Perfect Bayesian Equilibria in Repeated Sales with Multiple Buyers

Kaito Nakatani

May 7, 2021

Version: Final Draft

Advisor: Professor Dianna Xu

Abstract

In an auction with repeated sales of the same item, sellers try to learn buyers' values of the item and buyers try to hide it. At a Perfect Bayesian Equilibrium, sellers maximize their profit against buyers' strategies and buyers maximize their utilities against sellers' strategies. This thesis first explores the Perfect Bayesian Equilibrium structures in different auction settings that were studied in existing publications. Then, we add our original work to extend the current understanding of Perfect Bayesian Equilibrium of repeated sales to the multi-buyer settings where buyers' values for items are drawn from different distributions. We prove that under certain conditions, there is a Perfect Bayesian Equilibrium between a seller and two buyers with arbitrary distributions.

ii

Acknowledgements

First, I would like to thank my advisor, Professor Dianna Xu at Bryn Mawr College for agreeing to be my thesis advisor, letting me choose a research topic of my choice, giving me insightful advice and feedback even the topic is outside of her research area, and, most of all, allowing me to spend time on job search so that I do not have to worry about my senior thesis in April. Without her generosity, I would not have been able to figure out my post-graduation plan. I would also like to thank Professor Giri Parameswaran from Economics Department at Haverford College for giving me advice on how to study the background for auction theory. Furthermore, I would like to thank professors in the computer science department at Bi-Co for providing me with both valuable feedback on this thesis and great courses for the past three years, especially Professor Sorelle Friedler and Professor Dave Wonnacott. I would like to thank my family as well who made it possible for me to study and succeed at Haverford. Finally, I would like to show my gratitude to my friends, especially Shreya Kishore, who made me feel welcome at Haverford and supported me throughout my undergraduate career.

iii

1

Introduction

1.1 Introduction and Motivation

In this era of COVID-19 pandemic, more people use online platforms for purchasing goods (Charm et al. 2020). As a result, repeated sales is becoming a common auction format more and more in our daily lives. Some people purchase their daily goods from Amazon Prime every month; some people get a delivery for groceries every week from their favorite supermarkets. In this circumstance where buyers meet the same sellers repeatedly, we can easily imagine that sellers try to learn buyers' behavior to increase their profits. However, strategic buyers will try to exploit this behavior of sellers so that the seller's strategies work in their own favor.

In this thesis, using a game-theoretical perspective, we investigate the equilibrium structure in repeated sales. Economics literature have investigated the equilibrium structure in an one-round version of the auction (Karlin and Peres 2017) (Krishna 2009). Furthermore, a recent study has explored the equilibrium structure of n-round sales when there is a single buyer and a single seller (Devanur et al. 2015). In this thesis, we extend the current result to discuss the equilibrium structure of repeated sales of a single item when there are multiple buyers and a single seller. In this extension, we will mainly use the techniques used in (Immorlica et al. 2017).

Following this introductory paragraph, we conduct a literature review of the research area in Chapter 2. Specifically, this paper deals with a topic in algorithmic game theory, which requires knowledge of economics. For unfamiliar readers, in Section 2.1, we introduce the necessary economics background to understand the setting of this study. Then, in Section 2.2, we introduce and discuss previous literature in the area, followed by a more detailed discussion of the selected previous studies in

Section 2.3. In Chapter 3, we add our original work to extend the current work. We first propose our research question and then describe our method in Section 3.1. We then explore the method and prove the existence of an equilibrium under certain set of assumptions in Section 3.2, 3.3, 3.4, and 3.5. Finally, we conclude the paper and give some directions for future researches in Chapter 4.[1]

[1]This work was completed as a part of senior thesis that is required to complete computer science major at Haverford College.

3

2

Literature Review

In this chapter, we conduct a literature review on the equilibrium structures in repeated sales by introducing necessary background, a survey of the relevant literature, and details for the closely related studies.

2.1 Background

In this section, we will introduce the necessary background in algorithmic game theory to understand the previous studies in Section 2.2. This section is heavily based on the motivational example (Example 2.1.1) in (Devanur et al. 2015) and the information explained in (Karlin and Peres 2017) and (Krishna 2009). While this is a computer science thesis, this topic requires a lot of economics background. First, in Subsection 2.1.1, we study the settings of the study in this thesis. Then, in Subsection 2.1.2, we discuss how to analyze the strategies taken by buyers and sellers in an auction.

2.1.1 The Setting of the Study

In this subsection, we define terms required to understand the setting of the study in this paper. This section heavily relies on (Karlin and Peres 2017) for general algorithmic game theoretical concepts and (Devanur et al. 2015) for more specific background in repeated sales.

We begin by giving a motivational example using the problem illustrated in (Devanur et al. 2015).

Example 2.1.1 (Devanur et al. 2015, The Fishmonger Problem). Let s be a seller who sells a fresh fish to a buyer b for n days. Suppose b values the fish at value v, a random value in $[0, 1]$. On a given day i, without knowing this evaluation v, s sets the price p_i. In this scenario, we are interested in the following question: "What are the best strategies for the buyer and the seller if they are trying to achieve respectively

$$b : \max_{I_i} \sum_i I_i(v - p_i)$$

$$s : \max_{p_i} \sum_i I_i p_i$$

where I_i is an indicator random variable whether b buys the fish on day i?" In other words, when should the buyer b buy the fish to maximize their value gained and what price should the seller s set to maximize their profits each day? Naturally, the seller s will try to learn the buyer b's evaluation v of the fish over the course of n days. On the contrary, the buyer b will try to exploit this seller s's behavior.

The solution to this example can be applied to many different real-world markets. We will revisit Example 2.1.1 in Example 2.1.12 to find the solution. However, we will first study the necessary background to formally define the problem and the solution scheme.

This Fishmonger Problem is a problem in the field called auction theory. There are many different types of prevalent auctions such as English auctions where prices increase until there is only one bidder left or Dutch auctions where prices decrease until there is some bidder who wants to buy the product (Krishna 2009). While it is hard to write down a universal definition of what an auction is, there are always (at least) two agents in auctions: a buyer and a seller. Attempting to quantify these two agents, we begin by defining a basic economic concept, utility.

Definition 2.1.2. (Debreu 1954) Let A be an agent of interest and X be the set of possible result of an action. The function $u : X \to \mathbb{R}$ is a *utility function* if for all $x, y \in X$, A prefers x to y if and only if $u(x) > u(y)$.

While Definition 2.1.2 requires the definition of preference, we leave the precise definition to (Debreu 1954). For now, we assume that preference is the linear ordering of all possible actions.[1]

[1] Linear order is defined by reflexivity, anti-symmetry, and transitivity.

Definition 2.1.2 indicates that the utility is simply the ordered representation of preferences. We will illustrate a simple economic principle in the following example.

Example 2.1.3. Suppose an agent A has a utility function of $u(x) = -x^2 + 8x + 1$ where x represents the number of goods purchased. We can see that when A buys 4 goods, the utility is maximized, and therefore that state is most preferred. If we try to solve this economically, we can consider the marginal utility of buying another good, which can be obtained by

$$\frac{\partial u(x)}{\partial x} = -2x + 8.$$

With an implicit assumption of $u(\emptyset) = 0$, we can determine that at $x = 4$, the marginal utility of buying another good goes below the utility of doing nothing. This illustrates the foundational concept of resource allocation based on marginal utilities.

In this paper, we assume that utility is linear to the value gained. Using the definition of utilities, we can now define an auction by defining the motivations of two agents:

Definition 2.1.4. Suppose seller(s) are trying to sell item(s) to buyer(s). An *auction* satisfies the following properties:

1. Buyer(s) try to maximize their utilities;

2. Seller(s) try to maximize their profits.

Based on this definition, we consider the strategies taken by buyers and sellers to achieve their goals under design-specific settings. We can now define the auction-specific terms to specify our setting of the study. We first define the private value.

Definition 2.1.5 (Karlin and Peres 2017, Definition 14.1.2)**.** Let b be a buyer in an auction. The *value* that the buyer b holds for an item i is the evaluation of how much i is worth to themselves and is drawn from a publicly known distribution $\mathcal{F}$. This value is denoted by v_i (or v when there is only one item). We call v_i *private* if it is hidden from other parties in the auction.

Assuming utility is linear to the value, this enables us to obtain buyers' utility given their utility functions.

Remark 2.1.6. Suppose the utility function equals to the value gained in an auction. The utility corresponding to the item i with a price p_i is

$$u(i) = \begin{cases} v_i - p_i & \text{if the buyer buys the item } i \\ 0 & \text{otherwise} \end{cases}.$$

We use this result throughout the paper. Now we have enough tools to explain the details of our study.

Definition 2.1.7 (Basic Setting)**.** We define the single-item repeated-sales posted-price auction as follows: let s be a seller of a single item i and β be the set of buyers. Suppose the auction consists of n rounds. Each buyer in β holds the private value that is fixed for all n rounds. The auction proceeds as follows each round:

1. The seller s decides the price p for the round.
2. The buyers decide if they accept the offer or not.
3. If there are multiple buyers who accept the offer, the item i will be randomly given to someone among those accepted.

In this thesis, we will refer to the setting in Definition 2.1.7 as the basic setting. Recall from Definition 2.1.4 that sellers decide the price and buyers decide their strategies to maximize their utilities or profits respectively. Now that we understand the setting of our study, we move on to discuss the optimal strategies in auctions.

2.1.2 Equilibrium Structures

We now have enough tools to talk about the equilibrium structures in an auction. In this paper, we mainly use the concept of Perfect Bayesian Equilibrium (PBE). Intuitively, in PBE, the strategy of one party should be optimal given the expected strategy of the other party. We define it formally as follows:

Definition 2.1.8 (Devanur et al. 2015, p.6)**.** Let s be a seller and b be a buyer of the basic setting. Suppose the seller has their belief μ_v about the buyer's private value v. The auction is at *Perfect Bayesian Equilibrium* if the strategies of s and b satisfy the following condition:

1. The seller assumes that the buyer's strategy is the PBE strategy. After every round, the seller updates μ_v using the Bayes' rule;

2. After every round of the game, given the belief and the current history of the strategy, the buyers maximize their expected utilities and the sellers maximize their expected profits against the PBE strategy by the other agent.[2]

Our study focuses on obtaining PBE in different auction settings. However, there are so many strategies that can be taken. To limit the domain of the strategies, we define the following classes of strategies and PBE.

Definition 2.1.9 (Karlin and Peres 2017, Section 2.2)**.** Consider any auction. If the strategy taken by an agent is deterministic (not probabilistic), we call the strategy a *pure strategy*.

Definition 2.1.10 (Devanur et al. 2015, p.6)**.** Consider an auction and let b be a buyer of the auction. Suppose b decides whether to buy or not based on a threshold t i.e. b buys the item if $v \geq t$ and rejects if $v < t$. We say that b is following a *threshold strategy*. Note that a threshold strategy is always a pure strategy.

Definition 2.1.11 (Devanur et al. 2015, p.6)**.** Given an auction, suppose a Perfect Bayesian Equilibrium is achieved with a threshold strategy by the buyer and a pure strategy by the seller. Then, we call this PBE a *pure strategy threshold PBE*.

Definition 2.1.9 and 2.1.10 limit the domain of the buyers' and sellers' strategies. Therefore, only considering a pure strategy threshold PBE can simplify our exploration of a PBE. We now illustrate a pure strategy threshold PBE in the following example by revisiting the fishmonger example (Example 2.1.1).

Example 2.1.12 (Fishmonger Problem, Revisited). We note that this example fits the basic setting described in Definition 2.1.7, a repeated sales with a single item. Here, there is only a single buyer i.e. $\beta = \{b\}$ and the private value v of b is drawn from $\mathcal{D} \sim U[0,1]$.

1. First let us consider the case where the auction only goes for one day. That is, $n = 1$. Consider b's strategy. Since the utility U_b can be calculated by

$$U_b = I_1(v - p), \tag{2.1.1}$$

it is clear that

[2]This condition is often called Bayes-Nash Equilibrium or Subgame Perfect Equilibrium in the literature. However, we omit introducing another term for the sake of simplicity.

a) $I_1 = 1$ (i.e. b buys the item) when $v \geq p$ (in which case $U_b = v - p$); and

b) $I_1 = 0$ (i.e. b does not buy the item) when $v < p$ (in which case $U_b = 0$).

This is true regardless of the seller's strategy since the seller will have no time to update their beliefs when $n = 1$.

We now consider s's strategy. We can calculate the expected profit P_s by

$$P_s = E[I_1]p = (1 - p)p. \tag{2.1.2}$$

This is because s expects the strategy above and we can obtain

$$E[I_1] = P(v \geq p) = 1 - p$$

since v is drawn from $U[0, 1]$. Therefore, P_s can be maximized when $p = \frac{1}{2}$. Since there is no update of belief required, we can easily verify that this is at PBE for $n = 1$ case.

2. We now consider $n = 2$. First consider modifying the strategy above with the assumption that b's strategy stays similar. Then, after a similar derivation with a Bayesian update of the belief, we obtain the following strategy:

a) s sets $p = \frac{1}{2}$ for Day 1; and

b) s sets $p = \frac{1}{4}$ if $I_1 = 0$ and $p = \frac{3}{4}$ otherwise for Day 2.

However, we now show that this set of strategy does not give us a PBE. Suppose $v = \frac{5}{8}$. The strategy above gives b the total utility of $\frac{1}{8}$ since b buys the item on Day 1 and does not buy the item on Day 2. However, by not buying the item on Day 1, b can gain the total utility of $\frac{5}{8} - \frac{1}{4} = \frac{3}{8}$ after the second round. Therefore, we can conclude that the game is not at PBE with this set of strategy.

Instead, we now consider the following set of strategy $\sigma = (\sigma_s, \sigma_b)$ where σ_s is the strategy for the seller and σ_b is the strategy for the buyer;

σ_s : In round 1, the seller sets $p_1 = 0.3$. If $I_1 = 1$, then $p_2 = 0.6$. Otherwise, $p_2 = 0.3$.

σ_b : In round 1, if $v \geq 2p_1$ and $p_1 < 0.5$, buy the item. Otherwise, don't buy the item. In round 2, buy the item if $v \geq p_1$.

This auction is indeed at Perfect Bayesian Equilibrium and we prove it in the following proof.

Proof that σ is at PBE. Recall from Definition 2.1.8 that we need to maximize the expected utility for the buyer and the expected profit for the seller.

1. First show that the strategy maximizes the expected utility at each round for the buyer b against σ_s given the value v, beliefs, and the current history of the strategies. For the second round, this choice is optimal as seen in the case for the $n = 1$ case. It remains to show that the strategy is optimal for the first round. Let t be the threshold. Note that we can easily obtain that

   a) given $v \leq 0.3$, b should not buy the item on the first day to achieve an optimality;

   b) given $v \geq 0.6$, b should buy the item on the first day to achieve an optimality;

   c) given $0.3 \leq v \leq 0.6$, b is indifferent between buying and not buying based on σ_s.

   Therefore, we need to have a threshold of

   $$0.3 = p_1 \leq t \leq 2p_1 = 0.6.$$

   Thus, σ_b maximizes the expected utility against σ_s.

2. Now show that the strategy σ_s maximizes the expected profit at each round for the seller s against σ_b, given beliefs and the current history of the strategies. Let p_1 be the price for the first round, p_2 be the price for the second round when $I_1 = 1$, and P_2 be the price for the second round when $I_1 = 0$. We now consider three cases.

   a) When $p_1 < 0.25$. Notice that by the update of the beliefs and the structure of σ_b, the belief of the buyer's distribution will be updated to $U[2p_1, 1]$ if $I_1 = 1$ and $U[0, 2p_1]$ if $I_1 = 0$. For the second round, we perform a similar calculation as the $n = 1$ case (Equation 2.1.2) with the updated distribution belief and can obtain that $p_2 = 0.5$ and $P_2 = p_1$. Then, the expected profit P_s can be computed as

$$\begin{aligned} P_s &= E[I_1]p_1 + P(I_1 = 1)E[I_2|I_1 = 1]p_2 + P(I_1 = 0)E[I_2|I_1 = 0]P_2 \\ &= (1 - 2p_1)p_1 + 0.5 * 0.5 + p_1 * p_1 \\ &= p_1 - p_1^2 + 0.25 \\ &\leq 0.4375. \end{aligned}$$

b) When $p_1 > 0.5$. The buyer will never buy the item in the first round. Then, the expected profit P_s will be same as the single-round auction because the update of belief does not happen as well. That is,

$$P_s = 0.25.$$

c) When $0.25 \leq p_1 \leq 0.5$. By the update of the beliefs and the structure of σ_b, the belief of the buyer's distribution will be updated to $U[2p_1, 1]$ if $I_1 = 1$ and $U[0, 2p_1]$ if $I_1 = 0$. For the second round, we perform a similar calculation as the $n = 1$ case (Equation 2.1.2) with the updated distribution belief and can obtain that $p_2 = 2p_1$ and $P_2 = p_1$. Then, the expected profit P_s can be computed as

$$\begin{aligned} P_s &= E[I_1]p_1 + P(I_1 = 1)E[I_2|I_1 = 1]p_2 + P(I_1 = 0)E[I_2|I_1 = 0]P_2 \\ &= (1 - 2p_1)p_1 + (1 - 2p_1)2p_1 + p_1 * p_1 \\ &= 3p_1 - 5p_1^2. \end{aligned}$$

We can obtain $p_1 = 0.3$ as the maximizer, which gives us $P_s = 0.45$.

Hence, we obtain $p_1 = 0.3$, $p_2 = 0.6$, and $P_2 = 0.3$. Thus, σ_s maximizes the expected utility against σ_b.

Therefore, we have shown that $\sigma = (\sigma_b, \sigma_s)$ is at PBE. □

11

2.2 Related Literature

The auction theory is a very popular field in algorithmic game theory. Researchers in many different areas including economics, operations research, and computer science have studied various topics in the auction theory such as equilibrium structures in different auction structures, the complexity of computing equilibrium, or handling privacy in the auctions.

In this section, we review such literature to give an overview of the related field to our study. In Section 2.2.1, we introduce several different topics in the auction theory and give relevant literature.

2.2.1 The Equilibrium in Auctions

The equilibrium structures in various types of auctions have been studied extensively both in economics and computer science literature (Vickrey 1961) (Karlin and Peres 2017)(Krishna 2009). We surveyed relevant literature to this paper and summarized them below.

Perfect Bayesian Equilibrium in Repeated Sales. Many studies have studied an equilibrium structure of an auction with imperfect information. In this design of a single-item auction, buyers know their own hidden values for an item and sellers try to learn them. Myerson is one of the first to study the optimal auction design for a one-shot single-item auction that achieves Bayes-Nash Equilibrium (Myerson 1981). Many studies have attempted to extend this result to repeated sales that achieve Perfect Bayesian Equilibrium. In (Hart and Tirole 1988) and (Schmidt 1993), the existence of a PBE in a single-buyer finite horizon repeated sales with discrete buyer distributions were shown. The two papers studied the effects of contracts on the PBE to study the power of commitment.

Directly building upon these papers, (Devanur et al. 2015) showed that there is no PBE in a single-buyer finite horizon repeated sales with continuous buyer distribution. It also explored the power of commitment for the seller to obtain revenue in both a finite horizon auction and an time-discounted infinite horizon auction. This study was further extended in (Immorlica et al. 2017), which considered a multi-buyer auction. By assuming the buyers' distribution are independent and identically distributed, it showed an existence of a PBE in the multi-buyer case. Our original work in this study contributes to this line of literature by investigating the frontier of the multi-buyer case more in detail by not assuming that buyers' value distributions are independent and identically distributed.

Learning in Auction and Privacy. A group of literature has emphasized on how repeated auctions entail seller learning the private information of buyers and how this data can be potentially hidden. The study in (Conitzer et al. 2012) investigated the equilibrium structure in repeated games with an option for buyers to anonymize at a cost. The study establishes that the existence of a costly option of anonymization can be beneficial for buyers. On the contrary, (Abernethy et al. 2019) designs an auction that is near-optimal and also incentivizes buyers to behave truthfully to some extent by using differential privacy. Some studies assume that the distributions of the buyers are unknown to sellers in repeated sales and the buyers draw a new value for each round. In (Amin et al. 2013), the authors show that under some assumptions, near-optimality is achieved. In (Hummel 2018), it is shown that buyers benefit from using a distribution that is lower than it actually is. Handling private information in a single-shot auction is also studied commonly such as (Ghosh and Roth 2011), which considers selling privacy in auctions.

Effectiveness of Simple Pricing Scheme. While some studies attempt to explore a more general existence and the properties of equilibrium structures, others limit their domain of strategies to a very simple pricing scheme and investigate the revenue by imposing some heavy conditions both on repeated sales and non-repeated sales. For instance, (Balcan et al. 2008) (Balcan and Constantin 2010) show that under the assumption of unlimited supply, a random simple price in a market with a sequence of random buyers achieves a strong guarantee of the revenue. Some paper impose conditions on buyers' behaviors such as (Chawla et al. 2016), which explores the revenue guarantee in repeated sales under the assumption that the buyers' value evolve every round following the Martingale condition (which we will define in Section 2.3.3). In (Cheng et al. 2018), the authors show that even when there is a budget for the buyer and the budget is private, the simple pricing scheme for selling multiple items achieves near-optimal.

13

Various Assumptions on Buyers. Another line of study is the study of equilibrium structures and the optimality with different assumptions on buyers' distributions both in repeated sales and single-shot auctions. For example, (Azar et al. 2012) designs an auction model under an assumption that buyers do not know their distributions. In this case, the seller have to rely on the buyers' knowledge of their own value. The work in (Agrawal et al. 2018) designs an repeated auction for a single buyer that achieves a constant fraction of optimal revenue for different types of buyers. That is, some buyers think about the future more than others. Therefore, this auction is more robust against different behaviors and thus reflects the situation in the real-world better. Another idea in (Azar et al. 2013) shows that there is a single round auction that is both near-optimal and efficient just by knowing the median of the value distribution not the distribution itself. The work in (Chawla et al. 2014) explores an equilibrium structure when the buyers' values are inter-dependent with each other. In (Dhangwatnotai et al. 2010), the authors show that they can design a near-optimal auction by just sampling a value of one buyer in the multi-buyer case under an assumption that all buyers' distributions are independent and identically distributed.

2.3 Detailed Literature Review

In this section, we describe three of the papers introduced in Section 2.2 in detail. Some of these papers are connected to our original work more directly and closely than other studies, and studying them can prompt a better understanding on our research described in Chapter 3.

2.3.1 Perfect Bayesian Equilibrium in Repeated Sales with a Single Buyer

We first talk about the study in (Devanur et al. 2015). This paper investigates PBE in our basic setting (Definition 2.1.7) with a single buyer. In addition to that, the study imposes a condition on commitment power by the seller (which will be defined in Definition 2.3.1 below). The authors find different PBE structures depending on the degree of the commitment. We start by defining a price commitment:

Definition 2.3.1. Suppose an auction is in the basic setting. When the seller s makes a *partial commitment*, s is not going to raise price in any consecutive rounds once a purchase is made while possessing the freedom to lower prices. When the seller s makes a *full commitment*, s is not going to change price in any consecutive rounds once a purchase is made.

In short, a commitment made by a seller is a promise regarding the future pricing of the item. In the paper, the authors first discuss the existence of PBE with no commitment. Then they move onto discuss the existence of PBE with a partial commitment. Finally, they extend the argument to discuss the existence of PBE with a partial commitment in time-discounted future. We will follow the paper's structure to discuss the results in each case.

The first main result shows that there is no PBE for $n > 2$ when there is no commitment by the seller.

Theorem 2.3.2 (Devanur et al. 2015, Theorem 1)**.** *Given an $n(> 2)$-round auction with a single buyer, for any continuous distribution F of buyer's value supported in $[l, \mathfrak{h}]$, a pure strategy threshold PBE never exists.*

Outline of the proof for Theorem 2.3.2. The authors proves Theorem 2.3.2 by proving that there is no PBE when $n = 3$ and extending the result. The paper assumes for the sake of contradiction that there is a pure threshold PBE. By using a detailed case analysis on the relationship between the buyer's value and the threshold, it concludes that there is no PBE. □

This result gives an insight to readers regarding the power of commitment in maximizing the revenue. The paper goes on to discuss PBE with a partial commitment and finds that there is a unique PBE as follows:

Theorem 2.3.3 (Devanur et al. 2015, Theorem 2)**.** *Given an auction with a single buyer for n rounds with partial commitment, for any continuous distribution F of buyer's value supported in $[l, \mathfrak{h}]$, a pure strategy threshold PBE exists.*

Outline of the proof for Theorem 2.3.3. The paper proves Theorem 2.3.2 by using induction. The argument uses the fact that the seller s will keep the price p once the buyer buys the product because of the partial commitment. The buyer simply needs to find a price where the utility of accepting and rejecting equals each other. □

The paper works out an example to demonstrate Theorem 2.3.3. The following remark summarizes the result:

Remark 2.3.4 (Devanur et al. 2015, Theorem 3). Given an auction with a single buyer with the value distribution $F \sim U[0,1]$ and partial commitment, there exists a unique pure strategy threshold PBE that obtains a revenue of $\sqrt{\frac{n}{2} + \frac{\log n}{8} + O(1)}$.

Similarly to the proof for Theorem 2.3.3, this remark is shown by an induction on n, the number of rounds.

The authors further discuss the existence of PBE in time-discounted infinity horizon. To discuss it, we must first define time-discounted infinity horizon.

Definition 2.3.5. Suppose we are given an auction is in the basic setting with a single buyer. In the *time-discounted infinity horizon* setting, the following properties hold:

1. the game is played infinitely i.e. $n \to \infty$

2. the utility of both the buyer b and the seller s decrease by a factor of $(1-\delta)$ every round. That is, the utility after ith round is multiplied by δ^{i-1}.

Before moving onto the last part, the paper discusses the connection between a PBE in a time-discounted infinite horizon repeated sales and a bargaining game. We will not discuss this result in detail in this literature review because the result is not closely related to our proposed study. In the last part, the authors discuss the PBE in time-discounted infinite horizon repeated sales auction. The result shows the following bound for the revenue in the case of $U[0,1]$.

Theorem 2.3.6 (Devanur et al. 2015, Theorem 7)**.** *Given an auction with time-discounted infinite horizon partial commitment, let the buyer's distribution be* $F \sim U[0,1]$*. If the seller* s*'s strategy is restricted to scale-invariant strategies, then there exists a unique PBE such that as* $\delta \to 0$*, the revenue approaches* $\frac{4}{3+2\sqrt{2}}$ *fraction of* $\frac{1}{4\delta}$*.*

Outline of the proof for Theorem 2.3.6. The paper proves Theorem 2.3.6 in almost an identical way as in the proof of Theorem 2.3.3. The proof uses an induction on n and find the price where the utility of accepting and rejecting equals each other. Note that the utility in this case will be time-discounted. □

The paper establishes the role of price commitment in PBE structures in a single-buyer repeated sales auction. The freedom of commitment allows the seller to obtain more revenue using a simple pricing strategy. In conclusion, the authors indicate that extending this result to multiple buyers will be a good future direction of research.

2.3.2 Perfect Bayesian Equilibrium with Multiple Buyers with i.i.d Value Function

This section describes the study in (Immorlica et al. 2017). This paper extends the work done in (Devanur et al. 2015) to multi-buyer auctions in the basic setting where the buyers' value are drawn from independent and identically distributed (i.i.d) distribution.

The first half of the paper (Immorlica et al. 2017, Section 3, 4) motivates the second half of the study (Immorlica et al. 2017, Section 5, 6). The discussion first introduces the work in (Devanur et al. 2015) and reinforces the results in Theorem 2.3.2 that there is no non-trivial simple threshold PBE (Immorlica et al. 2017, Theorem 4.2). In this review, we focus on the second half of the paper.

Instead of imposing conditions on the commitment power, this paper attempts to obtain a PBE by increasing the number of buyers. The authors claim that by increasing the number of buyers, each buyer is less important to the seller, and therefore, buyers should be less afraid of the seller exploiting their values once it's known. In order to discuss the equilibrium structure for the multi-buyer case, the authors first find the equilibrium structure for the two-buyer case (Section 5) and then extend the argument to the general k-buyer case (Section 6).

The main result shows that there is an equilibrium in the multi-buyer case:

Theorem 2.3.7 (Immorlica et al. 2017, Section 5.1, Appendix D & E)**.** *Suppose an auction has $k > 1$ buyers. Then, there is a Perfect Bayesian Equilibrium that consists of two phases:*

1. Exploration phase*: the seller offers pricing that will be rejected with positive probability until at least $k-1$ buyers reject.*

2. Exploitation phase*: the seller offers pricing as follows:*

a) If $k-1$ buyers reject in the exploration phase, the seller offers the item at a price that equals to the bottom of the support of the beliefs for the remaining buyer.

b) If k buyers reject in the exploration phase, the seller offers the item at a price that equals to the bottom of the union of the support of the beliefs for the set of buyers who rejected most recently.

Algorithm 5-10 in Appendix D & E give a comprehensive description of the buyers' and the seller's strategies. However, in this literature review, we only give the intuition of the algorithm and how to obtain it. We extend this work in Chapter 3. The algorithms described in Section 3.3 is very similar to the algorithm in Appendix D.

This auction will be conducted in the basic setting. That is, even if the item is accepted by multiple buyers, only one buyer gets the item. Therefore, in the exploitation phase, the seller ignores the buyers who reject the offer and tries to find out the buyer with the highest value distribution. Then, once the seller knows the buyer(s) with the highest value, in the exploitation phase, the seller targets them using the bottom of their support of the beliefs. On the contrary, the buyer simply follows their threshold strategy. This threshold for buyers is non-trivial because of competing buyers (otherwise, they will not be able to buy the item).

The threshold and the price will be decided by finding a threshold where the utility of accepting and rejecting equals each other similarly to the proof for Theorem 2.3.3 For the two-buyer cases, these threshold equations can be found in (Immorlica et al. 2017, Equation 2, 3). For the $k(>2)$-buyer case, the equation for the buyers can be found in (Immorlica et al. 2017, Equation 5). Note the fact that the buyers' distributions are i.i.d is used here to prove the optimality of the algorithm.

Finally, the paper shows how optimal the PBE is in the following theorems for the two-buyer case and the $k(>2)$-buyer case respectively:

Theorem 2.3.8 (Immorlica et al. 2017, Theorem 5.1)**.** *Given the two-buyer version of the infinite-horizon auction, let F be the value distribution for two buyers. Suppose F maintains monotone hazarad rate i.e. $\frac{f(v)}{1-F(v)}$ and $\delta \geq \frac{2}{3}$. In the equilibrium, the lower bound of the seller revenue is $\frac{1}{3e^2}$ of the optimal revenue.*

Theorem 2.3.9 (Immorlica et al. 2017, Theorem 6.2, Modified)**.** *Given the $k(>2)$-buyer version of the infinite-horizon auction, let F be the value distribution for two* *buyers. Suppose F maintains monotone hazard rate i.e. $\frac{f(v)}{1-F(v)}$ and $\delta \geq \frac{n}{n+1}$. In the equilibrium, the lower bound of the seller revenue is approximately* 0.1202 *of the revenue of the revenue-optimal auction for F.*

Note that Theorem 2.3.9 is only the weaker version of (Immorlica et al. 2017, Theorem 6.2). To avoid introducing unnecessarily many definitions and to make clear contrast between Theorem 2.3.8 and Theorem 2.3.9, we gave a stricter condition on F.

The outline of the proofs for Theorem 2.3.8, Theorem 2.3.9. To prove Theorem 2.3.8, the authors show that there always exists a sequence of prices that induces a certain threshold. By using a sequence that induces the optimal price for single-buyer monopoly market, the seller can achieve the revenue stated above. The proof for Theorem 2.3.9 follows the same technique as the two-buyer case (Appendix F). □

This shows that by imposing some conditions on buyers' values, we can achieve a near-optimal auction in the multi-buyer case without relying on the power of commitment. In applications in the real-world where there are multiple buyers that seem to value the good in a similar way, sellers can use this result as a general direction for the pricing.

2.3.3 Repeated Sales with Evolving Values

We finally introduce the work in (Chawla et al. 2016). The study designs a different setting compared to previous studies to capture different aspects of repeated sales. While the paper investigates the revenue-maximizing pricing mechanism, the striking difference between the previous papers in Subsection 2.3.1 and Subsection 2.3.2 and this paper is that the buyers' values evolve with time and usage.

We introduce the paper because this setting is more reflective of some real-world situation; for example, to evaluate the value of Netflix subscription service, people review the value of their usage every month depending on how much they enjoyed it in the previous month. One important assumption the paper uses to realizes this is to assume that the buyers' values evolve as Martingale i.e., the expected value of the buyers' value in the new time period given all the history of their values equals to the value of the previous time period. We will now formally define the assumption of buyers' value evolution:

Definition 2.3.10 (Chawla et al. 2016, Section 1.1)**.** Let V_t be the value of a buyer b at time period t. Suppose the initial value V_0 is positive with $V_0 \in (0, 1)$ and $V_t \in [0, 1]$. Then, the study assumes the following:

1. *Absorbing Condition.* If $V_t = 0$, then $V_{t+1} = 0$.

2. *Martingale Condition.* If $V_t \in (0, 1)$, then $E[V_{t+1}|V_0, \ldots, V_t] = V_t$.

3. *Bounded Step Size Condition.* For every t, there exists a small constant $\epsilon > 0$ such that $\epsilon \geq |\Delta_t|$.

4. *Minimum Variance Condition.* There exists $\delta \in (0, \epsilon)$ such that $V_t \in (0, 1]$ implies $E[\Delta_t^2|V_0, \Delta_0, \ldots, \Delta_{t-1}] \geq \delta^2$.

These conditions are tied to the real-world phenomena. The first condition describes that once a buyer loses their interest, they will not regain it. The second condition describes that their new value is only based on the most recent evaluation from the previous time period. The third and fourth condition limit the magnitude of the change.

In this paper, the authors try to find a approximation of the revenue for sellers using a simple pricing scheme instead of finding the optimal strategy. In this auction, the seller employ a simple pricing scheme as follows:

1. Offer a free trial for the first T time period;

2. Thereafter, set a constant price c per time period.

This should be fairly familiar to readers in the form of common subscription services such as Netflix and Spotify. However, it should be intuitive that a more complicated dynamic pricing scheme dependent on buyers' value would yield more revenue.

Now, we assume that an extremely risk-averse buyer b that buys the item at time period t only if V_t exceeds c. Then, we can obtain the following approximation for the revenue $\mathcal{R}_{T,c}$ of the seller.

Theorem 2.3.11 (Chawla et al. 2016, Theorem 1.1). *Suppose the buyer b's value evolution satisfies all conditions described in Definition 2.3.10 and $\frac{\epsilon^2}{\delta}$ is sufficiently small. Let $\mathcal{C}(v)$ be the cumulative value of the buyers from the all buyers' period. Then, there exist constants T and c as a function of ϵ and δ such that given $V_0 = v$,*

$$\mathcal{R}_{T,c|V_0=v} = \Omega\left(\frac{\delta^3}{\epsilon^3}\mathcal{C}(v)\right)$$

where $\mathcal{R}_{T,c}$ is the expected revenue for the risk-averse buyer.

This result implies that simple pricing mechanism obtains constant fraction of the cumulative value of the buyer. While we will not give details in this literature review, the paper shows each condition in Definition 2.3.10 is necessary to obtain the bound in Theorem 2.3.11 (Chawla et al. 2016, Section 3).

Outline of the proof for Theorem 2.3.11. The paper proves Theorem 2.3.11 using a series of Lemmas quantifying the relationship between the expected revenue and three defined time periods:

1. τ_0: the first time that $V_t \geq \frac{\delta}{\epsilon}$ or $V_t = 0$;

2. τ_1: the first time after τ_0 that $V_t = 1$ or $V_t = c$, the item's price (conditional on $V_{\tau_0} \geq \frac{\delta}{\epsilon}$ since the value stays at zero otherwise);

3. τ_2: the first time after τ_1 that $V_t = c$ (conditional on $V_{\tau_0} \geq \frac{\delta}{\epsilon}$ and $V_{\tau_1} = 1$).

These time periods enable us to show that there is large enough probability of the buyer's value reaching 1 eventually and if that happens, there is a long enough time period that the buyer's value stays above the item's price. Specifically, first, Lemma 4.1, Lemma 4.2, and Lemma 4.3 obtain the probability of the buyer's value reaching $\frac{\delta}{\epsilon}$ before reaching 0, and show that, conditional on reaching $\frac{\delta}{\epsilon}$ before 0, τ_0 is relatively small. Then, the authors prove Lemma 4.4 and Lemma 4.5, conditional probability of the buyer's value reaching above $\frac{\delta}{\epsilon}$, to show that the probability of $V_{\tau_1} = 1$ and the expected value of $\tau_2 - \tau_1$ are large enough.

Furthermore, Lemma 4.6 shows the upper bound for $\mathcal{C}(v)$. Then, by choosing T carefully $\left(T = \frac{6}{5} = \mathbf{E}\left[\tau_0 | V_{\tau_0} \geq \frac{\delta}{\epsilon}\right]\right)$, the authors obtain the lower bound of the expected revenue using the lemmas and show that it is large enough compared to $\mathcal{C}(v)$ (i.e. the constant fraction of $\mathcal{C}(v)$). □

This main result reinforces the effectiveness of the simple pricing scheme proposed above that is widely used in repeated sales with an evolving value. However, the authors also acknowledges that this result was obtained due to a heavy restriction imposed on the buyer's condition (Definition 2.3.10).

3

Perfect Bayesian Equilibria in a Repeated Auction with Multiple Buyers

In this chapter, we extend the scope of current study by considering an equilibrium structure in repeated sales with multiple buyers where their value distributions are not necessarily independent and identically distributed. In Section 3.1, we propose the method of our study in detail and explain the motivation behind it. Then, in Section 3.2, we formally define our extended problem and state our main result. We go on to describe the full algorithm in Section 3.3, followed by its proof of optimality in Section 3.4. Finally, in Section 3.5, we analyze the feasibility of the assumptions imposed on our algorithm.

3.1 The Motivation and an Overview of the Method

In this section, we explain the motivation of our study and an overview of the method of our study. Our original work directly extends the work done in (Devanur et al. 2015) and (Immorlica et al. 2017). Recall from Theorem 2.3.7 that there exists a Perfect Bayesian Equilibrium for repeated sales under the assumption that the buyers' value distributions are independent and identically distributed.

The ultimate goal of this line of study is to investigate the equilibrium structure in the n-buyer case without imposing any assumptions on buyers' distributions. Here, we only focus on the two-buyer case because

1. the two-buyer case is much easier to work on compared to the multi-buyer case; and

2. extending the two-buyer case to the multi-buyer case can be generally done by induction on the number of buyers.

By discussing an equilibrium structure in this scenario, we hope to give an insight to hypothesize a more general equilibrium structure for the multi-buyer case, which better reflects the diversity of buyers in the real-world. Furthermore, obtaining a general equilibrium may allow us to easily impose different conditions to restrict buyers' behaviors as well.

3.2 A Repeated Sales with Multiple Buyers

In this section, we formally state the multi-buyer extension of the problem stated in (Immorlica et al. 2017).

3.2.1 The Extended Problem and the Main Result

The auction we are exploring is almost identical to the auction described in (Immorlica et al. 2017). The only main difference is that the value distributions for two buyers are not necessarily i.i.d. but are arbitrary. We let F_1 and F_2 denote the value distributions for the buyer 1 and the buyer 2 respectively with the same initial support interval $[a, b]$. In the following sections, we explore if there is a PBE in this auction setting.

Our natural first instinct is to explore if there is an equilibrium when the higher threshold is used after each buyer separately calculates their respective threshold. Although we are not able to prove a general existence of PBE, we show that there is a PBE under certain set of assumptions and claim it as follows:

Claim 3.2.1. At a given round, let i be the buyer with the higher threshold and denote their threshold by t_i. Suppose the following equation holds:

$$-2\epsilon \leq (F_i(t_i) + 1)(t_i - p_k) - F_i(t_i)(t_i - a^*)\frac{\delta}{1-\delta} \leq 0 \tag{3.2.1}$$

where ϵ is a small enough positive constant. Then, there is a Perfect Baysian Equilibrium.

Our goal in the subsequent sections (Section 3.3 and 3.4) is to prove Claim 3.2.1 by illustrating an algorithm to obtain PBE and proving that it works. Assumption 3.2.1 is a sufficient condition to our proof in Section 3.4 for the existence of a PBE.

3.3 Full Algorithm

In this section, we give a full description of the algorithm used by each buyer and the seller. We keep the same structure of the "exploration" phase and the "exploitation" phase, which was used in (Immorlica et al. 2017). We first obtain the threshold for each buyer in Section 3.3.1 since our goal is to obtain a threshold equilibrium. Using the obtained thresholds, we explain our full algorithm in Section 3.3.2. The work in this section is adapted from the work in (Immorlica et al. 2017, Appendix D).

3.3.1 The Threshold Equations and the Revenue Equation

The threshold for this version of the problem is calculated in a very similar way as the original algorithm in (Immorlica et al. 2017, Appendix D). However, since the buyers (could) have different distributions, we need to calculate the threshold separately. Let $[a_1, b_1]$ and $[a_2, b_2]$ be the support interval for each buyer. The threshold equation for buyer 1 to calculate their threshold t_1 is:

$$(t_1 - p)\left(F_2(t_1) + \frac{1 - F_2(t_1)}{2}\right) = \frac{t_1 - a^*}{2} F_2(t_1) \frac{\delta}{1 - \delta} \tag{3.3.1}$$

where $a^* = \max(a_1, a_2)$ and p is the price. Similarly, we can obtain the threshold for buyer 2, t_2, as follows:

$$(t_2 - p)\left(F_1(t_2) + \frac{1 - F_1(t_2)}{2}\right) = \frac{t_2 - a^*}{2} F_1(t_2) \frac{\delta}{1 - \delta}. \tag{3.3.2}$$

We now explain the meaning of the threshold equations. It finds the threshold where for a buyer with value t_i, accepting the price and rejecting the price are equal to each other. The lefthand side indicates accepting the price. The utility gained when the buyer i is able to buy the item is $(t_i - p)$. However, it only happens either

1. when the other buyer j rejects the item; or

25

2. when the other buyer j also accepts the item but buyer i obtains the item by chance.

Case 1 happens with the probability $F_j(t_i)$ and Case 2 happens with the probability $\frac{1-F_j(t_i)}{2}$.

The righthand side indicates rejecting the price. The profit gained from that round is 0, but the probability that other buyer j rejects the price is $F_j(t_i)$, and we can expect the utility of $\frac{1}{2}(t_2 - a^*)\frac{\delta}{1-\delta}$ where the discount rate is δ since the buyer gains the item for half the time with the profit of $t_2 - a^*$.

Notice that the equation to obtain the threshold for buyer 1 uses the distribution for buyer 2 and vice versa. This is because buyer 1's behavior is restricted by the probabilistic behavior of buyer 2. This is one of the critical differences between this version of the problem and the original version in (Immorlica et al. 2017, Appendix D).

To obtain the price, we use the threshold obtained from Equation 3.3.1 and Equation 3.3.2 and solve the recurrence relationship. Let $T_i(a_i, b_i, p)$ be the classes of solutions for the threshold equation for buyer i. Now, suppose $\max(T_1(a_1, b_1, p)) \leq \max(T_2(a_2, b_2, p))$ without loss of generality. Similarly to (Immorlica et al. 2017, Appendix D), we use the following threshold for the calculation:

$$t(a^*, b^*, p) = \begin{cases} \infty & \text{if } T_1(a_1, b_1, p), T_2(a_2, b_2, p) = \emptyset \\ p^* & \text{if } p^* \in T_1(a_1, b_1, p) \text{ or } p^* \in T_2(a_2, b_2, p) \\ \max(T_2(a_2, b_2, p)) & \text{otherwise} \end{cases}$$

where p^* is the monopoly price (the profit maximizer) for the initial value distribution for each buyer, $a^* = \max(a_1, a_2)$, and $b^* = \min(b_1, b_2)$. Let $R(a^*, b^*, p)$ be the revenue of the seller. Then, by using this threshold, we obtain the following revenue equation where $t = t(a^*, b^*, p)$:

26

$$
\begin{aligned}
R(a^*, b^*, p) =&(1 - F_1(t))(1 - F_2(t))(p + \delta R(t, b^*, p)) \\
&+ F_1(t)(1 - F_2(t)) \left(p + \frac{\delta}{1-\delta} t\right) \\
&+ F_2(t)(1 - F_1(t)) \left(p + \frac{\delta}{1-\delta} t\right) \\
&+ F_1(t) F_2(t) \frac{\delta}{1-\delta} a^*.
\end{aligned}
$$

The first term is when both buyers accept. From the current round, the seller's profit is p and from the next round, the seller obtains $R(t, b^*, p)$ while discounted at the rate of δ. The second and third term are when only buyer 2 accepts and when only buyer 1 accepts, respectively. From the current round, the seller's profit is p and from the next round, the seller keeps posting the price of t for the rest of the game while discounted at the rate of δ. The fourth terms refers to the case where both buyers reject. The seller will post $a^* = \max(a_1, a_2)$ for the rest of the game. Hence, we can simply set p to be the maximizer of $R(a^*, b^*, p)$. By solving the threshold equations and the revenue equation recursively, we can obtain the threshold and the price.

3.3.2 A Description of the Full Algorithm

In this section, we specify the algorithm for each agent. That is, the pricing algorithm and the belief update algorithm for the seller and the decision making algorithm for each buyer. In the algorithms, we use $h^k = (h_1, \cdots, h_k)$ to refer to the history of behaviors taken by each buyer up to the k-th round where $h_i = \left(h_{(1,i)}, h_{(2,i)}\right)$ with $h_{(b,i)} \in \{Accept, Reject\}$ specifying the behavior taken by the buyer b on round i.

Similarly to the algorithm in (Immorlica et al. 2017, Appendix D), there are two phases in the algorithm:

1. *Exploration Phase:* In this phase, the seller keeps raising the price by adjusting the belief for each buyer every round until at least one of the buyer rejects the price. To decide the price, the seller maximizes the revenue equation $R(a^*, b^*, p)$ by solving the recurrence relationship. In a special case where $p^* \leq a^*$, we simply use a^* for the two buyers with distributions F_1 and F_2. A

buyer consults the threshold $t = \max(t_1, t_2)$, and decides if they should buy or not.

2. *Exploitation Phase:* After at least one buyer rejects, we enter this phase. The seller uses $a = \max(a_1, a_2)$ for the rest of the game after the belief update of the last round in the exploration phase. The buyer accepts the price if their value is higher than the price and rejects otherwise.

According to the description above, we first present the pricing algorithm for the seller.

Algorithm 1 The Seller's Pricing Algorithm for Round $(k+1)$ (Immorlica et al. 2017, Algorithm 5)

Input: The history of behaviors h^k and support intervals $[a_1^k, b_1^k]$ and $[a_2^k, b_2^k]$.
Output: Price p_{k+1} for the $(k+1)$-th round.

1: **if** $h^k == (Accept, Accept)^k$ **then**
2: $\quad a^* = \max(a_1^k, a_2^k)$; $b^* = \min(b_1^k, b_2^k)$
3: $\quad$ **if** $a^* \geq p^*$ **then** ▷ Already higher than the optimal revenue price.
4: $\quad\quad p_{k+1} = a^*$
5: $\quad$ **else**
6: $\quad\quad p_{k+1} = \arg\max_p R(a^*, b^*, p)$ such that $T_1(a_1^k, b_1^k, p) \neq \emptyset$ or $T_2(a_2^k, b_2^k, p) \neq \emptyset$
7: ▷ Solve the recurrence relation.
8: $\quad$ **end if**
9: **else** ▷ Already rejected by at least one buyer.
10: $\quad p_{k+1} = a^*$ ▷ Can't raise the price anymore.
11: **end if**

From Lines 1 to 7, we describe the pricing algorithm for the exploration phase where the seller keeps raising the price unless the support interval is already above the revenue optimal price. Below Line 8, we describe the exploitation phase where the seller keeps posting the same price.

We now illustrate the decision making algorithm for buyer i. Call the other buyer buyer j. The algorithm is very similar to (Immorlica et al. 2017, Algorithm 5) with a slight modification.

28

Algorithm 2 The Algorithm for Buyer i at Round k (Immorlica et al. 2017, Algorithm 6)

Input: The history of behaviors h^k, support intervals $[a_1^k, b_1^k]$ and $[a_2^k, b_2^k]$, value v_i, price p_k, and value distributions F_1 and F_2.
Output: Accept or Reject decision for buyer i in Round k.

1: $a^* = \max(a_1^k, a_2^k)$
2: $t^* = \max(t_1(a_1^k, b_1^k, p^*), t_2(a_2^k, b_2^k, p^*))$ ▷ Use the higher threshold.
3: **if** $h^k == (Accept, Accept)^k$ **then** ▷ In the exploration phase.
4:     **if** $a^* \geq p^*$ **then** ▷ Already higher than the optimal revenue price.
5:         Accept if and only if $p_k \leq a^*$ ▷ Accept if the price is lower than a^*.
6:     **else**
7:         Accept if and only if $v_i \geq t^*$ ▷ Use the threshold.
8:     **end if**
9: **else** ▷ In the exploitation phase.
10:     **if** $p_k \leq a^*$ **then**
11:         Accept if and only if $v_i \geq p_k$ ▷ Accept when the value is higher than p_k.
12:     **else**
13:         Reject ▷ Punish the off-path behavior by the seller.
14:     **end if**
15: **end if**

Notice that Algorithm 2 marks a difference between the analysis of the original algorithm and this version of the algorithm. Since buyer 1 and buyer 2 obtain different thresholds, t_1 and t_2 respectively, we simply use the bigger threshold. From Lines 3 to 8, we describe the buyer's behavior during the exploration phase. Similarly to the seller, the buyer changes their behavior based on the relationship between the revenue optimal price p^* and the support interval. Below Line 9, we describe the exploitation phase. Notice that the buyer punishes the off-path behavior by the seller by rejecting every price such that $p_k > a^*$.

We finally present the belief update algorithm for buyer i. Let the other buyer be buyer j. Note that this algorithm guarantees the satisfaction of Condition 1 in Definition 2.1.8 by using the Bayes' rule.

29

Algorithm 3 The belief update algorithm for buyer i after Round k (Immorlica et al. 2017, Algorithm 7)

Input: The history of behaviors h^k up to the k-th round, support intervals $[a_i^k, b_i^k]$, the initial support interval $[a, b]$, and threshold used in Algorithm 2, t^*.

Output: Updated support intervals $[a_i^{k+1}, b_i^{k+1}]$ for the $(k+1)$-th round.

1: **if** Buyer i was expected to reject p_k **then** ▷ Happens in the exploitation phase.
2:     **if** The buyer i rejected p_k **then**
3:         $[a_i^{k+1}, b_i^{k+1}] = [a_i^k, b_i^k]$
4:     **else**
5:         $[a_i^{k+1}, b_i^{k+1}] = [b, b]$ ▷ Punish the off-path behavior by buyer i.
6:     **end if**
7: **else if** Buyer i was expected to accept p_k **then**
8: ▷ Happens in the exploitation phase.
9:     **if** The buyer i accepted p_k **then**
10:         $[a_i^{k+1}, b_i^{k+1}] = [a_i^k, b_i^k]$
11:     **else**
12:         $[a_i^{k+1}, b_i^{k+1}] = [b, b]$ ▷ Punish the off-path behavior by buyer i.
13:     **end if**
14: **else** ▷ The exploration phase.
15:     **if** The buyer i accepted p_k **then**
16:         $[a_i^{k+1}, b_i^{k+1}] = [t^*, b_i^k]$
17:     **else**
18:         $[a_i^{k+1}, b_i^{k+1}] = [a_i^k, t^*]$
19:     **end if**
20: **end if**

Since both buyers use t^*, the higher threshold of t_1 and t_2, the belief update is performed using t^* during the exploration phase. Below Line 13, we update the belief depending on the buyer's behavior. Lines 1 to 12 describe the update during the exploitation phase where Line 1-6 correspond to Line 13 in Algorithm 2 and Lines 7 to 12 correspond to Line 11 in Algorithm 2. Note that during the exploitation phase, the buyer is expected to take the same action each round. If the buyer exhibits the same behavior as the expected one, we do not change the belief (Lines 2 to 3 and Lines 8 to 10). However, if buyer i takes an unexpected action, we punish them by simply changing the belief to a point mass at the top of the initial support interval for the rest of the game. By doing so, buyer i will not receive any profit for the rest of the game. We describe it in detail in the next section.

30

3.4 Equilibrium Proof

In this section, we attempt to prove that the set of algorithms presented above (Algorithm 1, 2, and 3) defines an equilibrium in this auction. The proof is given in Section 3.4.1. In this proof, Assumption 3.2.1 is crucial. We give a detailed analysis of the validity of the assumption in Section 3.5. The proof uses the main idea from (Immorlica et al. 2017, Appendix D) while requiring our own ideas to adapt to the new version of the auction.

3.4.1 Proof

Recall from Definition 2.1.8 that we need to prove that the algorithms used by each agent is optimal with respect to other agents' algorithm. The proof follows a similar structure to the proof in (Immorlica et al. 2017, Appendix D) by first proving that the off-path behavior is not optimal and then proving the optimality of the algorithm.

We begin by proving the ineffectiveness of an off-path action.

Lemma 3.4.1 (Immorlica et al. 2017, Lemma D.2.). *Suppose buyer i takes an off-path action that is expected to have zero probability in Round k. Then, there is an action for buyer i that gives more utility. That is, it is not the best strategy with respect to the other agents' strategies.*

Although (Immorlica et al. 2017, Proof of Lemma D.2.) gives a very short proof, in this paper, we give a comprehensive proof, which is our original work.

Proof of Lemma 3.4.1. There are two occasions that an off-path action can happen as follows:

1. In the exploitation phase, buyer i accepts the price although buyer i is expected to reject the price (Line 13, Algorithm 2).

2. In the exploitation phase, buyer i rejects the price although buyer i is expected to accept the price (Line 11, Algorithm 2).

31

We give a detailed analysis of each case. Denote the utility obtained by taking the on-path action by U_{on} and the utility obtained by taking the off-path action by U_{off}. It is clear that proving $U_{\text{on}} \geq U_{\text{off}}$ suffices to prove the lemma.

We begin with Case 1. In the exploitation phase, the seller keeps posting the price of $a^* = \max(a_1, a_2)$. If buyer i is expected to reject the price at Round k, we know that $v_i < p_k$. Since buyer i should keep rejecting the price (Line 11, Algorithm 2),

$$U_{\text{on}} = 0.$$

However, if buyer i accepts the price at Round k, from Line 5 in Algorithm 3, the belief is updated to $[b, b]$. Then, the seller will keep posting the price of b after this round. Notice that $v_i \leq b$ since b is the upper bound of the initial support interval. Therefore, regardless of actions taken by buyer i in rounds going forward,

$$U_{\text{off}} < (v_i - p_k) < 0.$$

Therefore, we obtain $U_{\text{on}} \geq U_{\text{off}}$.

We now move onto Case 2. This case is very similar to the analysis of Case 1. In the exploitation phase, the seller keeps posting the price of $a^* = \max(a_1, a_2)$. If buyer i is expected to accept the price at Round k, we know that $v_i \geq p_k$. Since buyer i should keep accepting the price (Line 11, Algorithm 2),

$$U_{\text{on}} = \frac{\delta}{1-\delta}(v_i - p_k) > 0$$

where δ is the discount rate. However, if buyer i rejects the price at Round k, from Line 12 in Algorithm 3, the belief is updated to $[b, b]$. Then, the seller will keep posting the price of b after this round. Notice that $v_i \leq b$ since b is the upper bound of the initial support interval. Therefore, regardless of actions taken by buyer i in rounds going forward,

$$U_{\text{off}} \leq 0.$$

Therefore, we obtain $U_{\text{on}} \geq U_{\text{off}}$.

We have proven that each off-path action results in a lower utility gained. Therefore, we have proven Lemma 3.4.1. □

We now move on to prove that the set of algorithms mentioned above is optimal. We begin with an analysis of the seller's strategy.

Lemma 3.4.2 (Immorlica et al. 2017, Modified Lemma D.4.). *Suppose the seller and the buyers follow the algorithms specified in Algorithm 1, 2, and 3. Then, the seller is best-responding to both buyers' strategies.*

Before proving Lemma 3.4.2, we state the following lemmas.

Lemma 3.4.3. *Suppose the seller updates their belief using the algorithm specified in Algorithm 3. Then, the update satisfies Condition 1 of Definition 2.1.8.*

Proof of Lemma 3.4.3. Recall that in Algorithm 2, buyers use the same threshold t_2 (when $t_2 \geq t_1$) for determining their behaviors during the exploration phase. We can easily see that Algorithm 3 is using the Bayes' rule of conditional probability to update the belief based on the buyers' behaviors and the threshold. □

Lemma 3.4.4. *Suppose the seller updates their belief using the algorithm specified in Algorithm 3. Then, using the updated belief is at least as good as using the old belief before the update.*

Proof of Lemma 3.4.4. Consider three different cases:

1. When both buyers reject the current price. Since the seller keeps posting the price of a^*, the update does not matter.

2. When one buyer rejects and one buyer accepts the current price. The seller will post the price of the current a^* without an update and t with an update. Since $t \geq a^*$, we know that updating the belief is at least as good as not updating.

3. When both buyers accept the current price. The seller will keep posting the price of the current a^* without an update and will be gaining at least t in all future rounds with an update. Since $t \geq a^*$, we know that updating the belief is at least as good as not updating.

In each case, updating the belief is at least as good as not updating. Hence, we have shown the statement. □

Although the proof of (Immorlica et al. 2017, Lemma D.4.) is given in the paper and we follow the same structure, we extend the proof to include all details. Hence, this is our original work.

Proof of Lemma 3.4.2. Similarly to (Immorlica et al. 2017, Proof of Lemma D.4.), we consider the action taken by the seller in the exploration phase and the exploitation phase.

1. We first consider the action taken by the seller in the exploration phase. Recall that during the exploration phase, the seller will post the price that yields the maximum revenue by construction (which maximizes $R(a^*, b^*, p)$), given that the threshold exists and $p^* > a^*$. We need to verify that the strategy is still optimal when $p^* \leq a^*$ and when the threshold does not exist.

   First, consider the case where $p^* \leq a^*$. We need to show that $p_k = a^*$ is the optimal pricing (Line 4, Algorithm 1). Recall that both buyers accept if and only if $p_k \leq a^*$ (Line 5, Algorithm 2). Therefore, when $p_k = a^*$, the utility $U_{=a^*}$ is

   $$U_{=a^*} = a^* + \delta U_{=a^*}$$

   since the support interval after a belief update is unchanged. Suppose $p_k > a^*$. Then, the utility $U_{>a^*}$ is

   $$U_{>a^*} = 0 + \delta U_{=a^*},$$

   so we can conclude that $U_{=a^*} > U_{>a^*}$. Similarly, suppose that $p_k < a^*$. Then, the utility $U_{<a^*}$ is

   $$U_{<a^*} = p_k + \delta U_{=a^*}.$$

   However, since $p_k < a^*$, we can conclude that $U_{=a^*} > U_{<a^*}$. Hence, if $p^* \leq a^*$, then the seller is achieving the optimal pricing by posting the price of $p_k = a^*$.

   We now consider the case where the seller posts the price where the threshold does not exist. Note that there is dependency because the price chosen by the seller affects the threshold equation. Suppose there is a threshold. Then, the utility gained U_{t} is

   $$U_{\text{t}} = U_{\text{cur}} + \delta U_{\text{new}} \geq \delta U_{\text{new}}$$

   where U_{cur} is the utility gained from the current round and U_{new} is the utility gained from the next round onward with the updated belief. On the contrary, suppose there is no threshold. Then, the utility gained $U_{\text{no-t}}$ is

   $$U_{\text{no-t}} = 0 + \delta U_{\text{t}}$$

   since both buyers always reject for the current round because the threshold is ∞, and there will be no update. By Lemma 3.4.4, we know that $U_{\text{new}} \geq U_{\text{t}}$, and hence $U_{\text{t}} \geq U_{\text{no-t}}$. Therefore, the seller should choose the price p_k with a non-empty set of thresholds. Note that this part of the algorithm is actually encoded in choosing a maximizer of the revenue equation, $R(a^*, b^*, p)$. If $t = \infty$, we can easily obtain that $R(a^*, b^*, p) = 0$, so the seller will not use such p.

2. We now consider the action taken by the seller in the exploitation phase. In the exploitation phase, the seller is expected to post $p_k = a^* = \max(a_1, a_2)$ (Line 10, Algorithm 1). Recall that there is no update of belief unless there is an off-path behavior taken by a buyer. Since we are considering the case where the buyer does not take an off-path behavior, we can assume that an update of the belief does not occur after Round k. In this case, at least one buyer is expected to accept the price (at least one buyer has a value more than a^*). Therefore, the utility $U_{=a^*}$ gained by offering p_k such that $p_k = a^*$ is

$$U_{=a^*} = a^* + \delta U_{k+1}$$

where U_{k+1} indicates the maximum utility obtained from $(k+1)$-th round onward.

Now suppose in Round k, the seller posts the price of $p_k > a^*$ trying to exploit the buyers more. Then, from Line 13 of Algorithm 2, both buyers reject the price. Therefore, the utility $U_{>a^*}$ gained by offering p_k such that $p_k > a^*$ is

$$U_{>a^*} = 0 + \delta U_{k+1}.$$

Therefore, we obtain $U_{=a^*} > U_{>a^*}$.

Similarly suppose in Round k, the seller posts the price of $p_k < a^*$. Then, from Line 11 of Algorithm 2, at least one buyer is expected to accept the price. Therefore, the utility $U_{<a^*}$ gained by offering p_k such that $p_k < a^*$ is

$$U_{<a^*} = p_k + \delta U_{k+1}.$$

However, notice that since $p_k < a^*$, we can easily obtain $U_{=a^*} > U_{<a^*}$. Therefore, the posting $p_k = a^*$ is optimal.

We have shown that the seller's behavior achieves an optimality given other agents' behaviors both in the exploration and the exploitation phase. Hence, we have proven the statement. □

35

We now prove that each buyer's behavior is an equilibrium-achieving behavior.

Lemma 3.4.5 (Immorlica et al. 2017, Modified Lemma D.3.)**.** *Suppose the seller and the buyers follow the algorithms specified in Algorithm 1, 2, and 3. Then, at a given round k, buyer i is best-responding to other agents' strategies for $\delta \geq \frac{2}{3}$ under Assumption 3.2.1 (δ is the time-discount factor).*

The proof of Lemma 3.4.5 has a very similar structure to (Immorlica et al. 2017, Proof of Lemma D.3.). However, our proof extends the argument for the multi-buyer scenario we are concerned with.

Proof of Lemma 3.4.5. Similarly to the proof of Lemma 3.4.2, we analyze the strategy for the exploration phase and the exploitation phase separately. Furthermore, we need to perform an analysis for each buyer. Without loss of generality, we can say that for Round k, the threshold for buyer 1 is smaller than the threshold for buyer 2. That is,

$$t^* = \max(t_1(a_1, b_1, p), t_2(a_2, b_2, p)) = t_2.$$

We first consider the exploitation phase. Let the support interval for each buyer be $[a_1^k, b_1^k]$ and $[a_2^k, b_2^k]$. Suppose that the price for this round is p_k.

1. Consider buyer 2. First assume that $v_2 \geq t_2$. The utility for buyer 2 accepting p_k, denoted by U_A, satisfies the following equation:

$$U_A \geq \left(F_1(t_2) + \frac{1 - F_1(t_2)}{2}\right)(v_2 - p_k) + \frac{\delta}{1-\delta}(v_2 - t_2)F_1(t_2). \qquad (3.4.1)$$

   We can obtain the first term $\left(F_1(t_2) + \frac{1-F_1(t_2)}{2}\right)(v_2 - p_k)$ by considering the utility from the current round: buyer 1 rejects the price with a probability of $F_1(t_2)$ and accepts the price with a probability of $(1 - F_1(t_2))$. If buyer 1 accepts the price, buyer 2 only obtains the item for half the time. The second term $\frac{\delta}{1-\delta}(v_2 - t_2)F_1(t_2)$ is the utility from the current round and the future round obtained when buyer 1 rejects the price in the current round. Since $v_1 \leq t_2$ and $a_n^* = t_2$ for all $n > k$ after the belief update, buyer 2 obtains the item for all future rounds. To calculate the utility obtained in the future rounds, we need to obtain the sum of geometric series with a multiplier δ, a time-discount factor. Notice that we are only using an inequality since we are not accounting for the future utility when both buyers accept the price for the current round.

36

Similarly, the following holds for the utility for buyer 2 rejecting p_k, denoted by U_R:

$$U_R \leq F_1(t_2)\frac{v_2 - a^*}{2}\frac{\delta}{1-\delta} + \delta(1 - F_1(t_2))\frac{v_2 - t_2}{2}. \tag{3.4.2}$$

We can obtain the first term $F_1(t_2)\frac{v_2-a^*}{2}\frac{\delta}{1-\delta}$ by considering the utility from the current round and the future round when buyer 1 rejects the price as well with a probability of $F_1(t_2)$. This means that a^* will not be updated after this round, and there is a competition between buyer 1 and buyer 2.

The second term $\delta(1-F_1(t_2))\frac{v_2-t_2}{2}$ is the total utility obtained from the current round and all future rounds when buyer 1 accepts the current price (while buyer 2 does not). Notice that buyer 2 will not get the item for this round, but buyer 2 can still participate in future rounds after the belief update happens. The seller will update the belief for buyer 1 to $[t_2, b_2]$. Since the auction enters the exploitation phase and $v_2 > a^* = t_2$, in the next round (Round $(k+1)$), buyer 2 accepts the price t_2. However, this is an off-path behavior by buyer 2. Hence, the seller will update the belief furthermore (Algorithm 3, Line 5), and there is no utility gained by buyer 2 in any future rounds.

We now need to show that $U_A - U_R \geq 0$ (Algorithm 2, Line 7). From Equation 3.4.1 and 3.4.2, we can compute:

$$\begin{aligned} U_A - U_R = (v_2 - t_2 + t_2 - p_k)\frac{1 + F_1(t_2)}{2} - (v_2 - t_2 + t_2 - a^*)\frac{F_1(t_2)\delta}{2(1-\delta)} \\ + (v_2 - t_2)\left(\frac{\delta F_1(t_2)}{1-\delta} - \frac{\delta(1 - F_1(t_2))}{2}\right). \end{aligned}$$

Here, we substitute Equation 3.3.2:

$$(t_2 - p)\left(F_1(t_2) + \frac{1 - F_1(t_2)}{2}\right) = \frac{t_2 - a^*}{2}F_1(t_2)\frac{\delta}{1-\delta}.$$

Specifically, we apply the left-hand side of Equation 3.3.2 to the first term. Then, we can simplify the equation as follows:

$$\begin{aligned} & U_A - U_R \\ & = (v_2 - t_2)\left(\frac{1 + F_1(t_2)}{2} - \frac{\delta F_1(t_2)}{2(1-\delta)} + \frac{\delta F_1(t_2)}{1-\delta} - \frac{\delta(1 - F_1(t_2))}{2}\right) \\ & = (v_2 - t_2)\left(\frac{\delta F_1(t_2)}{2(1-\delta)} + \frac{(1-\delta) + (1+\delta)F_1(t_2)}{2}\right). \end{aligned}$$

Since $v_2 \geq t_2$ and $0 < \delta < 1$, both terms are positive. Hence,

$$U_A - U_R \geq 0$$

Therefore, we have shown that accepting the price gives a higher utility than rejecting the price for buyer 2 when $v_2 \geq t_2$.

Now consider the case when $v_2 < t_2$. Then, we need to show that buyer 2 should reject the current price. That is, we need to prove that $U_R \geq U_A$. We first compute U_R, the utility when buyer 2 rejects the price:

$$U_R = \frac{F_1(t_2)}{2} \frac{\delta(v_2 - a^*)}{1 - \delta}. \tag{3.4.3}$$

Equation 3.4.3 is almost identical to Equation 3.4.2 for U_R when $v_2 \geq t_2$. However, since $v_2 < t_2$, buyer 2 does not obtain any positive utility for future rounds if buyer 1 accepts the current price (after an update of the belief, the seller will keep posting the price of t_2).

Similarly, we can compute U_A, the utility when buyer 2 accepts the price as follows:

$$U_A = \left(F_1(t_2) + \frac{1 - F_1(t_2)}{2} \right) (v_2 - p_k). \tag{3.4.4}$$

Equation 3.4.4 is almost identical to Equation 3.4.1 for U_A when $v_2 \geq t_2$. However, since $v_2 < t_2$, there is no future utility as a_{k+1} will be updated to t_2 based on Algorithm 3. Hence, we can compute from Equation 3.4.3 and 3.4.4 that:

$$\begin{aligned} U_A - U_R &= \frac{1 + F_1(t_2)}{2}(v_2 - p_k) - \frac{F_1(t_2)}{2} \frac{\delta(v_2 - a^*)}{1 - \delta} \\ &= (v_2 - t_2) \left(\frac{1 + F_1(t_2)}{2} - \frac{F_1(t_2)\delta}{2(1 - \delta)} \right) \end{aligned}$$

using the threshold equation (Equation 3.3.2). Furthermore, recall that Equation 3.3.2 states that:

$$(t_2 - p) \left(F_1(t_2) + \frac{1 - F_1(t_2)}{2} \right) = \frac{t_2 - a^*}{2} F_1(t_2) \frac{\delta}{1 - \delta}.$$

Notice that $p \geq a^*$ is true because if $p < a^*$, then the seller can achieve more revenue by posting a price of $p = a^*$. Then, since $t - p < t - a^*$, we can obtain:

$$\frac{1 + F_1(t_2)}{2} \geq \frac{F_1(t_2)\delta}{2(1 - \delta)}.$$

Then, since $v_2 < t_2$, we see that

$$U_A - U_R \leq 0.$$

Hence, we have verified that when $v_2 < t_2$, rejecting the current price is optimal.

2. Consider buyer 1. First assume that $v_1 \geq t_2$. The utility for buyer 1 accepting p_k, denoted by U_A, satisfies the following equation:

$$U_A \geq \left(F_2(t_2) + \frac{1 - F_2(t_2)}{2}\right)(v_1 - p_k) + \frac{\delta}{1-\delta}(v_1 - t_2)F_2(t_2) + U_{k+1}(1 - F_2(t_2)). \tag{3.4.5}$$

where U_{k+1} is the utility obtained in future rounds when both buyers accept the current price. Equation 3.4.5 is almost symmetric to Equation 3.4.1 for buyer 2. However, recall that the threshold $t = \max(t_1, t_2) = t_2$. Therefore, in both the first term and the second term, t_2 instead of t_1. Similarly, the following holds for U_R, the utility for buyer 1 after rejecting the current price:

$$U_R \leq F_2(t_2)\frac{v_1 - a^*}{2}\frac{\delta}{1-\delta} + \delta(1 - F_2(t_2))\frac{v_1 - t_2}{2}. \tag{3.4.6}$$

Equation 3.4.6 is almost symmetric to Equation 3.4.2 for buyer 2. However, again, we use t_2 instead of t_1.

Recall that we need to show that $U_A \geq U_R$ (Algorithm 2, Line 7). We can now compute:

$$\begin{aligned}U_A - U_R = (v_1 - t_2 + t_2 - p_k)\frac{1 + F_2(t_2)}{2} - (v_1 - t_2 + t_2 - a^*)\frac{F_2(t_2)\delta}{2(1-\delta)} \\ + (v_1 - t_2)\left(\frac{\delta F_2(t_2)}{1-\delta} - \frac{\delta(1 - F_2(t_2))}{2}\right) + U_{k+1}(1 - F_2(t_2)).\end{aligned}$$

Using Assumption 3.2.1, we can compute that:

$$\begin{aligned}&U_A - U_R \\ &\geq (v_1 - t_2)\left(\frac{1 + F_2(t_2)}{2} - \frac{\delta F_2(t_2)}{2(1-\delta)} + \frac{\delta F_2(t_2)}{1-\delta} - \frac{\delta(1 - F_2(t_2))}{2}\right) \\ &\qquad + U_{k+1}(1 - F_2(t_2)) - \epsilon.\end{aligned} \tag{3.4.7}$$

We now assume that ϵ is small enough and satisfy

$$U_{k+1}(1 - F_2(t_2)) - \epsilon \geq 0.$$

Then, we can see from Equation 3.4.7 that:

$$U_A - U_R \geq (v_1 - t_2)\left(\frac{\delta F_2(t_2)}{2(1-\delta)} + \frac{(1-\delta)+(1+\delta)F_2(t_2)}{2}\right).$$

Since $v_1 \geq t_2$ and $0 < \delta < 1$, both terms are positive, and so we can conclude that $U_A - U_R \geq 0$.

It remains to discuss the case when $v_1 < t_2$. We can calculate the utility for buyer 1 when the current price is rejected as follows:

$$U_R = \frac{F_2(t_2)}{2}\frac{\delta(v_1 - a^*)}{1-\delta}. \tag{3.4.8}$$

Again, Equation 3.4.8 is almost symmetric to Equation 3.4.3. However, instead of t_1, we use t_2 because t_2 is the threshold here. Similarly, we obtain the utility for buyer 1 when the current price is accepted:

$$U_A = \left(F_2(t_2) + \frac{1-F_2(t_2)}{2}\right)(v_1 - p_k). \tag{3.4.9}$$

Equation 3.4.9 is almost symmetric to Equation 3.4.4 except the threshold is kept the same.

We now need to show that $U_R \geq U_A$. We compute that:

$$\begin{aligned} U_A - U_R &= \frac{1+F_2(t_2)}{2}(v_1 - p_k) - \frac{F_2(t_2)}{2}\frac{\delta(v_1 - a^*)}{1-\delta} \\ &\leq (v_1 - t_2)\left(\frac{1+F_2(t_2)}{2} - \frac{F_2(t_2)\delta}{2(1-\delta)}\right) \end{aligned}$$

using Assumption 3.2.1. Similarly to the case for buyer 2, we can conclude that $\frac{1+F_2(t_2)}{2} - \frac{F_2(t_2)\delta}{2(1-\delta)} \geq 0$. Since $v_1 < t_2$, we see that

$$U_A - U_R \leq 0.$$

We have shown that this algorithm is optimal for buyer 1. Hence, we have shown that Lemma 3.4.5 is true for the exploration phase.

We now consider the exploitation phase. Recall that both buyers accept the price if and only if $v_i \geq p_k$ and $p_k \leq a^* = \max(a_1, a_2)$ (Algorithm 2, Line 10-11).

1. Consider the case when $p_k > a^*$. Then, we want to show that rejecting is better than accepting the price. Recall that if we accept the current price, the seller will update the belief to $[b, b]$ (Algorithm 3, Line 5). Hence, there will be no positive utility from future rounds. Therefore, we obtain that:

$$U_A = v_i - p_k. \tag{3.4.10}$$

On the contrary, if we reject the current price, we do not obtain any utility from the current round. However, we can compute the utility gained from future rounds as follows:

$$U_R \geq \frac{\delta}{1-\delta}\frac{v_i - a^*}{2}. \tag{3.4.11}$$

We obtain this because buyer i gets the item with the probability of at least $\frac{1}{2}$ (could be 1) where δ is a time-discount factor. Then, we can compute that:

$$\begin{aligned} U_R - U_A &\geq \frac{\delta}{1-\delta}\frac{v_i - a^*}{2} - (v_i - p_k) \\ &\geq (v_i - a^*) - (v_i - p_k) \text{ since } \delta \geq \frac{2}{3} \\ &= p_k - a^* \\ &> 0. \end{aligned}$$

Hence, we have shown that rejecting the current price is better than accepting the current price.

2. Consider the case when $p_k \leq a^*$. Then, we want to show that accepting is better than rejecting the price if and only if $v_i \geq p_k$. The utility for accepting the price, U_A is:

$$U_A = \frac{1}{1-\delta}\frac{v_i - p_k}{2} \tag{3.4.12}$$

where U_{k+1} indicates the utility gained from all future rounds. Furthermore, we obtain

$$U_R = 0 \tag{3.4.13}$$

since the seller updates the belief to $[b, b]$ after an off-path behavior taken by the buyer (Algorithm 3, Line 12). Notice here that $U_A \geq 0 = U_R$ if and only if $v_i \geq p_k$. Therefore, accepting is better than rejecting the price if and only if $v_i \geq p_k$.

41

We have now shown that the algorithm is optimal for buyers in the exploitation, and hence we have shown that the Algorithm 2 is optimal. □

With Lemma 3.4.5 and 3.4.2, we have shown that the algorithm achieves a PBE.

3.5 Analysis of Assumptions

In this section, we discuss the validity of Assumption 3.2.1:

$$-2\epsilon \leq (F_2(t_2)+1)(t_2-p_k) - F_2(t_2)(t_2-a^*)\frac{\delta}{1-\delta} \leq 0$$

when $t_2 \geq t_1$.

We first consider the validity of

$$(F_2(t_2)+1)(t_2-p_k) - F_2(t_2)(t_2-a^*)\frac{\delta}{1-\delta} \leq 0.$$

Recall that $F_2(t_2) \to 1$ as $t_2 \to b_2$ by the definition of CDF and also $\frac{\delta}{1-\delta} \geq 2$ since $\delta \geq \frac{2}{3}$. Then, we can compute that:

$$\begin{aligned}
&\lim_{t_2\to b_2}\left[(F_2(t_2)+1)(t_2-p_k) - F_2(t_2)(t_2-a^*)\frac{\delta}{1-\delta}\right] \\
&= 2(t_2-p_k) - \frac{\delta}{1-\delta}(t_2-a^*) && (\text{since } F_2(t_2)\to 1 \text{ as } t_2\to b_2) \\
&\leq (t_2-p_k)-(t_2-a^*) && \left(\text{since } \delta \geq \frac{2}{3}\right) \\
&\leq 0. && (\text{since } p_k \geq a^*)
\end{aligned}$$

Now, we claim that the result above shows that

$$(F_2(t_2)+1)(t_2-p_k) - F_2(t_2)(t_2-a^*)\frac{\delta}{1-\delta} \leq 0.$$

For the sake of contradiction, suppose

$$(F_2(t_2)+1)(t_2-p_k) - F_2(t_2)(t_2-a^*)\frac{\delta}{1-\delta} > 0.$$

Then, by the intermediate value theorem, we can say that there exists $t_3 > t_2$ that satisfies

$$(F_2(t_3) + 1)(t_2 - p_k) - F_2(t_3)(t_3 - a^*)\frac{\delta}{1-\delta} = 0.$$

However, this contradicts with the fact that t_1 is the biggest threshold that satisfies the threshold equation. Hence,

$$(F_2(t_2) + 1)(t_2 - p_k) - F_2(t_2)(t_2 - a^*)\frac{\delta}{1-\delta} \leq 0.$$

Therefore, we can conclude that $\delta \geq \frac{2}{3}$ simply is a sufficient condition for this part of the assumption.

We now discuss

$$-2\epsilon \leq (F_2(t_2) + 1)(t_2 - p_k) - F_2(t_2)(t_2 - a^*)\frac{\delta}{1-\delta}.$$

Specifically, this is a sufficient condition to ensure that $(F_2(t_2) + 1)(t_2 - p_k) - F_2(t_2)(t_2 - a^*)\frac{\delta}{1-\delta}$ is small enough in magnitude that accepting is the optimal choice for buyer 1 when $v_1 \geq t_2$. Notice that we do not need to satisfy the inequality for all $\epsilon > 0$. It suffices that the inequality is satisfied for an ϵ that is bigger than δU_{k+1} for relevant U_{k+1} in the exploration phase. Since we can check $\delta U_{k+1} \geq \epsilon$ for all rounds in the exploration phase, we can check if the assumption is satisfied for each auction.

We now interpret the meaning of this assumption. The assumption simply states that the difference between the utility of accepting the price and the utility of rejecting the price is small enough when $v_1 = t_2$. Recall that Equation 3.3.1 states the follows:

$$(t_1 - p)\left(F_2(t_1) + \frac{1 - F_2(t_1)}{2}\right) = \frac{t_1 - a^*}{2}F_2(t_1)\frac{\delta}{1-\delta}.$$

We can see that the right-hand side of the inequality is zero at t_1. If t_1 and t_2 are close to each other, then the assumption should be satisfied. Then, in the scenarios where the value distributions for buyer 1 and buyer 2 are similar, then the assumption is perhaps satisfied. However, this discussion does not have any proof currently, and require some future work.

43

4

Conclusion and Directions for Future Work

In this chapter, we summarize our main result in Section 4.1, and then, suggest directions for future work in Section 4.2.

4.1 Conclusion

Investigating the equilibrium structure of repeated games sees an important application in online sales these days. In this thesis, we have introduced relevant literature on the repeated games in auction theory. Although there is no Perfect Bayesian Equilibrium for repeated sales for more than 3 rounds as shown in (Devanur et al. 2015), previous work have shown that we can find interesting equilibrium structures or near-optimal auctions by imposing heavy conditions on either sellers' or buyers' behaviors.

In this thesis, we extended a previous result for multi-buyer repeated auction by removing an i.i.d assumption on buyers' distributions. Our main result (Claim 3.2.1) states that there is a Perfect Bayesian Equilibrium for two-buyer repeated auction under Assumption 3.2.1. We have justified this claim by giving a full description of the algrithms used by buyers and sellers and proving its optimality. This result shows an existence of an equilibrium in a more general setting, which can be applied to wider real-world situations.

44

4.2 Directions for Future Work

Although we have shown an existence of an equilibrium, Assumption 3.2.1 is a very heavy assumption. Some parts of the assumption is not intuitively interpretable, and it is not clear when it can be satisfied. Hence, we recommend that future work attempt to reduce the assumption and to interpret the assumption. Below, we give some different concrete steps that future work can take.

To reduce the assumption, there might be a need for a more complicated algorithm. We were not able to explore enough different algorithms taken by buyers and sellers for a possible equilibrium. However, it is possible that there is an equilibrium without an assumption, for instance, by using different thresholds.

Another approach that one could take is to impose a more interpretable assumption. The current assumption is hard to interpret. However, we could try imposing more interpretable assumptions and see if we could still prove an optimality. For example, we could try following approaches:

1. Use same distributions with different parameters such as mean μ or variance σ; or

2. Use named distributions such as a log-normal distributions or an exponential distribution.[1]

Although this imposes a heavier assumption on buyers' distributions, these assumptions are less heavier than the i.i.d assumptions, and the interpretation of these assumptions is more clear.

We could also take a completely different approach to extend the work in (Chawla et al. 2016). Recall from Theorem 2.3.11 that a simple pricing-scheme achieves a near-optimal revenue with the Martingale condition on buyers' evolving values. We could extend this result to the case when the buyers' values is monotonically decreasing. This value scheme will reflect some situations where more use of an item reduces the value of the item. Although this approach is not a direct extension

[1] Note that we propose a log-normal buyer distribution rather than a normal distribution because a log-normal distribution always takes a positive value while a normal distribution takes any real number.

of our work in this thesis, this extension could help us understand a more general equilibrium structure in a repeated auction with multiple buyers.

46

References

Abernethy, Jacob D, Rachel Cummings, Bhuvesh Kumar, Sam Taggart, and Jamie H Morgenstern (2019). "Learning Auctions with Robust Incentive Guarantees". In: *Advances in Neural Information Processing Systems*, pp. 11591–11601 (cit. on p. 13).

Agrawal, Shipra, Constantinos Daskalakis, Vahab S. Mirrokni, and Balasubramanian Sivan (2018). "Robust Repeated Auctions under Heterogeneous Buyer Behavior". In: *Proceedings of the 2018 ACM Conference on Economics and Computation* (cit. on p. 14).

Amin, Kareem, Afshin Rostamizadeh, and Umar Syed (2013). "Learning Prices for Repeated Auctions with Strategic Buyers". In: NIPS'14, pp. 622–630 (cit. on p. 13).

Azar, Pablo, Jing Chen, and Silvio Micali (2012). "Crowdsourced Bayesian Auctions". In: *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference* (cit. on p. 14).

Azar, Pablo, Silvio Micali, Constantinos Daskalakis, and S Matthew Weinberg (2013). "Optimal and efficient parametric auctions". In: *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*. SIAM'13, pp. 596–604 (cit. on p. 14).

Balcan, Maria-Florina and Florin Constantin (2010). "Sequential Item Pricing for Unlimited Supply". In: *Proceedings of the 6th International Conference on Internet and Network Economics*. WINE'10, pp. 50–62 (cit. on p. 13).

Balcan, Maria-Florina, Avrim Blum, and Yishay Mansour (2008). "Item Pricing for Revenue Maximization". In: *Proceedings of the 9th ACM Conference on Electronic Commerce*. EC'08, pp. 50–59 (cit. on p. 13).

Charm, Tamara, Becca Coggins, Kelsey Robinson, and Jamie Wilkie (2020). *The great consumer shift: Ten charts that show how US shopping behavior is changing*. URL: https://www.mckinsey.com/business-functions/marketing-and-sales/our-insights/the-great-consumer-shift-ten-charts-that-show-how-us-shopping-behavior-is-changing (cit. on p. 2).

Chawla, Shuchi, Hu Fu, and Anna Karlin (2014). "Approximate Revenue Maximization in Interdependent Value Settings". In: *Proceedings of the Fifteenth ACM Conference on Economics and Computation*. EC'14, pp. 277–294 (cit. on p. 14).

Chawla, Shuchi, Nikhil R. Devanur, Anna R. Karlin, and Balasubramanian Sivan (2016). "Simple Pricing Schemes For Consumers With Evolving Values". In: *Proceedings of the 2016 Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA'16, pp. 1476–1490 (cit. on pp. 13, 19–21, 45).

Cheng, Yu, N. Gravin, K. Munagala, and Kangning Wang (2018). "A Simple Mechanism for a Budget-Constrained Buyer". In: *International Conference on Web and Internet Economics* (cit. on p. 13).

Conitzer, Vincent, Curtis R. Taylor, and Liad Wagman (2012). "Hide and Seek: Costly Consumer Privacy in a Market with Repeat Purchases". In: *Marketing Science* 31.2, pp. 277–292 (cit. on p. 13).

Debreu, Gerard (1954). “Representation of a preference ordering by a numerical function”. In: *Decision processes* 3, pp. 159–165 (cit. on p. 5).

Devanur, Nikhil R., Yuval Peres, and Balasubramanian Sivan (2015). “Perfect Bayesian Equilibria in Repeated Sales”. In: *Proceedings of the 2015 Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA'15, pp. 983–1002 (cit. on pp. 2, 4, 5, 7, 8, 12, 14–17, 23, 44).

Dhangwatnotai, Peerapong, Tim Roughgarden, and Qiqi Yan (2010). “Revenue Maximization with a Single Sample”. In: *Proceedings of the 11th ACM Conference on Electronic Commerce*. EC'10, pp. 129–138 (cit. on p. 14).

Ghosh, Arpita and Aaron Roth (2011). “Selling Privacy at Auction”. In: *Proceedings of the 12th ACM Conference on Electronic Commerce*. EC'11, pp. 199–208 (cit. on p. 13).

Hart, Oliver D. and Jean Tirole (1988). “Contract Renegotiation and Coasian Dynamics”. In: *The Review of Economic Studies* 55.4, pp. 509–540 (cit. on p. 12).

Hummel, Patrick (2018). “Reserve Prices in Repeated Auctions”. In: *International Journal of Game Theory* 47, pp. 273–299 (cit. on p. 13).

Immorlica, Nicole, Brendan Lucier, Emmanouil Pountourakis, and Samuel Taggart (2017). “Repeated Sales with Multiple Strategic Buyers”. In: *Proceedings of the 2017 ACM Conference on Economics and Computation*. EC'17, pp. 167–168 (cit. on pp. 2, 12, 17–19, 23–31, 33, 34, 36).

Karlin, A.R. and Y. Peres (2017). *Game Theory, Alive*. Miscellaneous Book Series. American Mathematical Society (cit. on pp. 2, 4, 6, 8, 12).

Krishna, Vijay (2009). *Auction Theory*. Elsevier Science (cit. on pp. 2, 4, 5, 12).

Myerson, Roger B. (1981). “Optimal Auction Design”. In: *Mathematics of Operations Research* 6.1, pp. 58–73 (cit. on p. 12).

Schmidt, Klaus M. (1993). “Commitment through Incomplete Information in a Simple Repeated Bargaining Game”. In: *Journal of Economic Theory* 60.1, pp. 114 –139. ISSN: 0022-0531 (cit. on p. 12).

Vickrey, William (1961). “Counterspeculation, Auctions, and Competitive Sealed Tenders”. In: *The Journal of Finance* 16.1, pp. 8–37 (cit. on p. 12).

48