1})(4\beta_{n+1}^2-2n\beta_{n+1}\gamma_{n+1}+n\gamma_{n+1}^2)}{(8\beta_{n+1}^3-4(n-1)\beta_{n+1}^2\gamma_{n+1}+n\beta_{n+1}\gamma_{n+1}^2-(n-1)n\gamma_{n+1}^3)}$$

where the first fraction is exactly equal to the original fee in (3.5), while the second is less than one as $\beta_{n+1}-n\gamma_{n+1} \geq 0$. $\square$

3.B Modeling Firm Entry with a QQUM

This Section highlights the properties of quasi-linear quadratic utility (QQU) functions and how they are used to derive a linear demand system with an emphasis on the number of products in the market. Take the Spence (1976) formulation used in Section 3.2:

$$\max_q \mathcal{L} = \sum_i (a-p_i)q_i - s\sum_i\sum_{j>i} q_i q_j - \frac{b}{2}\sum_i q_i^2 \tag{3.B.3}$$

A consumer taking all prices as given optimizes its consumption so that it consumes non-negative amounts of each variety $q_i \geq 0$. This leads to the Karush-Kuhn-Tucker (KKT) conditions, namely:

$$q_i\frac{\partial\mathcal{L}}{\partial q_i} \stackrel{!}{=} 0 \Leftrightarrow q_i(a-p_i-bq_i-s\sum_{j\neq i} q_j) \stackrel{!}{=} 0 \tag{3.B.4}$$

If prices are not prohibitively high, so that all $q_i > 0$ one gets indirect demand for all i

$$p_i = a - bq_i - s\sum_{j\neq i} q_j \tag{3.B.5}$$

which can be inverted to direct demand

$$q_i = \frac{a}{(b+(n-1)s)} - \frac{b+(n-2)s}{(b-s)(b+(n-1)s)}p_i + \frac{s}{(b-s)(b+(n-1)s)}\sum_{j\neq i} p_j \tag{3.B.6}$$

When all prices are set to zero, the demand for a single product is $\frac{a}{b+(n-1)s}$ and the total demand for all products is $\frac{na}{b+(n-1)s}$, which increases in n.

Suppose that the price of all products p_j for $j = 2, ..., n$ are so high that $q_j = 0$. From the KKT condition (3.B.4), this implies that the only F.O.C. left is:

$$q_1(a-p_1-bq_1-s\sum_{j\neq 1} q_j) = 0 \Leftrightarrow a-p_1-bq_1 = 0 \tag{3.B.7}$$

$$\Leftrightarrow p_1 = a - bq_1 \Leftrightarrow q_1 = \frac{a-p_1}{b} \tag{3.B.8}$$

In this scenario, the prices of all products except product 1 are prohibitively high, the firm that controls product 1 is a monopolist.

In their formulation Shubik and Levitan (1980) assume that the total market size/demand at zero prices is fixed at a for any number of firms n, and achieve this by setting $b = n - (n-1)s$. This assumption implies that having more firms splits the market, but does not grow the market. Choné and Linnemer (2020) call this "shutting down variety effects". In the fully symmetrical case, all firms in a market set a price p^*, so one can easily verify that the utility of a consumer is given by $\frac{a-p^*}{2}$ and is fully independent of n. The only difference entry in such a framework would make is to decrease prices through competition effects.

However, as pointed out by Höffler (2008), researchers used to make a mistake by adjusting the number of sellers (n) in a market after excluding one of them. This adjustment automatically changes the initial utility function by altering the parameter b. For instance, if one looks at the optimization problem in (3.B.3) and (3.B.4), if from the beginning there was only one firm in the market with $b = 1$, then demand for this firm is $q_1 = a - p_1$. In the Shubik and Levitan (1980) specification the same result holds. If you have from the beginning 2 firms and the price of firm 2 is prohibitively high, in Spence demand for product 1 is still $q_1 = a - p_1$. However, in Shubik and Levitan (1980) having two firms in the market would mean $b = 2 - s$ and would imply that for a prohibitively high p_2 demand for product 1 should be $q_1 = \frac{a-p_1}{2-s}$ if one optimizes according to the same utility function. Thus, by varying n in Shubik and Levitan (1980) framework, one does not just vary the number of firms, but also the parameter b, which then changes the utility function.

One potential solution is to set N as the maximum possible number of varieties, with $b = N - (N-1)s$, and let the number of products in the market, n, approach the value of N to achieve convergence to the market size a. However, this restores the variety effects in Shubik and Levitan (1980), so that if $n < N$, then consumer surplus and total demand will depend on n for given p^*, and the qualitative results won't be different from the ones in Spence (1976).

Bibliography

Amir, R., Erickson, P., & Jin, J. (2017). On the microeconomic foundations of linear demand for differentiated products. *Journal of Economic Theory*, *169*, 641–665.

Anderson, S., & Bedre-Defolie, Ö. (2022). Online trade platforms: Hosting, selling, or both? *International Journal of Industrial Organization*, *84*, 102861.

Anderson, S. P., & Bedre-Defolie, Ö. (2023). *Hybrid platform model: Monopolistic competition and a dominant firm* (tech. rep.). The RAND Journal of Economics.

Arya, A., Mittendorf, B., & Sappington, D. E. (2007). The bright side of supplier encroachment. *Marketing Science*, *26*(5), 651–659.

Arya, A., Mittendorf, B., & Sappington, D. E. (2008). Outsourcing, vertical integration, and price vs. quantity competition. *International Journal of Industrial Organization*, *26*(1), 1–16.

Bontems, P., Monier-Dilhan, S., & Requillart, V. (1999). Strategic effects of private labels. *European Review of Agricultural Economics*, *26*(2), 147–165.

Cai, G. G. (2010). Channel selection and coordination in dual-channel supply chains. *Journal of Retailing*, *86*(1), 22–36.

Choné, P., & Linnemer, L. (2020). Linear demand systems for differentiated goods: Overview and user's guide. *International Journal of Industrial Organization*, *73*, 102663.

Crawford, G. S., Courthoud, M., Seibel, R., & Zuzek, S. (2022). *Amazon entry on amazon marketplace*. Centre for Economic Policy Research.

Das Varma, G., & De Stefano, M. (2020). Equilibrium analysis of vertical mergers. *The Antitrust Bulletin*, *65*(3), 445–458.

Domnenko, G. B., & Sibley, D. S. (2020). Simulating vertical mergers and the vertical guppi approach. *Available at SSRN 3606641*.

Etro, F. (2021). Product selection in online marketplaces. *Journal of Economics & Management Strategy*, *30*(3), 614–637.

Etro, F. (2023a). The economics of amazon. *Available at SSRN 4307213*.

Etro, F. (2023b). Hybrid marketplaces with free entry of sellers. *Review of Industrial Organization*, 1–30.

Gautier, A., Madio, L., & Shekhar, S. (2021). *Platform duality and network externalities* (Working paper).

Ha, A. Y., Tong, S., & Wang, Y. (2022). Channel structures of online retail platforms. *Manufacturing & service operations management*, *24*(3), 1547–1561.

Höffler, F. (2008). On the consistent use of linear demand systems if not all varieties are available. *Economics Bulletin*, *4*(14), 1–5.

Kang, Z. Y., & Muir, E. V. (2022). Contracting and vertical control by a dominant platform. *Unpublished manuscript, Stanford University.*

Lam, W. M. W., & Liu, X. (2021). *Data usage and strategic pricing: Does platform entry benefit independent traders?* (Working paper).

Madsen, E., & Vellodi, N. (2022). Insider imitation. *Available at SSRN 3832712*.

Padilla, J., Piccolo, S., & Vasconcelos, H. (2022). Business models, consumer data and privacy in platform markets. *Journal of Industrial and Business Economics*, 1–36.

Pouyet, J., & Thomas, T. (2021). The competitive effects of vertical integration in platform markets.

Ritz, R. A. (2018). Oligopolistic competition and welfare. In *Handbook of game theory and industrial organization, volume i* (pp. 181–200). Edward Elgar Publishing.

Salinger, M. A. (2021). The new vertical merger guidelines: Muddying the waters. *Review of industrial organization*, *59*(2), 161–176.

Shopova, R. (2023). Private labels in marketplaces. *International Journal of Industrial Organization*, 102949. https://doi.org/https://doi.org/10.1016/j.ijindorg.2023.102949

Shubik, M., & Levitan, R. (1980). *Market structure and behavior*. Harvard University Press.

Spence, M. (1976). Product differentiation and welfare. *The American Economic Review*, *66*(2), 407–414.

Tremblay, M. J. (2021). The limits of marketplace fee discrimination.

Tsay, A. A., & Agrawal, N. (2004). Channel conflict and coordination in the e-commerce age. *Production and operations management*, *13*(1), 93–110.

Weyl, E. G., & Fabinger, M. (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy*, *121*(3), 528–583.

Zennyo, Y. (2022). Platform encroachment and own-content bias. *The Journal of Industrial Economics*, *70*(3), 684–710.

Zhu, F., & Liu, Q. (2018). Competing with complementors: An empirical look at amazon.com. *Strategic Management Journal*, *39*(10), 2618–2642.

Abstract

This dissertation contributes to the field of industrial organization by examining three key topics. The first chapter, joint work with Maarten Janssen and Daniel Garcia, focuses on markets where firms sell time-dated products and possess private information about unsold capacities. We show that competition under private information explains increased price dispersion and higher expected prices as the deadline approaches. Private information limits firms' market power, which creates incentives for unilateral disclosure and espionage, which then increases firms' profits.

The second chapter addresses regulators' concern about the distortion of competition by online marketplace operators when they introduce their own private label products. The chapter highlights how online marketplaces differ from traditional retailers in pricing dynamics. The introduction of a private label does not incentivize marketplace operators to distort competition or exclude outside sellers. On the contrary, operators reduce fees charged to outside sellers and vertically differentiate their own products to protect the seller's channel. Nevertheless, compared to traditional retailers, online marketplaces offer lower-quality products at higher consumer prices, potentially resulting in less improvement in consumer surplus.

The third chapter builds upon the findings of the second chapter by examining a horizontal differentiation setting with multiple firms. By studying this setting, the chapter further validates the conclusions drawn in the second chapter, by further emphasizing the importance of both variety effects and the level of competition in the decision to reduce the fee to outside sellers and accommodate them.

In summary, this thesis explores markets with private information, the impact of private labels in online marketplaces, and the role of variety effects in the platform's decisions to foreclose outside firms. These findings enhance understanding of competition and market dynamics, providing valuable insights for researchers and policymakers.

Zusammenfassung

Diese Dissertation trägt zum Bereich der Industrieorganisation bei, indem sie drei zentrale Themen untersucht. Das erste Kapitel, zusammen mit Maarten Janssen und Daniel Garcia verfasst, konzentriert sich auf Märkte, auf denen Unternehmen zeitdatierte Produkte verkaufen und über private Informationen zu unverkauften Kapazitäten verfügen. Wir zeigen, dass der Wettbewerb unter privaten Informationen eine erhöhte Preisstreuung und höhere erwartete Preise mit sich bringt, wenn der Stichtag näher rückt. Private Informationen begrenzen die Marktmacht der Unternehmen, was Anreize für einseitige Offenlegung und Spionage schafft, die wiederum die Gewinne der Unternehmen steigern.

Das zweite Kapitel befasst sich mit den Bedenken der Regulierungsbehörden hinsichtlich der Wettbewerbsverzerrung durch Betreiber von Online-Marktplätzen, wenn sie Eigenmarkenprodukte einführen. Das Kapitel zeigt auf, wie sich Online-Marktplätze in Bezug auf Preisdynamik von traditionellen Einzelhändlern unterscheiden. Die Einführung einer Eigenmarke veranlasst Marktplatzbetreiber nicht, den Wettbewerb zu verzerren oder externe Verkäufer auszuschließen. Im Gegenteil, die Betreiber reduzieren die Gebühren für externe Verkäufer und differenzieren ihr eigenes Produkt vertikal, um den Vertriebskanal des Verkäufers zu schützen. Dennoch bieten Online-Marktplätze im Vergleich zu traditionellen Einzelhändlern Produkte von geringerer Qualität zu höheren Preisen an, was möglicherweise zu einer geringeren Verbesserung der Konsumentenrente führt.

Das dritte Kapitel baut auf den Ergebnissen des zweiten Kapitels auf, indem es eine Einstellung der horizontalen Differenzierung mit mehreren Unternehmen untersucht. Durch die Untersuchung dieser Einstellung validiert das Kapitel die in Kapitel zwei gezogenen Schlussfolgerungen weiter, indem es die Bedeutung von Vielfalteffekten und dem Wettbewerbsniveau in der Entscheidung, die Gebühren für externe Verkäufer zu reduzieren und sie zu unterstützen, weiter betont.

Zusammenfassend untersucht diese Arbeit Märkte mit privaten Informationen, den Einfluss von Eigenmarken in Online-Marktplätzen und die Rolle von Vielfalteffekten bei den Entscheidungen der Plattform, externe Unternehmen auszuschließen. Diese Erkenntnisse verbessern das Verständnis von Wettbewerb und Marktdynamik und liefern wertvolle Einblicke für ForscherInnen und EntscheidungsträgerInnen.