



How Efficient Is Vinmonopolet's Pricing Policy?

A Structural Estimation of the Norwegian Wine Auctions

Emil Mathias Strøm Halseth

Ola Kristoffer Nestvold

Supervisors: Malin Arve and Morten Sæthre

Master's Thesis in Economic Analysis

NORWEGIAN SCHOOL OF ECONOMICS

This thesis was written as a part of the Master of Science in Economics and Business Administration at NHH. Please note that neither the institution nor the examiners are responsible – through the approval of this thesis – for the theories and methods used, or results and conclusions drawn in this work.

Acknowledgements

Writing this thesis has been an exciting and educational process, and we are grateful for the opportunity to combine several interests of ours in this project. We have both enjoyed the journey of learning more about electronic auctions and how to apply advanced econometric methods.

We are grateful for our advisors Malin Arve and Morten Sæthre for giving us valuable advice throughout this endeavour. Their commitment to and knowledge of auction theory and empirical methods have been highly inspiring, and paramount to our process of writing this master's thesis.

Finally, we would like to thank Vinmonopolet for giving us access to the relevant data. Without it, we would not have had the opportunity to carry out this master's thesis. We would like to express our gratitude to Sara Frimandslund in particular for her very helpful expertise on wine and Vinmonopolet's wine auctions.

Norwegian School of Economics

Bergen, June 2019



Emil Mathias Strøm Halseth



Ola Kristoffer Nestvold

Abstract

This paper uses an empirical approach to investigate how Vinmonopolet's current pricing policy affects their electronic wine auctions by looking at how the binding reserve price affects the auction outcomes and how the current policy fares against the theoretical optimal reserve price. We use comprehensive data from Vinmonopolet, containing information on all electronic wine auctions hosted in Norway since 2013. By estimating the bidders' willingness to pay and bidder participation, we are able to conduct counterfactual analyses on how changes in the reserve price and bidder participation affects both expected revenue and allocative efficiency. We find it is implausible that Vinmonopolet determines the reserve price on either criterion of revenue maximization or optimal allocation, since the reserve price is not conditioned on the sellers' valuation. However, Vinmonopolet's current pricing policy – setting the reserve price equal to 80% of their value assessment – fares well against the theoretical optimal reserve price in the majority of the auctions. This is largely due to reserve prices having a small effect on the wine auctions in general. Finally, we show that bidder participation is the main driver of the expected revenue of an auction, and increasing bidder participation reduces the probability of the reserve price affecting the auction outcomes. Hence, increasing bidder participation serves as a possible remedy for inefficient reserve prices.

Keywords – Electronic Auctions, Policy Analysis, Mechanism Design, Structural Estimation, Wine Economics

Contents

1	Introduction	1
1.1	Outline	3
2	Background	4
2.1	Vinmonopolet	4
2.2	Vinmonopolet's Electronic Wine Auction	5
2.2.1	Participation and Vinmonopolet's Value Assessment	5
2.2.2	The Bidding Process	6
2.2.3	The Transaction	6
2.3	Why Auctions?	7
3	Wine Auctions Data	8
4	The Theoretical Model of Vinmonopolet's EAs	11
4.1	The Model Environment	11
4.2	Equilibrium Bidding	12
4.3	Expected Revenue	14
4.4	Binding Reserve Price	15
5	The Econometric Model of Vinmonopolet's EAs	17
5.1	Structural Estimation	17
5.2	Notations and Definitions	18
5.3	Model Specification	19
5.4	Identification of the Model	20
5.4.1	Observation of the Second-Order Statistic	20
5.4.2	Dealing with Unobserved N	21
5.4.3	Imposing Structure on $F_v(\cdot)$ and $F_N(\cdot)$	21
5.5	Defining the Log-Likelihood Function	24
6	Analysis	25
6.1	Obtaining the Results	25
6.1.1	Parameter Estimation	25
6.1.2	Interpreting Magnitudes and Illustrating the Distributions	28
6.2	Vinmonopolet's Value Assessment	31
6.3	Counterfactual Analysis	32
6.3.1	Calculating the Optimal Reserve Price	32
6.3.2	The Impact of the Reserve Price on Expected Revenue	34
6.3.3	The Impact of Bidder Participation on Expected Revenue	39
7	Reflections and Concluding Remarks	43
7.1	Conclusion	43
7.2	Reflections On The Model	44
7.2.1	Bid Increments	44
7.2.2	Bidders' Arrival Process	45
7.2.3	One Shot, Static Game Assumption	45

References	47
Appendix	49
A1 Testing the IPV Assumption	49
A2 Simulations	51
A3 Regressing the Number of Observed Bidders on Vinmonopolet's Value Assessment	53
A4 Supplementary Graphs: Revenue Expectation and Variance as a Function of the Reserve Price	55

List of Figures

3.1	Number of Auctions and Amount of Revenue Generated per Round in NOK.	8
3.2	Shares of Observations for the Three Biggest Regions.	9
3.3	Box-Plots of the Number of Observed Participants and the Number of Auctions per Round.	10
6.1	Expected Potential Participation Across Time.	28
6.2	Effects of Changing the Scale (η) and Shape (β) Parameters of the Weibull Distribution.	29
6.3	Estimated Weibull Densities for Bordeaux, Burgundy and Piedmont Wine.	30
6.4	Estimated Poisson Probabilities for All Three Regions.	31
6.5	Expected Utility (EU) and the Standard Deviation of EU as a Function of the Reserve Price for a Bordeaux Wine Auction with $V_v = 20000$	35
6.6	Expected Utility (EU) and the Standard Deviation of EU as a Function of the Reserve Price for a Bordeaux Wine Auction with $V_v = 50000$	36
6.7	Expected Utility as a Function of the Poisson Parameter (λ) for a Bordeaux Wine with $V_v = 20000$ and $V_v = 50000$ in Round 23.	40
A1.1	Non-Parametrically Fitted Graph Between Winning Bids and Active Participants.	50
A2.1	Estimated Densities with 30 000 Observations.	51
A2.2	Estimated Densities with 2000 Observations.	52
A3.1	Vinmonopolet's Value Assessment Regressed on the Number of Observed Bidders.	54
A4.1	Expected Utility (EU) and the Standard Deviation of EU as a Function of the Reserve Price for a Bordeaux Wine with $V_v = 5000$	55
A4.2	Expected Utility (EU) and the Standard Deviation of EU as a Function of the Reserve Price for a Burgundy Wine with $V_v = 5000$	56
A4.3	Expected Utility (EU) and the Standard Deviation of EU as a Function of the Reserve Price for a Piedmont Wine with $V_v = 5000$	57
A4.4	Expected Utility (EU) and the Standard Deviation of EU as a Function of the Reserve Price for a Burgundy Wine with $V_v = 50000$	58

List of Tables

3.1	Descriptive Statistics for Bordeaux, Burgundy and Piedmont.	9
3.2	Increment Classes and the Number of Observations in Each Class for the Three Regions.	10
6.1	Estimates of the Shape (β), Scale (η) and Poisson (λ) Parameters.	26
6.2	Optimal Reserve Prices for Bordeaux and Burgundy.	33
6.3	Optimal Reserve Prices for Piedmont.	34
A1.1	Estimates for the CV <i>vs.</i> IPV Test.	49
A2.1	The Results from Testing Our Model with 30000 Observations.	51
A2.2	The Results from Testing Our Model with 2000 Observations.	52
A3.1	Regression of the Number of Observed Bidders on Vinmonopolet's Value Assessment.	53

1 Introduction

Vinmonopolet, a state-owned enterprise, was in 2012 awarded the exclusive rights to facilitate all sale of alcoholic beverages between consumers in Norway using auctions.¹ It chose to use an auction design that allowed consumers to buy wines from other consumers via an e-commerce website, an auction format often referred to as Electronic Auctions (EAs). Vinmonopolet's EAs were launched in 2013, and was immediately a success in the Norwegian market.² ³ Since its launch, Vinmonopolet has held 24 rounds of auctions and has sold objects for over NOK 62 million during the November 2013 to November 2018 period. Most of the objects put up for auction ends up being sold, with a success rate of 97.9%.

This study aims to investigate how Vinmonopolet's current pricing policy affects the electronic wine auctions. More specifically, we seek to investigate how the minimum amount required to bid on a wine – commonly referred to as the binding reserve price – influences the expected revenue of an auction, and what other consequences setting a binding reserve price has on the auctions. For example, is the current pricing policy scheme aiming to allocate the wine to the consumer who has the highest valuation of the wine, or is Vinmonopolet trying to maximize the expected revenue generated by the auction? How is the current pricing policy – setting the binding reserve price equal to 80% of Vinmonopolet's value assessment – faring against the theoretical optimal reserve price? In other words, what are the consequences of having a reserve price that is too high or too low relative to the theoretical optimal reserve price?

To answer our proposed questions, we need information about the fundamentals in the electronic wine auctions: the number of potential participants, and how much these participants are willing to pay for the good. However, the bidders' willingness to pay for the wine is not directly observed. Thus, we will utilize the data we have available on bids, participation and reserve prices to infer the bidders' underlying valuation distribution, which captures the potential participants' willingness to pay. Additionally – as is typical in

¹Consumers refers to all consumers without a license to sell alcohol. Licenses are typically granted to restaurants, bars, social events etc.

²<https://www.nrk.no/kultur/ma-utvide-polets-vinauksjon-1.11146075>

³<https://www.dn.no/enorm-interesse-for-historisk-auksjon/1-1-2030616>

electronic auctions – the number of potential bidders in the wine auctions is an unknown measure that needs to be estimated as well. Using the two distributions, we will be able to conduct a counterfactual analysis of Vinmonopolet’s pricing policy.

To briefly outline what a counterfactual analysis entails, consider Vinmonopolet’s pricing policy in the electronic wine auctions. Vinmonopolet sets a minimum amount required to bid on a bottle of wine, i.e. the reserve price. In order to evaluate how the reserve price affects the expected revenue of an auction, it is necessary to hold the other primitives fixed – such as the latent valuation distribution and distribution of potential participants. Holding the primitives fixed allow us to compare and infer what would happen to the expected revenue if Vinmonopolet had used a different reserve price instead of its current reserve price.

Our motivation to write this thesis is based on three elements. First, we are interested in understanding electronic auctions in more depth because of its popularity in the world economy. EAs’ ability to reduce frictions by connecting a large pool of buyers and sellers, and mitigating traditional physical constraints are some of the factors driving the popularity of EAs as a market mechanism in the economy. To get a perspective on the relevance of EAs; consider eBay – an e-commerce company which is the biggest host of EAs in the world. It averaged \$23.64 billion per quarter in gross merchandise volume for the period Q4 2017 - Q4 2018.⁴ Another example is Google, which uses EAs for its advertising platform. Google reported \$28.9 billion in advertising for Q3 2018. Other well known examples of companies utilizing EAs as part of their business model includes Twitter and Facebook.

Furthermore, we are curious to investigate how Vinmonopolet’s pricing policy affects sellers and consumers in the electronic wine auctions due to the central role of Vinmonopolet in the Norwegian alcoholic beverages market, and since no real alternative exists for private individuals to trade alcoholic beverages among themselves in Norway. Thus, the decisions Vinmonopolet makes in designing the auctions will potentially have large implications.

Finally, websites that host EAs are data-rich environments, and the playing rules of the auctions are clear-cut and common knowledge. This makes it possible to precisely

⁴<https://www.statista.com/statistics/242267/e-bays-quarterly-gross-merchandise-volume-by-sales-format/>

define the expected equilibrium behavior in the auctions without imposing too strong assumptions, while at the same time mapping the theoretical predictions to the data. Hence, it is feasible to study economic concepts such as market efficiency and pricing policy by the means of counterfactual analysis.

In this research paper we show that it is implausible that Vinmonopolet sets the reserve price based on either criterion of revenue maximization or optimal allocation since the reserve price is not conditioned on the seller's valuation. However, Vinmonopolet's current pricing policy – setting the reserve price equal to 80% of their value assessment – fares well against the optimal reserve price in the majority of the auctions. This is largely due to reserve prices having a small effect on the wine auctions in general. Finally, we show that bidder participation is the main driver of the expected revenue of an auction, and increasing bidder participation reduces the probability of the reserve price affecting the auction outcomes. Hence, increasing bidder participation serves as a possible remedy for inefficient reserve prices.

1.1 Outline

This master's thesis will be organized as follows: Chapter 2 provides a brief overview of Vinmonopolet's purpose and responsibilities, explains the participation process of the EAs and discusses why auctions are an appropriate market mechanism for wine in the regulated alcoholic beverage market in Norway. Chapter 3 presents the data used in the estimation. In Chapter 4 we lay out the modeling environment, discuss the assumptions necessary for a tractable model of the auction environment, and present relevant theory regarding the bidding function, calculation of expected revenue and the implications of a binding reserve price. Chapter 5 outlines our methodological approach, derives the econometric model and proposes an appropriate parameterization. In Chapter 6 we present our results and investigate how Vinmonopolet's pricing policy and bidder participation affects expected revenue by the means of counterfactual analysis. Finally, in Chapter 7 we will close with some concluding remarks and reflect upon the validity of our model.

2 Background

In this chapter, we will provide a brief overview of Vinmonopolet's purpose and responsibilities, explain the participation process of Vinmonopolet's electronic wine auctions and discuss why auctions are an appropriate market mechanism for wine in the regulated alcoholic beverage market in Norway.

2.1 Vinmonopolet

Vinmonopolet is a government-owned alcoholic beverage retailer with a monopoly on all retail sale of alcohol beverages stronger than 4.75 percent in Norway. Vinmonopolet is subject to the Ministry of Health and Care Services, and its mission is to ensure responsible distribution of alcoholic beverages in Norway, primarily through limited access and high cost. Through its responsibility of being the only distributor of alcoholic beverages in Norway, it has played a pivotal role in removing the private economic profit motive in the sale of alcohol in Norway.

The Alcohol act, written in 1927, is the law Vinmonopolet abides by. Prior to 2012, it was illegal for consumers without a license to sell alcoholic beverages to trade among themselves, and consumers had to use international auctions to sell their alcoholic beverages. However, in 2012 the parliament of Norway broadened Vinmonopolet's scope by changing the Alcohol act to also include responsibilities regarding all sale of alcohol between consumers. From 2013 and onward, Vinmonopolet engaged a third party – Blomqvist – to carry out wine auctions using an electronic auction format. The electronic auction format closely resembles how eBay auctions work. See Lucking-Reiley (2000) for a guide for economists to electronic auctions. We will provide an overview of the auction format in the next section.

Vinmonopolet has hosted a total of 24 auction rounds in the period from November 2013 to November 2018, and on average, 757 objects were put up for auction in each auction round.⁵ Vinmonopolet hosted a total of 18156 auctions and sold objects worth a total of

⁵Several auctions are hosted in different time periods during a year, and we will refer to an auction as the sale of one specific object while an auction round refers to a period of time where several auctions are hosted simultaneously.

NOK 62.4 million in the aforementioned period. Most of the objects put up for auction ends up being sold, with a success rate of 97.9%. The wine auctions' popularity exceeded all expectations, and the third party had to expand its number of auction spots.^{6 7} In its infancy year just two auction rounds were held whereas in 2018 five auction rounds were conducted with a total number of 4017 auctions.

2.2 Vinmonopolet's Electronic Wine Auction

In the following section, we will explain the participation and bidding process of Vinmonopolet's electronic wine auction, as well as how the transaction is handled.

2.2.1 Participation and Vinmonopolet's Value Assessment

In order to participate as a buyer in the wine auctions, you must register a profile online from which you manage your bids and prospective items. The buyer can follow objects of interest and get notifications about new objects for sale. To sell an object on Vinmonopolet's wine auction you must be a private citizen or anyone with no other legal channel available to sell your wine.

Furthermore, the seller must get a value assessment of the wine undertaken by experts at Vinmonopolet, which is free of charge and demands no commitment from the seller should he nevertheless change his mind about selling the good. The value assessment serves as an indication of the wine's quality, which reflects characteristics like geographic origin down to the specific vineyard, brand, bottle size, vintage, storage, and bottle condition. Notably, the value assessment aggregates information that might not be readily available for potential buyers, which will be further discussed in chapter 4. The experts at Vinmonopolet has also stated that they adjust their valuation based on international prices in both auction and retail markets. Furthermore, they adjust their valuations based on historical prices on similar wines. Finally, the assessment will serve as the foundation for the current pricing policy of setting a binding reserve price equal to 80% of the value assessment.

⁶<https://www.nrk.no/kultur/ma-utvide-polets-vinauksjon-1.11146075>

⁷<https://www.dn.no/enorm-interesse-for-historisk-auksjon/1-1-2030616>

2.2.2 The Bidding Process

In line with well-known electronic auctions like the eBay auction, the participants are asked to submit a maximum bid that ideally represents their maximum willingness to pay for the object. The bid values are chosen from a discrete list of options that form a bid interval that starts from a lower bound (effectively a reserve price) and rises in discrete steps that increases at specific points.⁸ At any time, you are constrained to bidding at least one increment above the current standing price. A bidding algorithm places bids on the buyer's behalf, overbidding any standing price below the buyers stated maximum value, but never above. The first bid placed, pushes the standing price to the lower bound. Subsequent bids will initially push the standing price up with one increment, and if someone's maximum bid is one increment above the current standing price, the algorithm will keep overbidding. When the bidding reaches the second highest maximum bid plus a bid increment the standing price settles. It is possible to re-submit bids, and participants are notified if someone overbids them. Bids within the final 3 minutes of the auction extend the time by 3 minutes. When the auction ends the good is awarded to the highest placed maximum bid at the final standing price. In the event that multiple participants have placed the same maximum bid, the earliest submission trumps the subsequent ones.

2.2.3 The Transaction

Proceeds from an eventual sale are awarded to the seller after subtracting a 15% fee that is equally distributed between Vinmonopolet and Blomqvist. The buyer can either retrieve the wine directly from the storage facility in Oslo free of charge, or he or she can have the package delivered to one of Vinmonopolet's outlets in Bergen, Sandefjord, Stavanger, Trondheim, Ålesund or Hamar for a shipment fee of NOK 150 for every 12th bottle.⁹

⁸See Chapter 3.

⁹Up until April 2018, buyers could only retrieve the wine from the storage facility in Oslo.

2.3 Why Auctions?

One could ask; why use auctions to allocate the wine? In general, the benefit of auctions over alternatives like posted prices or pairwise negotiations in an unregulated market is that auctions function as an effective tool in overcoming two central challenges due to asymmetric information: who should be the buyer, and at which price?

In an environment where the seller wishes to maximize his or her expected revenue, he or she wants to identify the buyer with the highest willingness to pay. However, the true bidder valuations are the potential buyers' private information, and they have no incentive to reveal their true preferences. In an auction, the buyers are forced to at least partially reveal their true valuations of the good in order to win the auction. Since both the wine itself and the buyers' preferences on it is heterogeneous, these two issues are particularly prevalent in the wine market. Furthermore, in absence of an open wine market due to strict regulation in Norway, the wine auctions serve as an effective market mechanism that alleviates the allocative inefficiencies induced by the prohibition of private trade in alcoholic beverages.

Hence, it is also interesting to study the auction environment in order to uncover potential inefficiencies and investigate the effects of different policy measures such as changes in the reserve price and efforts to increase bidder participation. Due to the central role of regulation in the alcoholic beverages market in Norway, policy choices could have large impacts. However, to evaluate the counterfactuals of potential policy changes, we need to estimate the primitives of the auction environment – namely the latent valuation distribution and the potential participation distribution for Vinmonopolet's wine auctions. A natural starting point in this endeavor is the auction data that we retrieved from Vinmonopolet, which will be discussed in the following chapter.

3 Wine Auctions Data

The data set contains all EAs for alcoholic beverages held in Norway in the period of November 2013 - November 2018. We used an IT and software company to extract the data set, and no transformations were done on the data. The data covers 18156 EAs spread out over 24 rounds during the same period. On average, 757 objects were put up for auction in each round. Figure 3.1 displays the number of objects put up for sale in the period together with the amount of revenue raised for each round.




Figure 3.1: Number of Auctions and Amount of Revenue Generated per Round in NOK.

Wine is the dominant category in the data set, with 16787 observations or 92% of the total observations. In the wine category, it is predominantly red wine being put up for auction, yielding 11773 of the observations. After excluding all auctions where more than one object were put up for auction and filtering on bottles with a size of 0.75L, we are left with a data set of 3552 observations.

The wine in the data set originates from 12 different countries, but the two most significant ones are France and Italy with 76.5% and 15.9% of the observations, respectively. In a similar fashion, the three biggest regions constitute 82% of the data set, see Figure 3.2.




Figure 3.2: Shares of Observations for the Three Biggest Regions.

We allow the underlying valuation distribution and the distribution of potential participants to vary across regions, thus we exclude all regions with too few observations.¹⁰ This leaves us with the three biggest regions Bordeaux, Burgundy, and Piedmont for a total of 2781 observations as our working data set. Henceforth, any reference to the data set will be referencing the data set with 2781 observations of 0.75L single bottles of red wine from Bordeaux, Burgundy, and Piedmont. Note that the auction round in March 2017 is not represented, which leaves us with 23 auction rounds held in the period November 2013 - November 2018.

The data set contains information on the wine's vintage, district, and producer together with Vinmonopolet's value assessment of the wine (V_v). Table 3.1 displays descriptive statistics for the three regions Bordeaux, Burgundy, and Piedmont. Notably, Piedmont has a significantly shorter range of winning bids and V_v relative to Bordeaux and Burgundy.

	n	Vinmonopolet's Assessment				Winning Bid			
		Min	Median	Mean	Max	Min	Median	Mean	Max
Bordeaux	1492	100	2250	3417	80000	160	2500	3560	82000
Burgundy	946	200	2500	4987	80000	240	3200	6735	110000
Piedmont	343	200	1000	1399	10000	280	1300	1910	23000

Table 3.1: Descriptive Statistics for Bordeaux, Burgundy and Piedmont.

We have a detailed bid log for each auction, including the timing of bid submissions and the amount of each reported maximum bid, together with unique bidder IDs.¹¹ Figure

¹⁰More on this in Chapter 5.4.3.

¹¹Note that we only have access to the participants' maximum bids if they were overbid by other

3.3 depicts box-plots of the number of observed participants together with the number of auctions (observations) for each round in our data set.




Figure 3.3: Box-Plots of the Number of Observed Participants and the Number of Auctions per Round.

Denotes the number of auctions per round. Dots denote outliers.

As alluded to in Chapter 2, the auction bids are chosen from a discrete list of options that differ by a known bid increment, which increases as the price level increases. The increment steps in our data set is displayed in Table 3.2. The majority (82%) of the wine bottles have valuations that lie in the lowest price interval.

NOK	Price Intervals				
	(0, 5')	(5', 10')	(10', 20')	(20', 30')	(30', 100')
Increments (Δ)	100	200	500	1000	2000
Bordeaux	1238	198	47	6	3
Burgundy	701	179	47	0	19
Piedmont	339	4	0	0	0
Total	2278	381	94	6	22

Table 3.2: Increment Classes and the Number of Observations in Each Class for the Three Regions.

¹ The symbol ' denotes thousands.

participants in the same auction.

4 The Theoretical Model of Vinmonopolet's EAs

In this chapter, we are first going to lay out the modeling environment and discuss the assumptions required to produce a tractable model of the auction environment. Then, we derive the equilibrium bidding function for Vinmonopolet's electronic auctions, characterize the expected revenue of the seller, and discuss the role of a binding reserve price. We will derive the expected revenue formula and optimal reserve price under the assumption of a known number of potential participants N for the sake of exposition, since allowing N to be stochastic does not change the main insights proposed in this chapter.¹²

4.1 The Model Environment

We model the wine auction environment as a symmetric Independent Private Value (IPV), second-price auction with unit demand, risk neutral rational bidders and abstract away from bid increments.¹³ In a second-price auction, the winner of the auction pays the second highest bid as opposed to his or her own bid. This gives bidders an incentive to bid their true valuations, as we will see when we derive the equilibrium bidding function below. It has been a common assumption in the empirical literature to ignore the bid increments in online auctions; treating them as pure second-price auctions, and – in order to produce a tractable bidding model – we will do the same. The implications of this specification of the pricing rule will be discussed in chapter 7.

The participants in the wine auctions are assumed to have independent, private valuations of the wines represented by the stochastic variable V_i , which is identically and independently distributed (*i.i.d.*) for all i . The Cumulative Distribution Function (CDF) of V_i can be denoted $F_v(\cdot)$ and is common knowledge to all the auction participants. Each bidder knows how much he or she values the object for sale (bidder i 's realization of V_i is denoted v_i), but the information is private and independent of the other participants' valuations. Our IPV assumption is aligned with the expectations of the wine experts from Vinmonopolet;

¹²The counterfactual analysis in chapter 6 will take into account the stochastic nature of N .

¹³Introducing risk aversion does not change bidding behavior in a second-price auction.

the majority of the participants in the wine auctions are most likely going to consume the wine themselves rather than reselling it (Vinmonopolet's Podcast, 2018). See Appendix 1 for a formal test of the IPV assumption that corroborates our assumption in the sense that the participants seem to bid as if they have independent private values.

We can allow the parameters of the bidders' valuation distribution to be a function of various characteristics that are common knowledge to all the bidders, e.g. the wine's quality: $F_v(\cdot; \theta(q))$, where θ is the vector of parameters associated with the valuation distribution, and the quality component q is common knowledge to the bidders. Keep in mind that we are still within the IPV setting since there is no common value component to bidder i 's valuation v_i , and that since quality is common knowledge (i.e. known with certainty) there can not be inferred any new information from the rivals' bidding behavior or upon realization of the winning bid. The quality q can be interpreted as an objective piece of information that all the bidders are aware of that shapes the common beliefs on $F(\cdot)$, but how this quality component translates into bidder i 's valuation v_i is private information. We will come back to this quality component in chapter 5 when we propose an econometric model to estimate the distribution of V_i , and in particular, we will discuss how it relates to Vinmonopolet's value assessment.

Finally, we treat the auctions as one-shot games where the buyers' demand for any particular wine is independent of other wines up for auction and future auctions. This gives us a parsimonious and tractable model that allows us to focus on the key aspects of our research question. As is common in the auction literature, we thus abstract away from budget constraints and issues of substitution between goods and across time. The implications of these assumptions are discussed in chapter 7.

4.2 Equilibrium Bidding

In a second-price auction, bidders will find it optimal to bid their valuation, i.e. to tell the truth. To see how, consider the following setting for bidder i : Let as before v_i denote bidder i 's valuation, and let b_i denote bidder i 's bid and B denote the highest bid among the other participants. For bidder i , bidding v_i is better or just as good as any other bidding strategy b_i . This can be shown by looking at the two cases: (1) bidding b_i for

values $b_i > v_i$ and (2) bidding b_i for values $b_i < v_i$.

Case 1:

- $B > b_i > v_i$: Bidder i loses, and bidding b_i or v_i yields the same result.
- $b_i > v_i > B$: Bidder i wins and pays B . Bidding b_i or v_i yields the same result.
- $b_i > B > v_i$: Bidder i wins and pays B , but obtains a negative payoff as the payment is bigger than bidder i 's valuation. Bidder i is better off by bidding v_i .

Case 2:

- $B > v_i > b_i$: Bidder i loses, and bidding b_i or v_i yields the same result.
- $v_i > b_i > B$: Bidder i wins and pays B . Bidding b_i or v_i yields the same result.
- $v_i > B > b_i$: Bidder i loses. Bidder i would have been better off bidding v_i , wherein i would have won and obtained a positive payoff.

As we can see, in both cases bidder i is as well or better off by bidding his or her valuation. In other words, bidding your true valuation stochastically dominates any other bidding strategy, and the optimal strategy becomes:

$$b(v_i) = v_i, \text{ when } r \leq v_i \forall i$$

$$b(v_i) = 0, \text{ when } r > v_i \forall i$$

When we introduce a binding reserve price this argument holds for those bidders that have a valuation at or above the reserve price, while those that do not, abstain from participating.

As mentioned in chapter 2, there is a window of time in which the participants can place their bids, and they are allowed to resubmit their bids as many times as they wish as long as they bid above the current standing price. The key point is that all of the observed final bids are the true valuations of the bidders. This holds if we invoke the following conditions, as used by Song (2004):

- No participant ever submits a cutoff price greater than his or her valuation.
- At his or her final submission time, if his or her valuation is greater than the current

standing price, bidder i will submit a final bid equal to his or her valuation if he or she has not yet done so.

In principle this allows a range of equilibrium behaviors such as bidding your valuation immediately after arriving at the auction, waiting to bid your valuation at the last possible moment, or placing a low bid early on and eventually resubmitting a bid equalling your valuation at a later stage in the auction. However, under these conditions we can without loss of generality interpret the optimal bidding strategy $b(v_i) = v_i$ as the final bid placed in the auction by bidder i .

4.3 Expected Revenue

Presumably, the seller is interested in maximizing the expected revenue gained from selling the wine. In our case, the revenue is awarded to the seller after subtracting a 15% fee that is equally distributed between Vinmonopolet and Blomqvist. Riley and Samuelson (1981) show that within the symmetric IPV environment with risk neutral bidders, where potential bidders participate if their valuation is higher than the reserve price, the highest bidder wins the auction, and the equilibrium bidding functions are symmetric and strictly increasing, the expected revenue is shown to be:

$$N \int_r^{\bar{v}} [x f_v(x) - (1 - F_v(x))] F_v(x)^{N-1} dx$$

We can see that expected revenue is a function of the reserve price r , the underlying valuation distribution $F_v(x)$ and the number potential participants N . It can be shown that the expected revenue is strictly increasing in the number of participants N , which in the general framework can be attributed to increased competition and an increased chance of someone having a high valuation. Thus, any efforts to increase participation could potentially increase the expected revenue of the wine auction. It is shown in Bulow and Klemperer (1996) that a seller that runs an English auction with no reserve price with $N + 1$ symmetric bidders will earn more in expectation than a seller who can hold an auction with an optimal reserve price with N bidders. This implies that efforts to increase participation are paramount, which could be done by e.g. information diffusion

and advertisement, or reductions in participation costs.

4.4 Binding Reserve Price

The wine auctions hosted by Vinmonopolet all have a known and binding starting price. Hitherto, the reserve prices are determined by a static rule: 80% of the value assessment performed by Vinmonopolet's experts. A salient question then becomes: what is the optimal reserve price? Since the seller retains the wine if it goes unsold, we must add the utility of retaining the object to the aforementioned expected revenue formula, leaving us with the following formula for the combined expected utility of the seller, Vinmonopolet and Blomqvist:¹⁴

$$v_0 F_v(r)^N + N \int_r^{\bar{v}} [x f_v(x) - (1 - F_v(x))] F_v(x)^{N-1} dx$$

Maximizing the expected utility with respect to the reserve price yields:

$$r^* = v_0 + \frac{1 - F_v(r^*)}{f_v(r^*)} \quad (4.1)$$

We can see that the optimal reserve price is a function of the seller's valuation and the underlying valuation distribution of the bidders, and that the optimal reserve price is independent of N . Thus, with knowledge about the underlying preferences of the agents in the auction, it is possible to make general policy prescriptions that are invariant to the number of potential participants.

However, in the case of Vinmonopolet's wine auctions, it is not clear whose interests are to be considered in the choice of the reserve price. Since Vinmonopolet and Blomqvist take a percentage fee of the revenues, all parties would appreciate a higher expected revenue. Yet, in the case when the wine goes unsold it is returned to the seller, and Vinmonopolet and Blomqvist are left with nothing. Thus – if the seller has a non-zero valuation of the wine – there is a conflict of interest since the seller would be more or less willing to depart from

¹⁴This holds for risk neutral agents. We will briefly discuss the effects of risk aversion on the optimal reserve price in chapter 6.

the wine based on his preferences and would like the reserve price to reflect that. In the general auctions hosted by Blomqvist, the reserve price is determined jointly by Blomqvist and the seller. However, for Vinmonopolet's wine auctions the reserve price is determined by Vinmonopolet according to the 80% rule. If Vinmonopolet aims to maximize allocative efficiency, the optimal reserve price would be a reserve price that ensures that the good is awarded to the agent that values it the most (Dasgupta and Maskin, 2000). This is the case when the reserve price equals the seller's valuation: $r = v_0$. However, v_0 is not known to Vinmonopolet, and the expected welfare maximizing reserve price for Vinmonopolet would have to consider endogenous participation also on the seller's side, since a reserve price too misaligned with the seller's preferences might induce the potential seller to abstain from participating.

It is beyond the scope of this paper to explore the implications of endogenous seller participation, but some conclusions can be drawn. Given v_0 , Vinmonopolet should never use a higher reserve price than the total expected utility maximizing one if they aim for allocative efficiency, since this would yield strictly worse expected outcomes for both selling parties and buyers. Expected utility would be strictly lower for the claimants on the revenue, and the risk of forgoing a mutually beneficial trade would increase. In chapter 6 we will calculate optimal reserve prices in terms of the total utility of the seller, Blomqvist and Vinmonopolet – abstracting away from conflicts of interest between the selling parties – and we will investigate how Vinmonopolet's static pricing rule compares to the optimal reserve price in terms of total expected utility and allocation.

5 The Econometric Model of Vinmonopolet's EAs

In this chapter, we will outline our methodological approach of structural estimation and provide motivation for why the structural approach is advantageous in our specific application. Next, we will introduce the necessary notation, specify the econometric model, discuss issues of identification and propose a suitable parameterization of said model. Finally, in order to estimate the model parameters by the means of maximum likelihood estimation, we define the log-likelihood function.

5.1 Structural Estimation

In the words of Holmes and Sieg (2015): "A structural estimation is a methodological approach in empirical economics explicitly based on economic theory". In other words, a structural estimation can be viewed as a theory-based estimation, and the goal is to estimate the structural parameters of an internally consistent model. This contrasts with the reduced form approach that to a larger degree is based on statistical assumptions and is only implicitly based on economic theory. There are pros and cons to both approaches; the structural approach relies more on the validity of theoretical assumptions, whereas the reduced form approach is to a greater degree divorced from economic theory. The main endeavor of the reduced form approach is to construct research designs that can uncover causal effects in the sense of contrasting counterfactual outcomes.

For our purposes of investigating how the reserve price affects the expected revenue of an auction, the most viable method seems to be the structural approach. Common reduced-form methods are randomized treatments, making use of exogenous variation in the data and the difference in differences technique, which in our case could translate to randomly assigning different reserve prices to different treatment groups, using some exogenous change in the reserve price as an instrument or exploiting some natural experiment on the reserve price. However, we do not command such control over the reserve price, know of any exogenous changes in Vinmonopolet's 80% rule or have knowledge about any subgroup

of auctions exposed to some naturally occurring change in the reserve price.¹⁵ Furthermore, such an analysis would not leave us with an explicit link between the observable outcomes and the underlying economic structures such as the valuation and potential participation distributions. Instead, we seek to identify the causal effects of changes in the reserve price through restrictions implied by auction theory. Since we explicitly model the mechanisms through which expected revenue is determined, we can infer the consequences of changes in the reserve price even when no such change in the reserve price is observed in the data. We will follow the typical steps in structural estimation; model specification, identification and estimation, and policy analysis for the rest of this paper.

5.2 Notations and Definitions

The CDF of the potential bidders' valuations is denoted $F_v(\cdot)$ and defined to be *i.i.d.* with its corresponding Probability Density Function (PDF) denoted as $f_v(\cdot)$. Let N_t denote the number of potential participants in auction t and n_t denote the number of active participants for each auction t . Let $F_N(\cdot)$ denote the CDF of N_t and $P(N_t)$ denote the Probability Mass Function (PMF) since the number of participants is an integer value. For each auction, $y_t^{1:N_t}, \dots, y_t^{k:N_t}$ are the order statistics of the potential bidders' valuations with $y_t^{k:N_t}$ denoting the k th order statistic. The CDF of $y_t^{k:N_t}$ is denoted $F_y^{k:N_t}(\cdot)$ with the corresponding PDF denoted as $f_y^{k:N_t}(\cdot)$. We assume that $F_v(\cdot)$ is distributed according to the Weibull family and that $F_N(\cdot)$ follows a Poisson process.

The model of Vinmonopolet's electronic auction is composed of three elements: the PDF of the second-order statistic $f_y^{2:N_t}(\cdot)$, the conditional probability of n_t given N_t , $P(n_t|N_t)$, and the probability of N_t , $P(N_t)$. Together with the functional form specifications on $F_v(\cdot)$ and $F_N(\cdot)$, the model is identified whenever the second-order statistic, the binding reserve price and the number of active participants are observed.

¹⁵Keep in mind that any variation in the reserve price in our data set would be confounded by variation in the value assessment due to the 80% rule.

5.3 Model Specification

Since we only observe a truncated sample n_t of the potential bidders N_t , we use the second-order statistic of the truncated sample in our model. The PDF of the truncated second-order statistic $f_y^{2:n_t}(\cdot)$ is obtained from the underlying distribution of bidders' valuations $f_v(\cdot)$, where

$$f_y^{2:n_t}(y_t|n_t; \boldsymbol{\theta}) = n_t(n_t - 1)F_v(y_t|Y > r; \boldsymbol{\theta})^{n_t-2}(1 - F_v(y_t|Y > r; \boldsymbol{\theta}))f_v(y_t|Y > r; \boldsymbol{\theta}) \quad (5.1)$$

with $\boldsymbol{\theta}$ representing the vector of parameters associated with the underlying valuation distribution.

The truncated CDF and truncated PDF are defined as

$$F_v(y_t|Y > r; \boldsymbol{\theta}) = \frac{F_v(y_t; \boldsymbol{\theta}) - F_v(r; \boldsymbol{\theta})}{1 - F_v(r; \boldsymbol{\theta})}, \quad f_v(y_t|Y > r; \boldsymbol{\theta}) = \frac{f_v(y_t; \boldsymbol{\theta})}{1 - F_v(r; \boldsymbol{\theta})} \quad (5.2)$$

Dropping parameter notation and subscripts for simplicity, substituting 5.2 in 5.1 yields

$$f_y^{2:n}(y|n) = \frac{n(n-1)(F_v(y) - F_v(r))^{n-2}(1 - F_v(y))f_v(y)}{(1 - F_v(r))^n} \quad (5.3)$$

The probability of observing any number of active participants given the number of potential participants can be modeled as a binomial distribution. Thus, it is defined as

$$P(n_t|N_t; \boldsymbol{\theta}) = \binom{N_t}{n_t} (1 - F_v(r; \boldsymbol{\theta}))^{n_t} F_v(r; \boldsymbol{\theta})^{N_t - n_t} \quad (5.4)$$

To complete the model for auction t , the probability of each N_t is given by the Poisson distribution

$$P(N_t; \boldsymbol{\lambda}) = e^{-\boldsymbol{\lambda}} \cdot \frac{\boldsymbol{\lambda}^{N_t}}{N_t!} \quad (5.5)$$

where $\boldsymbol{\lambda}$ represents the vector of parameters associated with the underlying distribution of potential participants.

5.4 Identification of the Model

To uniquely determine the parameters of the underlying model that generates the observable data, one can use the variation from the data to map and estimate the set of parameters in our model. The vector θ is identified by mapping the PDF of the truncated second-order statistic to the number of active bidders n_t and the reserve price r_t , together with the Weibull parameterization of the underlying valuation distribution $f_v(\cdot)$.

To achieve identification, we must also be able to uniquely determine the parameters of the underlying distribution of potential bidders λ . We do this by mapping the probability of observing n_t active bidders to the reserve price and valuation distribution, together with the Poisson parameterization of the underlying distribution of potential bidders $P(N_t)$. This is feasible under the assumption that the only reason not to participate is that the reserve price exceeds the bidder's valuation.

Two main challenges arise in our effort to identify θ and λ : we need to establish that the second-order statistics are observed and deal with the fact that the number of potential bidders N is unobserved.

5.4.1 Observation of the Second-Order Statistic

To explain why we observe the second-order statistics for auctions with $n_t \geq 2$, we will utilize the established equilibrium bidding strategy for second-price auctions and the properties of the bidding algorithm used in Vinmonopolet's electronic auctions. Participants telling the truth $b(v_i) = v_i$, implies that all of the observed final bids in the auctions are the true valuations of the bidders. A key feature of the bidding algorithm, as discussed in chapter 3, is the iteration process. This ensures that the algorithm will iterate to the second highest bid plus the bid increment. Combining these two elements, we can claim that the transaction price for auction t with $n_t \geq 2$ is the second-order statistic $y_t^{2:N_t}$. For auctions with $n_t < 2$, there will either be no transaction ($n_t = 0$) or there will be a sale with the transaction price being the reserve price ($n_t = 1$). Thus, the transaction price can not be interpreted as the second-order statistic in these cases, which

will be accounted for in the log-likelihood function below.

5.4.2 Dealing with Unobserved N

A challenge in the structural estimation of electronic auctions is the fact that the number of potential bidders N is unobserved. There are two general reasons for this. First, some potential participants will realize that the starting price of the auction exceeds their valuation and choose not to participate. This issue is handled in our model by modeling the probability of participation as a function of the reserve price and the underlying valuation distribution.

Second, as participants place their bids during the auction, the standing price will adjust up to the current second highest maximum bid. Depending on the timing and sequence of placed bids, someone that has a valuation above the reserve price might not be able to place a bid if by the time of their arrival they have already been outbid by at least two other participants. The standing price would already have exceeded their valuation. Thus, we implicitly assume that the sequence of bidders' arrivals is in such a way that everyone that has a valuation above the reserve price is able to place their bid. The implications of this assumption are discussed in chapter 7.

5.4.3 Imposing Structure on $F_v(\cdot)$ and $F_N(\cdot)$

Our data contains binding reserve prices, which only allows us to observe a truncated sample of the valuations generated by the Weibull distribution; specifically those valuations that exceed the reserve price. To allow for identification of the part of the distribution whose realizations are truncated away, we impose a functional form on the underlying distribution of valuations and potential bidders so that we can extrapolate into the domain of the valuation distribution that is truncated by the reserve price. This permits counterfactual analyses on both increases and decreases in the reserve price as opposed to only increases, which could be feasible without functional form assumptions since the realizations from the domain of the distribution above the reserve price are observed in the data. Ex-ante we do not know whether or not the reserve price is too high or too low, so we believe the parametric assumptions are justified for our purposes.

We assume that the distribution of the valuations V_i is from the Weibull family, which is only defined for strictly positive outcomes. Thus, let $F_v(\cdot)$ and $f_v(\cdot)$ be denoted

$$F_v(x) = 1 - e^{-\left(\frac{x}{\eta}\right)^\beta}$$

$$f_v(x) = \frac{\beta}{\eta} \left(\frac{x}{\eta}\right)^{\beta-1} e^{-\left(\frac{x}{\eta}\right)^\beta}$$

where η and β are named the *scale* and *shape* parameters respectively. As is indicative of their names, the shape parameter determines the shape of the Weibull distribution – uniquely determining its skewness and coefficient of variance – while the scale parameter roughly speaking determines the magnitude of the valuations. Weibull is convenient in that it has few parameters and is flexible. It is also in accordance with studies using parametric estimation in the empirical literature (Harry J. Paarsch, 1992) (Athey et al., 2004) (Susan Athey, Dominic Coey, Jonathan Levin, 2011) (Canals-Cerdá and Percy, 2013).

We allow the scale parameter η and shape parameter β to be functions of a set of covariates:

$$\eta = \eta_0 + \eta_V \cdot V_v + \eta_B \cdot \mathbb{1}\{B\} + \eta_P \cdot \mathbb{1}\{P\} + \eta_{VB} \cdot V_v \cdot \mathbb{1}\{B\} + \eta_{VP} \cdot V_v \cdot \mathbb{1}\{P\}$$

$$\beta = \beta_0 + \beta_V \cdot V_v + \beta_B \cdot \mathbb{1}\{B\} + \beta_P \cdot \mathbb{1}\{P\} + \beta_{VB} \cdot V_v \cdot \mathbb{1}\{B\} + \beta_{VP} \cdot V_v \cdot \mathbb{1}\{P\}$$

where V_v is Vinmonopolet's value assessment, and $\mathbb{1}\{B\}$ and $\mathbb{1}\{P\}$ are dummy variables that indicate whether or not the wine originates from Burgundy or Piedmont respectively. Note that Bordeaux is the base group for the estimates. The covariates capture any differences in the demand for red wine between geographic origins, and that Vinmonopolet's value assessment provides information about the wine's quality. In the absence of a pure quality measure, we attempt to capture the quality component by using the value assessment as a proxy, even though the value assessment is influenced by factors beyond objective quality. The interaction terms allow the value assessment to affect the various regions differently, implying that the way quality is inferred from the value assessment can vary between regions. In principle, it is possible to make the analysis more granular by introducing additional characteristics like vintage, producer, branding and so on. However, for the purposes of our analysis, such granularity is not necessary.

We assume that the distribution of the number of potential participants N_t follows a Poisson process – see Equation 5.5 – which is characterized by a single parameter λ . The Poisson distribution is a simple and discrete probability distribution with valuable properties such as being parsimonious and non-negative, allowing us to model the probability of a given number of N_t participants. Let λ be a function of the following covariates:

$$\lambda = \lambda_0 + \lambda_B \cdot \mathbb{1}\{B\} + \lambda_P \cdot \mathbb{1}\{P\} + \lambda_t \cdot R_t$$

Like the parameterization of the valuation distribution, we allow the underlying distribution of potential participants to vary across regions. Additionally, we include a round indicator R_t , which allows the underlying distribution of potential participants to vary across the 23 rounds of wine auctions included in our sample. This opens up for changes in participation over time.

We choose to not include V_v because the quality of a good should determine a bidder's willingness to pay, but it is not clear how it should affect how many potential bidders there are. It is hard to pinpoint what mechanism determines the number of potential bidders in an electronic auction since there is free entry throughout the auction round, and we have abstained from discussing the issue beyond the fact that such a number exists. Furthermore, the value assessment draws on a complex set of information beyond the objective quality of the wine, so if the value assessment is related to the number of potential participants in any way, the variable would be difficult to interpret and not of direct interest to the counterfactual analyses that we intend to perform. In Appendix 3, we provide a regression of Vinmonopolet's value assessment V_v on the number of active participants n_t , and there is no indication that the value assessment affects the number of participants. Thus, we do not think that our counterfactual analyses will be adversely affected by omitting the value assessment from the covariates of the Poisson parameter.

If we were to allow the value assessment to determine the distribution of the number of potential bidders, it could interact with the estimation of the valuation distribution, potentially allowing changes in one of the distributions to be offset by changes in the other. For example, a high realized second-order statistic could both be the product of a large number of potential participants or high valuations on a small number of bidders. If a high

value assessment is associated with both high valuations and a large numbers of potential participants, it would not be clear through which mechanism the value assessment captures the data generating process that generated the data we observe.

5.5 Defining the Log-Likelihood Function

Finally, to estimate the structural parameters $\boldsymbol{\theta}$ and $\boldsymbol{\lambda}$ by the means of Maximum Likelihood Estimation (MLE), we will define the Log-Likelihood (LL) function. Since our estimation requires knowledge of the second highest bid, we consider only auctions where there are a minimum of two active participants when calculating the PDF of the truncated second-order statistic. We set the truncated PDF equal to 1 for $n_t \in [0, 1]$ because the transaction price can not be interpreted as the second-order statistic in these cases. In the LL function the log density nets out to zero for those two instances.

Incorporating the definitions above together with Equation 5.3, 5.4 and 5.5, the likelihood function for a particular auction t is defined as

$$L_t(\boldsymbol{\theta}, \boldsymbol{\lambda} | y_t, n_t) = \begin{cases} \prod_{N_t=n_t}^{\infty} \{P(n_t | N_t; \boldsymbol{\theta}) \cdot P(N_t; \boldsymbol{\lambda})\}, & n_t \in [0, 1] \\ f_y^{2:n_t}(y_t | n_t; \boldsymbol{\theta}) \cdot \prod_{N_t=n_t}^{\infty} \{P(n_t | N_t; \boldsymbol{\theta}) \cdot P(N_t; \boldsymbol{\lambda})\}, & n_t \in [2, +\infty) \end{cases} \quad (5.6)$$

Taking the log of Equation 5.6 and summing across all auctions t yields

$$LL(\boldsymbol{\theta}, \boldsymbol{\lambda}) = \sum_t \ln \{L_t(\boldsymbol{\theta}, \boldsymbol{\lambda} | y_t, n_t)\} \quad (5.7)$$

The estimates of $\boldsymbol{\theta}$ and $\boldsymbol{\lambda}$ are obtained by maximizing Equation 5.7 with respect to said parameters.¹⁶ We use a MLE method from the package *bbmle* developed by Bolker (2017). The package uses the method of limited-memory modification of the quasi-Newton method, also known as a variable metric algorithm. The method was simultaneously developed by Fletcher (1970), Goldfarb (1970) and Shanno (1970). Simply put, the technique uses function values and gradients to build a picture of the "surface" to be optimized (R Core Team).

¹⁶R code for the model can be provided on request.

6 Analysis

In this chapter, we are going to present and interpret our estimated parameters and provide graphical illustrations of the resulting probability distributions. Subsequently, we investigate how the total expected utility of the selling parties is affected by changes in the reserve price and bidder participation through counterfactual analysis. Throughout the chapter, we will refer to expected revenue and expected utility interchangeably. Keep in mind that the calculated expected revenue in this chapter amounts to the total amount accrued to Vinmonopolet, Blomqvist and the seller.

6.1 Obtaining the Results

6.1.1 Parameter Estimation

We estimated the parametric model based on the data set described in chapter 3 with 2781 observations.¹⁷ Relevant for the estimation, 9.7% (270) of the observations had zero or one participant, and one round is not represented in the sample (date 2017-03).¹⁸ The estimates of the coefficients for the shape, scale and lambda parameters are displayed in Table 6.1. Note that Vinmonopolet's value assessment has been divided by 1000 and the coefficients should be interpreted in thousands.

The shape parameter β displayed in column 1 of Table 6.1 (also known as the Weibull slope) is above one for all specifications of the covariates, implying that the Weibull distribution starts in origo and has a global maximum. An increase in the value assessment of NOK 1000 is associated with a change in the shape parameter of 0.02 for Bordeaux, 0.01 for Burgundy and -0.13 for Piedmont. This implies that the dispersion and skewness of the valuations of Bordeaux and Burgundy wines diminishes as the value assessment increases, while the opposite is true for Piedmont.¹⁹ The intercepts of Bordeaux, Burgundy and

¹⁷See Appendix 2 for performance tests of the model.

¹⁸Some of the transactions in the four last auction rounds in our sample might have incurred a NOK 150 shipping fee due to the increased number of outlets, but we do not have data on which transactions – if any – included shipping.

¹⁹Specifically, the ratio of the standard deviation to the mean – the coefficient of variation $\frac{\sigma}{\mu}$ – decreases as the value assessment increases for Bordeaux and Burgundy.

Table 6.1: Estimates of the Shape (β), Scale (η) and Poisson (λ) Parameters.

	<i>Parameter:</i>		
	Shape	Scale	Poisson
	β	η	λ
	(1)	(2)	(3)
Value Assessment	0.02** (0.01)	747.77*** (0.07)	
Burgundy	-0.50*** (0.04)	-322.09*** (5.16)	3.65*** (0.23)
Piedmont	0.39*** (0.09)	-145.96*** (0.27)	-1.58*** (0.20)
Value Assessment $\times$ Burgundy	-0.01 (0.01)	30.70*** (0.20)	
Value Assessment $\times$ Piedmont	-0.15*** (0.03)	303.97*** (0.03)	
Round Indicator			0.02 (0.01)
Constant	1.62*** (0.03)	273.28*** (5.87)	6.70*** (0.18)
Observations	2781	2781	2781

¹ *p<0.05; **p<0.01; ***p<0.001.

² Standard errors in parenthesis.

³ Numbers are rounded to two decimal places.

Piedmont are 1.62, 1.12 and 2.01 respectively. We can see that the valuations of Bordeaux wine are relatively less dispersed and skewed compared to Burgundy wine for any value assessment since the shape parameter of Bordeaux is strictly larger than for Burgundy. The slopes and constants are significantly different from zero at the 1% level for all the regions, but the difference in the slopes between Bordeaux and Burgundy is only significant at the 10% level (p-value = 0.09).

As expected, we see from column 2 of Table 6.1 that a higher value assessment is associated with a higher scale parameter for all regions; an increase of NOK 1000 in Vinmonopolet's value assessment is associated with an increase in the scale parameter of 747.77, 778.47 and

1051.74 for Bordeaux, Burgundy and Piedmont respectively. The intercepts of Bordeaux, Burgundy, and Piedmont are 273.28, -48.81 and 127.32 respectively. This is consistent with the value assessment capturing some quality component of the wines, as an increase in the value assessment is associated with an increase in the bidders' willingness to pay for the wine. The slopes, constants, and their differences are significantly different from zero at the 0.1% level for all the regions.

The Poisson parameter λ is displayed in column 3 of Table 6.1, and there seem to be substantial differences in participation between the regions, with the average number of potential participants being 7.16, 10.81 and 5.58 in the last round for Bordeaux, Burgundy, and Piedmont respectively. The constants and their differences are significantly different from zero at the 0.1% level for all the regions. The coefficient of the round indicator is only statistically different from zero at the 10% level ($p\text{-value} = 0.07$), so the number of potential participants per auction seems to have remained stable across auction rounds in the sample period. However, when we account for the total number of auctions being held each round, and the number of rounds being held each year, the yearly participation in the wine auctions seems to have increased over time.




Figure 6.1: Expected Potential Participation Across Time.

Figure 6.1 depicts how total participation has evolved over auction rounds (B) and years (A). Keep in mind that the number of distribution outlets was increased from one (Oslo) to seven (Oslo, Bergen, Sandefjord, Stavanger, Trondheim, Ålesund, and Hamar) in April 2018. Looking at B in Figure 6.1, there seems to be an overall increase in the per round participation around the time in 2018 when it became an option to retrieve the wine at outlets outside of Oslo, which at least is consistent with an increase in the number of outlets allowing more individuals to partake in the auctions.

6.1.2 Interpreting Magnitudes and Illustrating the Distributions

To get some sense of magnitude on the parameters in Table 6.1, we have illustrated how the Weibull distribution changes when either the scale (η) or shape parameter (β) increases in Figure 6.2. The change in the shape parameter β is equivalent to increasing the value

assessment by NOK 10000, and the change in the scale parameter η is equivalent to increasing the value assessment by only NOK 1000, both for a Bordeaux wine (evaluated at $V_v = 1000$).




Figure 6.2: Effects of Changing the Scale (η) and Shape (β) Parameters of the Weibull Distribution.

A higher shape parameter (β) will shift the mass of the Weibull distribution to the right and make the distribution relatively less dispersed (the coefficient of variation decreases) and skewed, while the right tail becomes shorter and thinner, as Figure 6.2 depicts in A. The mean and standard deviation changes from ($\mu = 914$, $\sigma = 572$) with $\beta = 1.64$ to ($\mu = 907$, $\sigma = 511$) with $\beta = 1.84$, holding the scale constant at $\eta = 1021$. Even though the mean and standard deviation decreases in the shape parameter, the differences are small as $\beta = 1.62 \rightarrow \beta = 1.82$ holding everything else fixed.

A higher scale parameter (η) is associated with increasing the mean and standard deviation of the Weibull distribution, depicted in B for Figure 6.2. The mass of the Weibull distribution shifts to the right and becomes more balanced. The mean and corresponding standard deviation increases from ($\mu = 914$, $\sigma = 572$) with $\eta = 1021$ to ($\mu = 1583$, $\sigma = 990$) with $\eta = 1769$, holding the shape parameter fixed $\beta = 1.64$. The differences are

large, implying that the main channel through which changes in the value assessment affects the mean and standard deviation of the valuations is through the scale parameter.

We have also illustrated the valuation distribution for two cases for each region. The graphs are divided into two buckets according to the regions' range of V_v . Piedmont has observations of $V_v \in [200, 10000]$, whereas Bordeaux and Burgundy have observations of $V_v \in [100, 80000]$ and $V_v \in [200, 80000]$ respectively. We choose to use $V_v = 1000$ and $V_v = 20000$ for Bordeaux and Burgundy, and $V_v = 500$ and $V_v = 5000$ for Piedmont. The range $V_v \in [1000, 20000]$ for Bordeaux and Burgundy, and $V_v \in [500, 5000]$ for Piedmont represents the majority of observations in our data set. Thus, using those ranges allow us to infer something about a majority of the wine auctions.²⁰




Figure 6.3: Estimated Weibull Densities for Bordeaux, Burgundy and Piedmont Wine.

As expected, increasing the value assessment for Bordeaux and Burgundy reduces the skewness and the dispersion of the distributions depicted in Figure 6.3. However, the relationship is reversed for Piedmont, which becomes more skewed and has more dispersion when the value assessment increases due to the lower shape parameter.

²⁰We will use the same cases when we estimate the optimal reserve prices in chapter 6.3.1.

Finally, Figure 6.4 displays the Poisson probabilities for all three regions. Notice how Burgundy has a much higher Poisson parameter relative to Bordeaux and Piedmont, which directly corresponds to a larger mean and variance for the potential participation distribution.




Figure 6.4: Estimated Poisson Probabilities for All Three Regions.

6.2 Vinmonopolet's Value Assessment

As discussed in chapter 2.2.1, the value assessment that Vinmonopolet performs on the wines considers various quality characteristics such as vineyard, brand, bottle size, vintage, storage and bottle condition that are not fully observable by the bidders. Our estimates confirm that higher value assessments are associated with higher buyer valuations, which is consistent with the value assessment capturing the quality component of the bidders' valuations discussed in chapter 4. However, we also know that Vinmonopolet draws on other sources of information such as own historical prices and international prices of similar wines when forming their valuations. Hence, there could be an issue of reverse causality in the sense that past auction outcomes driven by the bidders' preferences inform

Vinmonopolet's value assessment. Thus, we are not able to isolate the causal effect of quality on the bidders' willingness to pay. Nonetheless, the value assessment serves as a useful tool in predicting the parameters of the valuation distribution, and it allows policymakers to make policy decisions on a range of different wines conditioned on a single aggregating measure instead of a large set of covariates.

6.3 Counterfactual Analysis

Now that we have estimated the structural parameters of the underlying valuation distribution and potential participation distribution, we can proceed to calculate the optimal reserve prices, discuss how Vinmonopolet's pricing policy fares against the optimal reserve price, and conduct counterfactual analyses on the expected revenue. All of the calculations in this subsection are estimated by simulating 100 000 auctions for each individual case using the parameters derived in chapter 6.1.1. Note how the potential participants are treated differently in the counterfactual analysis relative to chapter 4.3 and 4.4. In our model, N is a stochastic variable, and simulating the potential participants instead of assuming a known N allows us to account for this feature.²¹

6.3.1 Calculating the Optimal Reserve Price

Using the estimated parameters of our model, we can calculate the optimal reserve prices from the perspective of a revenue-maximizing seller. Relevant for the discussion below, notice the role of v_0 in Equation 4.1. The higher utility the seller enjoys from consuming the wine him or herself, the higher the optimal reserve price. The seller does not want to accept an offer from a bidder who has a lower valuation than the utility the seller receives from consuming it him or herself. Hence, the reserve price is increasing in the agent's utility of retaining the wine.

For the sake of exposition, we are going to calculate the reserve prices for the aforementioned cases and let the utility of retaining the wine either be zero or 80% of Vinmonopolet's

²¹The assumption of a known N in chapter 4 was done only to develop an intuition for the expected revenue and the reserve price.

valuation, formally $v_0 = 0$ and $v_0 = 0.8 \cdot V_v$.²² The scenario with $v_0 = 0$ could represent a scenario where the seller inherited a wine cellar, but does not drink wine him or herself. The scenario with $v_0 = 0.8 \cdot V_v$ is included to represent those instances where the seller enjoys drinking the wine him or herself. This is also the case where Vinmonopolet's 80% pricing rule coincides with the optimal allocation rule of $r = v_0$.

	$V_v = 1000$		$V_v = 20000$	
	Bordeaux	Burgundy	Bordeaux	Burgundy
Optimal Reserve Price, $v_0 = 0$	755.2	654.9	10752.2	12576.6
Vinmonopolet's Reserve Price	800.0	800.0	16000.0	16000.0
<i>Implied v_0 from the Pricing Policy</i> ¹	<i>72.2</i>	<i>162.0</i>	<i>8831.6</i>	<i>4355.9</i>
Optimal Reserve Price, $v_0 = 0.8 \cdot V_v$	1326.6	1393.6	21343.0	25972.0

Table 6.2: Optimal Reserve Prices for Bordeaux and Burgundy.

¹ The sellers' implied v_0 given revenue maximization.

Table 6.2 displays the optimal reserve prices for Bordeaux and Burgundy using the cases presented in chapter 6.1.2, along with Vinmonopolet's reserve price. It is evident that Vinmonopolet's pricing policy – setting the reserve price equal to 80% of V_v – will be too high in the case where $v_0 = 0$ and too low when $v_0 = 0.8 \cdot V_v$ for Bordeaux and Burgundy wine if the goal is to maximize total expected utility of the selling parties. Since we do not know how the sellers' valuations are distributed, it is hard to say which case is more likely. Indications from Vinmonopolet's wine experts point to several cases where a good proportion of the wine sellers are people that inherited a wine cellar, but it is unknown to us whether those sellers enjoy the wine themselves or not. Setting a reserve price which lies above the revenue-maximizing reserve price is strictly worse for both the selling parties and the buyers. Thus, if Vinmonopolet believes v_0 is low ($v_0 \rightarrow 0$), there could be potential gains in allocative efficiency in setting a lower reserve price relative to the current pricing policy. The current pricing policy Vinmonopolet employs suggest that the implied v_0 for Bordeaux and Burgundy are ($v_0 = 72$, $v_0 = 162$) for $V_v = 1000$, and ($v_0 = 8832$, $v_0 = 4356$) for $V_v = 20000$.

Table 6.3 depicts Vinmonopolet's and the optimal reserve prices for Piedmont. Interestingly, the table indicates that the optimal reserve price for $v_0 = 0$ is higher than Vinmonopolet's

²²By construction, optimal reserve prices with $v_0 > 0$ are strictly greater than optimal reserve prices with $v_0 = 0$.

price, which means if Vinmonopolet's pricing policy were to correspond with the revenue-maximizing reserve price, this would imply a negative v_0 . No reserve prices for a Piedmont wine below 464 for $V_v = 500$ or 4296 for $V_v = 5000$ would be optimal if the seller has any utility in retaining the wine, i.e. $v_0 > 0$. From Table 6.3, the current pricing policy Vinmonopolet employs suggests that the implied v_0 for Piedmont is $v_0 = -134$ for $V_v = 500$, and $v_0 = -408$ for $V_v = 5000$. This gives strong indications that Vinmonopolet is not setting the total expected utility maximizing reserve prices since it is implausible that anyone would pay any substantial amount to have the wine removed.

	$V_v = 500$	$V_v = 5000$
Vinmonopolet's Reserve Price	400.00	4000.00
<i>Implied v_0 from the Pricing Policy</i> ¹	<i>-133.8</i>	<i>-408.0</i>
Optimal Reserve Price, $v_0 = 0$	464.0	4296.1
Optimal Reserve Price, $v_0 = 0.8 \cdot V_v$	710.3	7513.1

Table 6.3: Optimal Reserve Prices for Piedmont.

¹ The sellers' implied v_0 given revenue maximization.

The static rule that Vinmonopolet employs – which does not consider the seller's actual valuation – runs the risk of setting inefficient reserve prices. This is the case regardless of the goal being revenue maximization or allocative efficiency since both depend on knowledge about v_0 . This stands in contrast to how Vinmonopolet's third party – Blomqvist – operates with auctions of arts and other expensive items. The general Blomqvist model is to allow the seller to set the reserve price in cooperation with them, making it possible to set the reserve price conditional on v_0 .

6.3.2 The Impact of the Reserve Price on Expected Revenue

We have established that Vinmonopolet's pricing rule does not coincide with either the total revenue or welfare maximizing reserve prices, but the implications of these findings hinge on to which degree deviating from said optimal prices yield different outcomes. In this section, we will investigate how the optimal reserve price compares to alternative reserve prices such as Vinmonopolet's 80% rule in terms of expected revenue and its variance. We will use two specific cases to illustrate our findings in this section: a Bordeaux wine with a valuation of NOK 20000 and a Bordeaux wine valued at NOK

50000. Bordeaux is chosen due to its representativeness in the data sample. We use a valuation of NOK 20000 because most of the wines in our data set have a valuation below NOK 20000, and the main insights in the discussion below for $V_v = 20000$ hold for wines with lower value assessments across all three regions.²³ ²⁴ We include a valuation of NOK 50000 to study the effects of deviating from the optimal reserve price for wine bottles with very high value assessments, even though most wines are substantially less valued.²⁵




Figure 6.5: Expected Utility (EU) and the Standard Deviation of EU as a Function of the Reserve Price for a Bordeaux Wine Auction with $V_v = 20000$.

Figure 6.5 graphs how expected revenue and its uncertainty (represented by the standard deviation) varies with the reserve price for a Bordeaux wine with $V_v = 20000$. It indicates that for sellers with a low utility of retaining the object ($v_0 \rightarrow 0$) there is a small difference in expected revenue between setting the reserve price equal to zero, the optimal reserve price and Vinmonopolet's reserve price. More specifically, the difference between the optimal reserve price and a zero reserve price with $v_0 = 0$ is NOK 137. The difference between the optimal reserve price and Vinmonopolet's reserve price in terms of expected

²³See Table 3.2 for how observations are distributed according to V_v .

²⁴Appendix 4 contains similar analyses for Bordeaux, Burgundy and Piedmont for wine bottles with $V_v = 5000$.

²⁵Appendix 4 contains a similar analysis for Burgundy wine with $V_v = 50000$.

revenue is 544. For sellers with a high utility of retaining the object ($v_0 \rightarrow 16000$) a more distinct global maximum emerges from Figure 6.5. The difference between using the optimal reserve price and a zero reserve price for $v_0 = 16000$ is NOK 1943, and the difference between the optimal reserve price and Vinmonopolet's reserve price is NOK 826.

The graph depicting the standard deviation as a function of the reserve price in Figure 6.5, displays a reserve price with a potentially large effect on the uncertainty of the revenue for $v_0 = 0$ with standard deviations of 5013, 5104 and 6745 for the optimal reserve price, zero reserve price and Vinmonopolet's reserve price respectively. At really large reserve prices the standard deviation grows large – up to 12239 – which can be explained by a small probability of selling, but at a really high price in the event of a sale. The difference in standard deviation across reserve prices is much smaller for $v_0 = 16000$ with standard deviations of 3482, 4823 and 3462 for the optimal reserve price, zero reserve price and Vinmonopolet's reserve price respectively. This is intuitive since all the relatively extreme outcomes of no sale and zero utility is now replaced by the positive utility of retaining the wine, which is a relatively less extreme outcome compared to the sale outcomes.




Figure 6.6: Expected Utility (EU) and the Standard Deviation of EU as a Function of the Reserve Price for a Bordeaux Wine Auction with $V_v = 50000$.

Figure 6.6 depicts a similar graph as Figure 6.5, albeit for a Bordeaux wine with a valuation of NOK 50000. The aforementioned qualitative points hold for the more expensive Bordeaux wine as well. The difference in expected revenue between the optimal reserve price and a zero reserve price with $v_0 = 0$ is NOK 259, whereas the difference is NOK 2152 between Vinmonopolet's and the optimal reserve price. The standard deviation is 9891, 9367 and 15756 for a zero, optimal and Vinmonopolet's reserve price respectively. For $v_0 = 40000$ the difference between the optimal and no reserve price is NOK 3942, and between the optimal and Vinmonopolet's reserve price the difference is NOK 1535. The standard deviation is 9132, 5946 and 5962 for a zero, optimal and Vinmonopolet's reserve price respectively.

How does the reserve price affect the expected revenue of the wine auctions? Figures 6.5 and 6.6 paint a picture of the significance on expected revenue between setting an optimal reserve price relative to no reserve price: it depends on the level of v_0 . If the sellers utility of retaining the wine is low ($v_0 \rightarrow 0$), then there is a negligible difference in the expected revenue and its variance between a zero reserve price and the optimal reserve price for all valuations. However, if the sellers have high utility of retaining the wine ($v_0 \rightarrow 0.8 \cdot V_v$), the differences in expected utility and its variance are higher and increases with higher valuations. Increasing the valuation from NOK 20000 to NOK 50000 is associated with increasing the difference in expected utility between a zero reserve price and the optimal reserve price from NOK 1943 to NOK 3942.

How does Vinmonopolet's reserve price fare against the revenue maximizing reserve price? For low levels of v_0 , Vinmonopolet's pricing rule yields too high reserve prices, although not at a great cost for valuations below NOK 20000. Importantly, the difference in both expected revenue and its variance is decreasing in a lower valuation. Hence, for the majority of the wines auctioned by Vinmonopolet, the current pricing rule seems to fare reasonably well compared to the optimal reserve price if $v_0 \rightarrow 0$. However, for very high-valuated wines the performance of Vinmonopolet's pricing rule worsens, as is clear from 6.6. The difference in both expected revenue and its variance at $v_0 = 0$ and $V_v = 50000$ has become more distinct, and this difference keeps increasing if we increase V_v beyond NOK 50000.

For high levels of v_0 , Vinmonopolet's pricing rule yields too low reserve prices relative to

the revenue maximizing price, but at a low cost in terms of both revenue and its variance. The difference in expected revenue is increasing in the valuation, albeit at a decreasing rate. Furthermore, for almost all wines in the data set the difference is significantly smaller than NOK 826 in terms of expected revenue.²⁶ Thus, our analyses indicate that the current pricing rule fares relatively well against the optimal reserve price for high levels of v_0 .

In summary, what are implications from the analysis above on Vinmonopolet's pricing policy? It depends on Vinmonopolet's beliefs about v_0 and the value assessment of the wine. For low value assessments – below NOK 20000 – there seems to be little to gain for the selling parties by changing the current pricing policy in terms of either increasing expected revenue or lowering the uncertainty, for all relevant values of v_0 .²⁷ Likewise, for high value assessments, if Vinmonopolet believes that the sellers have high levels of v_0 there is little to gain by changing the current pricing policy across the board, viewed from the selling parties' perspectives. However, for high value assessments with low levels of v_0 , the inefficiencies incurred by Vinmonopolet's pricing rule are more substantial. Expected revenue and its variance could be significantly improved by adjusting the reserve price downwards.

Based on this, our policy recommendations regarding the pricing policy is the following: First, Vinmonopolet's current pricing policy seems to work relatively well for a majority of the auctions, viewed from the selling parties' perspectives. This is a robust finding that holds for all relevant levels of v_0 . Furthermore, for very-high valuated wines, Vinmonopolet should include the seller when deciding upon the reserve price. By including the seller, Vinmonopolet can form beliefs on the seller's v_0 , which could help Vinmonopolet set more efficient reserve prices. If the seller's utility of retaining the wine is low ($v_0 \rightarrow 0$), Vinmonopolet should set a lower reserve price relative to what their current policy implies. However, if the seller has a high utility of retaining the wine ($v_0 \rightarrow 0.8 \cdot V_v$), Vinmonopolet should set a higher reserve price relative to what their current policy implies.

Our above analysis has established that the selling parties' interests in many instances are maintained with Vinmonopolet's current pricing regime, but how about welfare in general? The consumers' payoffs are not explicitly modeled in this thesis, but we know

²⁶See Appendix 4 for simulations for Bordeaux, Burgundy and Piedmont wines valuated at $V_v = 5000$.

²⁷If the selling parties are risk averse, they would appreciate a reduction in the standard deviation.

that the optimal allocating rule is $r = v_0$. With the current pricing rule, this does not generally hold, and in principle Vinmonopolet runs the risk of mis-allocating in one of two ways: Setting the reserve price too high so the wine remains unsold even though a beneficial trade could have occurred with a lower reserve price, or setting the reserve price too low so the seller is forced to give up the wine to someone that values it less than him or her. The former scenario does not seem to be the case since it should manifest itself in a low success rate, which for the wine auctions is 97.9% in the sample period. The latter one will not be the case unless the seller has a really high valuation of the wine – specifically $v_0 > 0.8 \cdot V_v$ – and our analysis shows that the inefficiencies in expected revenue are relatively low in these cases.²⁸ Furthermore, if v_0 in fact is really high and there is some probability of mistakenly selling the wine to a bidder with a lower valuation than the seller, this could be mitigated, as we will discuss in the next section.

6.3.3 The Impact of Bidder Participation on Expected Revenue

In light of what the counterfactual analyses of the effect of changes in the reserve price on expected revenue revealed – that the difference between an optimal and zero reserve price was negligible in cases with $v_0 \rightarrow 0$ and significant for $v_0 \rightarrow 0.8 \cdot V_v$ – we wanted to identify the underlying driver of expected revenue for all auctions. It turns out that one of the primitives we previously held constant – namely the expected number of potential participants – plays a pivotal role in determining the expected revenue of an auction. Figure 6.7 graphs the expected utility as a function of the expected number of potential participants, represented by the Poisson parameter. The expected utility is evaluated at both the optimal reserve price and Vinmonopolet’s reserve price, and the figure includes all the previously analyzed combinations of $v_0 = 0$ and $v_0 = 0.8 \cdot V_v$, and $V_v = 20000$ and $V_v = 50000$ for a Bordeaux wine.²⁹

Figure 6.7 displays a clear relationship between the expected revenue and the expected number of potential participants λ . A higher λ is associated with a higher expected utility, albeit at a decreasing rate. In an auction of a Bordeaux wine with $V_v = 20000$ and $v_0 = 0$,

²⁸Presumably, the seller would not like to auction the wine if the reserve price is too misaligned with his interest, limiting the scope for having the seller give up the wine at extremely low transaction prices. However, we could forgo a mutually beneficial trade in this instance.

²⁹The results are similar for Burgundy and Piedmont.




Figure 6.7: Expected Utility as a Function of the Poisson Parameter (λ) for a Bordeaux Wine with $V_v = 20000$ and $V_v = 50000$ in Round 23.

Figure 6.7 indicates that increasing λ from 2 to 5 is associated with increasing the expected revenue by NOK 6557 and NOK 6822 for the optimal and Vinmonopolet's reserve price respectively. In the same auction, increasing λ from 10 to 13 is associated with increasing the revenue by NOK 1432 and NOK 1544 for the optimal and Vinmonopolet's reserve price respectively. For a Bordeaux wine with $V_v = 20000$ and $v_0 = 16000$, increasing λ from 2 to 5 is associated with increasing the expected utility by NOK 1708 and NOK 1581 for the optimal and Vinmonopolet's reserve price respectively. In the same auction, increasing λ from 10 to 13 is associated with increasing the revenue by NOK 1010 and NOK 1216 for the optimal and Vinmonopolet's reserve price respectively.

The discrepancy between the optimal and Vinmonopolet's reserve price is decreasing in the Poisson parameter, though at a faster rate for low levels of v_0 . This could be explained by the probability of the reserve price affecting the auction outcome being higher when it is high due to a high v_0 , and you need more potential bidders to make the probability negligible. Furthermore – for small values of λ in particular – the marginal value of attracting another bidder is less for a seller with a high v_0 since the no sale outcome

is relatively more likely when the reserve price is high due to a high v_0 , reducing the expected impact of increasing the probability of finding a bidder with an even higher willingness to pay.

Our analysis shows that bidder participation plays a pivotal role in determining the revenue raised in an auction, in contrast to some scenarios – e.g. when the seller has low levels of v_0 – where the choice between the optimal reserve price and a zero reserve price had a negligible effect on expected revenue. Our finding is consistent with the theoretical prediction discussed in chapter 4.3, that a seller who runs an English auction with no reserve price with $N + 1$ symmetric bidders will earn more in expectation than a seller who can hold an auction with an optimal reserve price with N buyers.

So what is the intuition for why bidder participation plays such an important role, and the reserve price plays a comparatively smaller role? Bidder participation is important because more bidders increase the probability of observing a large second-order statistic, resulting in a larger selling price. Our findings are consistent with Bulow and Klemperer (1996) that concludes that the value of exercising your market power through the reserve price is small relative to the value of attracting additional competition in general. However, we argue that the reserve price plays a particularly small role in the case of the wine auctions in the following paragraph.

The reserve price is more effective in removing bad outcomes for valuation distributions with a high mass on the left-hand side, i.e. low shape β and scale η parameters.³⁰ Consider the trade-off the auctioneer faces when setting the reserve price. The auctioneer is better off by removing some of the bad outcomes, but incurs a cost by lowering the probability of selling. A general intuition on the interplay between the distribution and the reserve price emerges: by increasing the reserve price, the seller eliminates bad outcomes, and the impact of this on expected revenue is determined by how probable those bad outcomes are. If these outcomes have a high probability, replacing them with a higher selling price has more impact on expected revenue. Thus, for distributions with a higher probability of bad outcomes (distributions with lower shape β and scale η parameters), the reserve price becomes a more effective tool in increasing expected revenue. In our case, the shape parameter is relatively high across our covariates which further explains why the effect of

³⁰We think of bad outcomes in terms of the second-order statistic being low, resulting in lower revenue.

the reserve price on expected revenue is negligible when $v_0 \rightarrow 0$.

We have already discussed how Vinmonopolet's current pricing policy seems to work relatively good for the selling parties in a majority of the auctions, and that it is likely that the pricing policy does not have adverse effects on welfare. Nonetheless, For the scenarios where Vinmonopolet's pricing policy seems misaligned – for very-high valuated wines from a revenue maximizing perspective or for very-high levels of v_0 in terms of welfare – increasing bidder participation will alleviate the inefficiencies in those instances. In general, a higher expected number of potential participants reduces the probability of the reserve price influencing the auction outcome and improves allocative efficiency because the probability of the wine being allocated to the agent with the highest willingness to pay increases. Thus, increasing bidder participation could be a sensible effort for Blomqvist, the seller and Vinmonopolet alike.

As noted in chapter 4.3, the main ways available to increase bidder participation are through reductions in the participation costs and increasing awareness of the wine Auctions. Vinmonopolet has already taken measures to reduce the participation cost with increasing the number of outlets to retrieve the wines. Furthermore, any measures that make the participation process more efficient and intuitive, or provide information that makes decision making easier, could increase participation.

Increasing awareness of the wine auctions is a delicate matter in the sense that Vinmonopolet might not want to promote increased alcohol consumption in general, but rather encourage trade through regulated channels such as the wine auctions specifically. This excludes any promotional features such as TV advertisements or sponsorships that directly encourages alcohol consumption, but any effort to shift wine trades to the wine auctions would be in line with the interests of all the involved parties. This could be achieved by promoting Vinmonopolet's wine auctions through channels where the consumers have already sought out alcoholic beverages such as Vinmonopolet's retail stores or information diffusion through Vinmonopolet's current platforms such as their podcast, newsletter and wine magazine.

7 Reflections and Concluding Remarks

In this final chapter, we are going to provide concluding remarks to the findings, discuss the validity of our modeling assumptions and propose potential extensions.

7.1 Conclusion

In this master’s thesis, we have performed a structural estimation of the underlying valuation and potential participation distributions for Vinmonopolet’s electronic wine auctions, using a simple model of auction participant behavior and data on the observed number of participants, winning prices and reserve prices in said auctions. We conditioned the distributions on the region of origin, time and Vinmonopolet’s value assessment. This allowed us to estimate the scale and shape parameters of the latent valuation distributions of the three wine regions Bordeaux, Burgundy, and Piedmont as a function of Vinmonopolet’s value assessment, as well as the differences in the expected number of potential participants between regions and across time. We find that the coefficients on Vinmonopolet’s value assessment is consistent with the value assessment capturing some quality component of the wines that influences the bidders’ willingness to pay. However, we are not able to isolate the causal effect of this quality component on the bidders’ valuations due to issues of reverse causality.

By calculating the optimal reserve price for a revenue-maximizing seller as a function of the underlying valuation distribution, we are able to compare the optimal reserve price with Vinmonopolet’s reserve price along the dimensions of region, value assessment and the sellers own valuation. We show that it is implausible that Vinmonopolet determines the reserve price on either criterion of revenue maximization or optimal allocation since the reserve price is not conditioned on the sellers’ valuation.

However, Vinmonopolet’s current pricing policy – setting the reserve price equal to 80% of their value assessment – fares relatively well against the optimal reserve price in the majority of the auctions in terms of expected revenue. This is due to reserve prices having a small effect on the wine auctions in general. Additionally, in most cases no excessive uncertainty is incurred. Furthermore, concerning the welfare loss induced by

Vinmonopolet's deviations from the optimal allocation rule, there is reason to believe that the consumers are not significantly harmed due to the high success rate in our sample period of 98%. Hence, Vinmonopolet's current pricing rule generally aligns with both the selling parties' and the buyers' interests, but Vinmonopolet should include the seller in the pricing process for very highly valuated wines.

Finally, we show that bidder participation is the most important driver of expected revenue in the wine auctions. Furthermore, the likelihood of incurring inefficient outcomes due to deviations from the optimal reserve price is mitigated by increasing bidder participation. The feasibility of increasing bidder participation is naturally limited by the costs of attracting more bidders, but measures already in place such as Vinmonopolet's podcast, newsletter, wine magazine and general media coverage could contribute to improved outcomes for all the parties involved in the wine auctions.

7.2 Reflections On The Model

Throughout this master's thesis, we have made implicit and explicit assumptions regarding the specification of our model, that is: the effect of bid increments on equilibrium bidding, the bidders' arrival process and the one-shot, static game assumption. In the following section, we will reflect upon the validity and consequences of these assumptions.

7.2.1 Bid Increments

We chose to abstract away from the bid increments in the auction environment, leaving us with a simple second-price bidding strategy. As we argued, this has been common practice in the empirical literature. A natural extension of this assumption would be to incorporate the increments in the equilibrium bidding strategy. Interestingly, Hickman (2010) studied the implications of incorporating the bid increments, and he found that bidders will shade their bids in equilibrium, arguing that an online auction with bid increments is essentially a hybrid version of a first-price and second-price auction.

In a follow-up paper, Hickman et al. (2017) studied the consequences of using a simple second-price bidding strategy relative to the hybrid version (developed in Hickman (2010))

in electronic auctions. They used a nonparametric approach with a data set containing homogeneous laptops auctioned on eBay with no binding reserve prices. Hickman et al. (2017) demonstrated through simulations that ignoring the bid increments introduced bias in the estimate of the latent valuation distribution. The bias lead to an under-prediction of the expected revenue. In our study, the average value assessment for the largest region is NOK 3400, which corresponds to a bid increment of NOK 100. Thus, the bid increments will in most cases only constitute approximately 3% of the value assessment. Due to the fundamental differences between our study and the study conducted by Hickman et al. (2017) regarding data and modeling assumptions, it is challenging to determine the magnitude of the bias in our case.

7.2.2 Bidders' Arrival Process

We implicitly assumed that the sequence of bidders' arrivals are in such a way that everyone that has a valuation above the reserve price can place their bid. By introducing this assumption, our model runs the risk of underestimating the expected number of potential participants, since it is assumed that the only difference between the observed bidders and the potential bidders is caused by the reserve price. For example, one effect of underestimating bidder participation is overestimating the bidders' valuations, as an observation of a second-order statistic is estimated to be more extreme than it really is. On the other hand, a more extreme second-order statistic implies high participation, which in turn could partially offset the the bias. Nonetheless – given the time-frame of this thesis – we view our modeling of bidder participation as a first-order approximation that works better than simply claiming that $N_t = n_t$.³¹ A natural extension of our model would be to model the arrival process with regards to participation timing and the current standing price at the time of arrival, as done by Canals-Cerdá and Percy (2013).

7.2.3 One Shot, Static Game Assumption

On the grounds of parsimony, we treated each auction as an independent static one-shot game and abstracted away from the potential effects of auctions being held repeatedly

³¹Harry J. Paarsch (1992) assumed that the number of potential bidders were equal to the number of observed bidders.

and simultaneously. It is common in the auction literature to disregard budget constraints and issues of substitution between goods – implying that the demand for a specific bottle of wine is independent of other bottles up for auction – which stands in stark contrast to standard consumer theory (Gentry et al., 2018). For the static auction model with no budget constraints or substitution to be a good approximation of reality, we need a low degree of substitution between the different wines, and for the wines to constitute a small share of the bidders’ budgets.

We argue that for red wine, each bottle can be considered relatively unique due to heterogeneity in the bottle specific characteristics like geographical origin, brand, vintage, and condition – implying a low degree of substitution between bottles. Furthermore, in a survey from SSB (2013), it is shown that alcoholic beverages and tobacco constitutes only 2.7% of Norwegian households’ total consumption expenditure, which might imply that the income effect is small. Thus, we suspect that the impact of substitution and income effects is low in this study. A natural extension to prove this argument would be to incorporate the budget and substitution considerations in the bidders’ decision problem, a potential avenue for future research.

Widening the lens of focus, how likely are the assumptions in section 7.2.1, 7.2.2 and 7.2.3 to severely impact the conclusions drawn in this master’s thesis? We believe that most of the findings in this master’s thesis are robust to small biases in the estimates. Vinmonopolet’s pricing rule will deviate from the optimal reserve price regardless, which follows from not considering the seller’s utility of retaining the wine. Biases in the estimates could affect the magnitudes of the inefficiencies induced by deviations from the optimal reserve price, but increasing participation diminishes said inefficiencies nonetheless.

References

- Athey, S., Levin, J., Seira, E., et al. (2004). *Comparing open and sealed bid auctions: Theory and evidence from timber auctions*. Fondazione ENI Enrico Mattei.
- Bolker, B. (2017). *bbmle: Methods and functions for fitting maximum likelihood models in R*. R package.
- Bulow and Klemperer (1996). Auctions versus negotiations. *The American Economic Review*.
- Canals-Cerdá, J. J. and Percy, J. (2013). Arriving in time: Estimation of english auctions with a stochastic number of bidders. *Journal of Business & Economic Statistics*, 31(2):125–135.
- Canals-Cerdá, J. J. and Percy, J. (2013). Arriving in time: Estimation of english auctions with a stochastic number of bidders. *Journal of Business & Economic Statistics*, 31(2):125–135.
- Dasgupta and Maskin (2000). Efficient auctions. *The Quarterly Journal of Economics*.
- Fletcher, R. (1970). A new approach to variable metric algorithms. *The computer journal*, 13(3):317–322.
- Gentry, M. L., Hubbard, T. P., Nekipelov, D., and Paarsch, H. J. (2018). Structural econometrics of auctions: A review. *Foundations and Trends® in Econometrics*, 9(2-4):79–302.
- Goldfarb, D. (1970). A family of variable-metric methods derived by variational means. *Mathematics of computation*, 24(109):23–26.
- Harry J. Paarsch (1992). Empirical models of auctions and an application to british columbian timber sales. *Department of Economics Library*.
- Hickman, B. (2010). On the pricing rule in electronic auctions. *International Journal of Industrial Organization*, 28(5):423–433.
- Hickman, B. R., Hubbard, T. P., and Paarsch, H. J. (2017). Identification and estimation of a bidding model for electronic auctions. *Quantitative Economics*, 8(2):505–551.
- Holmes, T. J. and Sieg, H. (2015). Structural estimation in urban economics. 5:69–114.
- Lucking-Reiley, D. (2000). Auctions on the internet: What’s being auctioned, and how? *The journal of industrial economics*, 48(3):227–252.
- R Core Team, s. *stats: General-purpose Optimization*.
- Riley and Samuelson (1981). Optimal auctions. *American Economic Review*.
- Shanno, D. F. (1970). Conditioning of quasi-newton methods for function minimization. *Mathematics of computation*, 24(111):647–656.
- Song, U. (2004). Nonparametric estimation of an ebay auction model with an unknown number of bidders. *University of British Columbia*.
- SSB (2013). *Expenditure per household per year, by commodity and service group*.

- Susan Athey, Dominic Coey, Jonathan Levin (2011). Set-asides and subsidies in auctions.
- Thaler, R. H. (1988). Anomalies: The winner's curse. *The Journal of Economic Perspectives*, 2(1):191–202.
- Vinmonopolet's Podcast (2018). *Vin på auksjon*.

Appendix

A1 Testing the IPV Assumption

To test the independent private values assumption, one can estimate the correlation between the size of the winning bids and the number of active participants. It is a well known result in auction theory that if the auction environment is characterized by a Common Value (CV) component in the bidders' valuations, the size of the winning bids in the auctions will decrease as the number of participants increases, a result driven by the winner's curse.³² As the number of participants grows, the bidders will increasingly shade their bids due to the fear of overpaying for the good. We will perform a linear regression to investigate if there is a positive, negative or insignificant relation between the size of the winning bid of an auction and the number of active participants. If the relationship is negative or statistically insignificant, the environment is most likely characterized by common values. However, if the coefficient is positive, the environment is most likely characterized by independent private values.

The results in Table A1.1 show a positive relation between the winning bid's size and the number of active participants. An increase of one participant is associated with an increase of NOK 668 in the winning bid, all else equal. The coefficient is statistically significant at the 1% level. This indicates that we are operating in an IPV environment.

Table A1.1: Estimates for the CV *vs.* IPV Test.

Dependent Variable: Winning Bid	
	Coefficients
Intercept	2209.6 (358.8)
Number of Active Participants	668.3 (97.5)

¹ Standard errors in parenthesis.


² $n = 2781$.

We also performed a non-parametric fit between the number of active participants and the size of the winning bids, depicted in Figure A1.1. The figure displays a strictly positive

³²See Thaler (1988) for an exposition on the winner's curse.

relation between the winning bids and the number of participants.

Figure A1.1: Non-Parametrically Fitted Graph Between Winning Bids and Active Participants.



A2 Simulations

To illustrate the performance of our estimation method, we simulated artificial data of both a set of 30000 and 2000 auctions and used MLE to retrieve the parameters we used to generate the data. The number of potential bidders N_t is drawn from a Poisson distribution with $\lambda = 5$. The bidder valuations v_{it} are drawn from a Weibull distribution with $\beta = 1.5$ and $\eta = 2$.

Table A2.1 shows the estimated parameters when using 30000 simulations alongside the true parameter values that generated the data. We can see that the estimated parameter values closely resemble the true parameters, and that the standard errors are low.

Table A2.1: The Results from Testing Our Model with 30000 Observations.

Simulations			Estimates		
β	η	λ	$\hat{\beta}$	$\hat{\eta}$	$\hat{\lambda}$
1.5	2	5	1.51	2.00	5.09
			(0.0000)	(0.0000)	(0.0001)

¹ Standard errors in parenthesis.

Figure A2.1 compares the true and estimated Weibull and Poisson distributions for 30000 simulations.

Figure A2.1: Estimated Densities with 30 000 Observations.




Table A2.2 shows the estimated parameters when using 2000 simulations alongside the true parameter values that generated the data. Once again, we can see that the estimated

parameter values closely resemble the true parameters, and that the standard errors are low.


Table A2.2: The Results from Testing Our Model with 2000 Observations.

Simulations			Estimates		
β	η	λ	$\hat{\beta}$	$\hat{\eta}$	$\hat{\lambda}$
1.5	2	5	1.49	2.01	5.08
			(0.0007)	(0.0005)	(0.0014)

¹ Standard errors in parenthesis.

Figure A2.2 compares the true and estimated Weibull and Poisson distributions for 2000 simulations. By looking at the estimated parameters and the graphed distributions, it seems like our method works as intended.

Figure A2.2: Estimated Densities with 2000 Observations.



A3 Regressing the Number of Observed Bidders on Vinmonopolet's Value Assessment

The regression in Table A3.1 indicates that the slope coefficient of the number of observed bidders is not significantly different from zero for any conventional significance levels (p-value = 0.601). The two variables are plotted against each other in Figure A3.1. There seems to be no linear relation between the number of observed bidders and Vinmonopolet's value assessment.

Table A3.1: Regression of the Number of Observed Bidders on Vinmonopolet's Value Assessment.

	<i>Dependent Variable:</i>
	Vinmonopolet's Value Assessment
Number of Observed Bidders	40.583 (77.769)
Constant	3,566.978*** (286.204)
Observations	2,781

*p<0.05; **p<0.01; ***p<0.001.




Figure A3.1: Vinmonopolet's Value Assessment Regressed on the Number of Observed Bidders.

A4 Supplementary Graphs: Revenue Expectation and Variance as a Function of the Reserve Price

Figure A4.1, A4.2 and A4.3 graphs expected revenue and its variance as a function the reserve price for wine bottles valued at $V_v = 5000$ for Bordeaux, Burgundy and Piedmont respectively. Additionally, Figure A4.4 graphs expected revenue and its variance as a function the reserve price for a Bordeaux wine valued at $V_v = 50000$. The main insights from the analysis on a Bordeaux wine valued at NOK 20000 seem to hold for the majority of the wines in our sample (83% of the wines across the three main regions are valued at NOK 5000 or below). For a value assessment of NOK 50000, the deviations in terms of expected revenue induced by deviations from the optimal reserve price for Burgundy seem to be even smaller than for Bordeaux.




Figure A4.1: Expected Utility (EU) and the Standard Deviation of EU as a Function of the Reserve Price for a Bordeaux Wine with $V_v = 5000$.




Figure A4.2: Expected Utility (EU) and the Standard Deviation of EU as a Function of the Reserve Price for a Burgundy Wine with $V_v = 5000$.




Figure A4.3: Expected Utility (EU) and the Standard Deviation of EU as a Function of the Reserve Price for a Piedmont Wine with $V_v = 5000$.




Figure A4.4: Expected Utility (EU) and the Standard Deviation of EU as a Function of the Reserve Price for a Burgundy Wine with $V_v = 50000$.