

Design and Evaluation of Procurement Combinatorial Auctions

Sang Won Kim

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2014

©2014

Sang Won Kim

All Rights Reserved

ABSTRACT

Design and Evaluation of Procurement Combinatorial Auctions

Sang Won Kim

The main advantage of a procurement combinatorial auction (CA) is that it allows suppliers to express cost synergies through package bids. However, bidders can also strategically take advantage of this flexibility, by discounting package bids and “inflating” bid prices for single-items, even in the absence of cost synergies; the latter behavior can hurt the performance of the auction. It is an empirical question whether allowing package bids and running a CA improves performance in a given setting.

Analyzing the actual performance of a CA requires evaluating cost efficiency and the margins of the winning bidders, which is typically private and sensitive information of the bidders. Thus motivated, in Chapter 2 of this dissertation, we develop a structural estimation approach for large-scale first-price CAs to estimate the firms’ cost structure using the bid data. To overcome the computational difficulties arising from the large number of bids observed in large-scale CAs, we propose a novel simplified model of bidders’ behavior based on pricing package characteristics. Overall, this work develops the first practical tool to empirically evaluate the performance of large-scale first-price CAs commonly used in procurement settings.

In Chapter 3, we apply our method to the Chilean school meals auction, in which the government procures half a billion dollars’ worth of meal services every year and bidders submit thousands of package bids. Our estimates suggest that bidders’ cost synergies are economically significant in this application ($\sim 5\%$), and the current CA mechanism achieves high allocative efficiency ($\sim 98\%$) and reasonable margins for the bidders ($\sim 5\%$). We believe this is the first work in the literature that

empirically shows that a CA performs well in a real-world application.

We also conduct a counterfactual analysis to study the performance of the Vickrey-Clarke-Groves (VCG) mechanism in our empirical application. While it is well known in the literature that the VCG mechanism achieves allocative efficiency, its application in practice is at best rare due to several potential weaknesses such as prohibitively high procurement costs. Interestingly, contrary to the recent theoretical work, the results show that the VCG mechanism achieves reasonable procurement costs in our application. Motivated from this observation, Chapter 4 addresses such apparent paradox between the theory and our empirical application. Focusing on the high procurement cost issue, we study the impact of competition on the revenue performance of the VCG mechanism using an asymptotic analysis. We believe the findings in this chapter add useful insights for the practical usage of the VCG mechanism.

Table of Contents

1	Introduction	1
2	Structural Estimation Approach for Large-Scale First-Price Sealed-Bid Combinatorial Auctions	
	Auctions	12
2.1	Introduction	12
2.2	Structural Estimation Approach for Combinatorial Auctions	15
2.2.1	A Structural Model for First-Price Sealed-Bid Combinatorial Auctions	16
2.2.2	A Two-Step Structural Estimation Method	20
2.2.3	Identification	22
2.2.4	Limitations in Large-Scale CAs	27
2.3	The Characteristic-Based Markup Approach for Large-Scale CAs	28
2.4	Specifying Markup Restrictions	31
2.4.1	Group-Based and Size-Based Markup Models	32
2.4.2	A Refinement of the Size-Based Markup Model	35
2.4.3	Further Requirements on the Package-Characteristics Matrix	37
2.5	Estimation	41

2.5.1	Estimating the Distribution of Competitors' Bids	42
2.5.2	Markup and Cost Estimation	47
2.5.3	Improving the Computational Efficiency	49
2.6	Conclusion	53
2.7	Appendix for Chapter 2 - Proofs	55
2.7.1	Notation	55
2.7.2	Proofs of Main Results	56
2.7.3	Proofs of Side Lemmas	63
3	Application to the Chilean School-Meal Auction	68
3.1	Introduction	68
3.2	The Chilean Auction for School Meals	70
3.2.1	Brief History	70
3.2.2	Auction Process	71
3.2.3	Winner Determination Problem Formulation	73
3.2.4	Description of the Data	75
3.2.5	Discussion on the Assumptions of the Structural Model	77
3.3	Estimates of the Bid Distribution Parameters	80
3.4	Cost and Markup Estimation Results	82
3.4.1	Small Auctions Estimation	84
3.4.2	Results for Large-scale CAs	88
3.5	Efficiency and Counterfactual Analysis	91
3.5.1	Performance of the First-Price CA	92

3.5.2	Supplier Diversification	94
3.6	Conclusions	96
3.7	Appendix for Chapter 3	98
3.7.1	Results for the 2005 Auction	98
3.7.2	Package Density Measures	99
4	Vickrey-Clarke-Groves Mechanism	101
4.1	Introduction	101
4.2	VCG Mechanism and the Core	104
4.2.1	Rules of the VCG Mechanism	104
4.2.2	Connection to Coalitional Games	107
4.2.3	Example of a VCG outcome that is not in the Core	109
4.2.4	Core-selecting Auctions.	111
4.3	VCG in the Chilean School Meals Program	112
4.4	Analysis	114
4.4.1	Unit-wise analysis	117
4.4.2	Allocation-wise analysis	123
4.5	Conclusions	126
4.6	Appendix for Chapter 4	128
4.6.1	Finding Core Payments	128
4.6.2	Proofs	131
5	General Conclusions	138

List of Figures

2.1	Average scale discounts in per-meal bid prices.	13
3.1	Markup estimates from the full-dimension and extended size-based model.	86
3.2	Overall bid discount and cost synergy curves.	92

List of Tables

2.1	Distributional parameters and the cost estimation.	24
2.2	Distributional parameters and the optimal bids.	25
3.1	Summary statistics for the 2003 auction.	76
3.2	Fraction of winning firms that won only one package bid for different auctions. . .	78
3.3	The first step regression results for the 2003 auction.	81
3.4	Correlation structure among regions from the second step regression.	83
3.5	Estimated markup levels and the winning probability levels.	89
3.6	Average cost synergies.	90
4.1	Summary of bids and the VCG outcome.	110

Acknowledgments

First and foremost, I would like to express my deep gratitude to my advisers, Marcelo Olivares and Gabriel Weintraub, for their consistent support and invaluable guidance. Throughout the years at Columbia Business School, they were my inspiring mentors and sometimes caring big brothers. I feel very fortunate to have had such great opportunities to interact with them which had a profound impact on me as a researcher. I cherish the memory of countless meetings and endless discussions which were always enjoyable and rewarding. Chapters 2 and 3 are the fruitful outcomes of our collaboration. I am also thankful to Gabriel for his guidance in developing Chapter 4 as well. I can proudly say that I am academically Chilean with Marcelo and Gabriel as my advisors.

In addition to my advisers, I am also grateful to the members of my dissertation committee, Jay Sethuraman, Assaf Zeevi, and Brett Gordon. I am grateful to Jay Sethuraman for his consistent and kind guidance, and for being a strong supporter since my years in MSFE program at the IEOR department. I really enjoyed attending his courses, having various discussions, and also collaborating with him – Chapter 4 is the result of the collaboration with Jay and Gabriel. I am also indebted to Assaf Zeevi for his tireless support. I feel fortunate to have been exposed to his critical thinking and enthusiasm which became essential ingredients for my doctoral study. Finally, I am grateful to Brett Gordon for his valuable comments and suggestions.

Chapters 2 and 3 were supported by the Deming Doctoral Fellowship, and I would like to acknowledge the Deming Center's generous financial support. I am also thankful to Nelson Fraiman for his support and encouragement as the director of the Deming Center.

I would like to mention the inspiring scholars at Columbia University who made a great impact

to my academic growth during my study. In particular, I am thankful to Sergei Savin for his support in the early stage of my doctoral study, and to Awi Federgruen for his encouragement throughout my stay at Columbia. Omar Besbes and Alireza Tahbaz-Salehi were very generous to provide invaluable comments and help during my job market period, and I truly appreciate their support. In addition, I would like to thank many of the faculty members at the IEOR department of Columbia University who have helped me develop my interest in the field of operations research and strengthen academic rigor. Especially, I am grateful to Ward Whitt, Donald Goldfarb, and Soulaymane Kachani for their advices and guidance in my early years at Columbia. Thanks to my fellow classmates at Columbia who made the long journey enjoyable. I am thankful to all of them but especially to Bar Ifrach, Damla Gunes, Roger Lederman and Denis Saure.

I would like to express my deep gratitude to my parents, Keun Sik Choi and Jong Soo Kim, for their endless love and support. Finally, I send my deepest love and gratitude to the most important person in my life – my wife, Jiwan Choi. She supported me with unconditional love and devotion, patiently waited, and always stood by me as I went through this long and sometimes tough academic journey. Without her love, nothing would have been possible. This dissertation is dedicated to her.

To my wife, Jiwan.

Chapter 1

Introduction

In many important procurement settings suppliers face cost synergies: for example, transportation service providers can lower costs by coordinating multiple deliveries in the same route, and producers can lower average costs by spreading a fixed cost across several units. Motivated by these types of settings, auction mechanisms that allow bidders to submit package bids for multiple units so that they can express their synergies, have received much recent attention in practice and the academic literature. Indeed, these multi-unit auctions, typically referred to as *combinatorial auctions* (CAs), have been implemented in many procurement applications. For example, Elmaghraby and Keskinocak (2004), Sandholm (2006), and Hohner et al. (2003) describe applications at The Home Depot, Procter & Gamble, and Mars Inc., respectively. This type of auctions have also been implemented in non-procurement settings, most notably in the auctions for wireless spectrum run by the Federal Communications Commission (FCC) (McDuff, 2003).¹

A central auction design question in multi-unit settings is how allowing bidders to submit bids for packages of units impacts the performance of the mechanism. From the perspective of an auction designer, there are typically two measures that are relevant when evaluating performance: (1) efficiency, which compares the actual bidders' costs realized in the auction allocation relative to the minimum possible cost allocation that can be achieved; and (2) optimality, which relates to

¹Cramton et al. (2006) provides an overview on CAs.

the total payments to the bidders by the auctioneer. The above design question is crucial because allowing for package bidding via a CA can have countering effects on the performance under these two measures, as we describe next.

On one hand, allowing package bids can enhance the performance especially in the presence of cost synergies. In many procurement applications, such as the examples mentioned above, bidders may have cost synergies due to *economies of scale*, which depend on the volume allocated to a given supplier, and *economies of density*, which depend on the proximity of the units in an allocation. If bidders were allowed only to submit bids for each unit separately, they would face the risk of winning some units but not others. This phenomenon, known as the *exposure problem*, can make the bidders less aggressive in expressing the economies of scale and density that arise from supplying multiple units. Enabling package bidding through a CA eliminates this risk, potentially leading to more efficient outcomes and lower procurement costs.

However, allowing package bids could also hurt the performance. As pointed out by Cantillon and Pesendorfer (2006b) and Olivares et al. (2012), when using a first-price rule, bidders can engage in strategic bundling in which they submit package discounts even in the absence of cost synergies. One motivation to do so may be to leverage a relative cost advantage in a unit (for which the bidder is the cost-efficient provider) into another unit (for which the bidder is not the efficient provider). The firm may attempt to win both units by submitting a “discounted” package bid for the bundle and “inflating” both single-unit bids. If the bidder wins the package, it will lead to an inefficient allocation in which a unit is not served by the lowest-cost supplier. In addition, package bidding can also lead to the *threshold problem*, in which “local” suppliers bidding for small packages free-ride on each other to outbid “global” suppliers submitting bids on larger packages; this free-riding can lead to less competitive bidding, higher margins and thereby higher payments for the auctioneer. Milgrom (2000) and Baranov (2010) provide examples of the threshold problem.

Given the aforementioned trade-off, we expect that a CA should enhance the performance, relative to auction mechanisms that preclude package bidding, if cost synergies are strong and the incentives for the types of strategic behavior mentioned above are weak. However, analyzing the

actual performance of a CA requires evaluating cost efficiency and the margins of the winning bidders, which is typically private and sensitive information of the bidders. Moreover, existing theory is not conclusive on how large the incentives for strategizing are in a specific application. Thus motivated, in Chapters 2 and 3 of this dissertation we develop and apply an empirical approach to evaluate the performance of first-price CAs based on observed bid data, and use it to inform the auction design.

To measure the performance of a CA, it is essential to identify bidders' supplying costs, which are not directly observable in the bid data. In a CA, bidders may place *discounted* package bids, which may reflect cost synergies. However, the presence of package discounts is not conclusive about the performance of the auction. Bid discounts could also be driven by the types of strategic behavior alluded to above; bidders could inflate their single-unit bids relative to package bids to increase the probability of winning larger packages with relatively high margins, even in the absence of cost synergies. Such strategic inflation of single-unit bids also results in package discounts. In Chapter 2, we propose a structural estimation approach that identifies the bidders' costs using actual bid data, and therefore disentangles whether the discounts observed in the bid data are driven by cost synergies or strategic markup adjustments. This distinction is important as one would expect a CA to perform well only if the discounts observed are mostly explained by cost synergies.

Our method is based on the influential work of Guerre et al. (2000) for single-unit auctions that was later extended by Cantillon and Pesendorfer (2006b) and was applied to the London bus routes CAs with two or three units. More specifically, Cantillon and Pesendorfer (2006b) conduct the structural estimation of first-price CAs in two steps. In the first step, a statistical distribution of the competitors' bids is estimated from bidding data. In the second step, the first-order conditions from the bidder's profit maximization problem are used to find the imputed costs that would rationalize the bids observed in the data. These first-order conditions involve beliefs about the competitors' bidding behavior, and the distribution estimated in the first step is used to sample competitors' bids and form these beliefs. The estimated costs enable the calculation of the cost efficiency and bidders' margins in the CA to evaluate its performance. They also allow us to evaluate alternative auction

designs. We show, however, that this approach cannot be directly applied to large-scale CAs with many units, due to the high dimensionality of the bid vectors. This is an important limitation for many real-world procurement applications of CAs; for example, Caplice and Sheffi (2006) and Bichler et al. (2006) report CAs for transportation and procurement of inputs that typically involve hundreds of units.

Consequently, an important methodological contribution of our work is the development of a novel approach to apply structural estimation to large-scale first-price CAs. We introduce a “simplified” version of the bidder’s problem where the markups charged on package bids are chosen based on a reduced set of package characteristics. With this simplification, the first-order conditions of the bidder’s problem become computationally and econometrically tractable. We impose reasonable restrictions to the structure of the markups that reduce the complexity of the bidders problem but still provide sufficient flexibility to capture strategic behavior that can hurt the performance of a CA. In addition, we introduce a parsimonious, yet flexible parametric description of the distributions of competitors’ bids for CAs that involve heterogeneous units, and scale and density discounts. This specification makes the estimation of the distributions of competitors’ bids tractable. Overall, these two simplifications make the use of the structural approach feasible in large-scale CAs.

We expect that our approach, based on pricing package characteristics, can be used in other real-world large-scale auctions. Despite the practical use of CAs like the ones mentioned above, their econometric analysis has been limited due to the complexity. Even though our method may need modification to accommodate different pricing and auction rules in those different settings, we believe that our idea of using characteristic-based pricing can be a useful starting point to reduce the complexity of econometric analysis in such large-scale settings.

In Chapter 3, we apply our method to the Chilean school meals CA in which the government procures half a billion dollars’ worth of meal services every year to feed 2.5 million children daily. This is one of the largest and most important social programs run by the Chilean government. The application fits well within the class of large-scale CAs: each auction has about 30 units and firms submit hundreds of bids (see Epstein et al. (2002) for more details on the auction). This application

serves as a template to illustrate how to apply our method and show how its results can provide managerial insights into the auction design. In particular, the government officials running this auction have considered revising its format and we use the structural approach developed here to inform this question. Moreover, to the best of our knowledge, this is the first work in the literature that empirically shows that a CA performs well in a real-world application.

While Olivares et al. (2012) provide some empirical evidence of the benefits of the CA design in the school meals application studied here, the structural approach contributes significantly to the performance evaluation of this auction. The reduced-form analysis of Olivares et al. (2012) provides a direct estimation of the bid discounts in this CA application – for the larger packages, the discounts can be as large as 6% of the average bid price. However, it could not directly distinguish between cost synergies and the strategic markup reductions from the package discounts. In contrast, the results of our structural estimation show that cost synergies account for most of these discounts (79-86%), and that cost synergies are significant and amount up to 5% of the average cost. The rest of the bid discounts are explained by strategic markup adjustments. We also use the estimation results to pin-point the package bids for which bidders engage in strategic bundling. In addition, the estimated costs reveal that the CA achieves a high efficiency (in the order of 98-99% of the efficient allocation) and reasonable margins for the bidders (in the order of 4-5%). Drivers of this result are the relatively large cost synergies and the high level of competition in the auction; there is a reasonable number of firms and most firms compete in all units and submit many package bids. For the latter reason, firms do not seem to have enough market power to significantly harm efficiency by using the flexibility that package bidding allows with strategic motivations. Overall, the structural approach provides a direct measure of the good performance of the CA in this setting, which cannot be done through a reduced-form approach. In summary, our results suggest that package bidding and running a CA seems appropriate in this application.

Furthermore, we use the estimates of the structural model to evaluate alternative mechanism designs for the school meals auction. In particular, we compare the total payments to bidders of the current design (a first-price sealed-bid auction) against the Vickrey-Clarke-Groves (VCG)

mechanism, which generalizes the second-price auction in multi-unit settings. It is well known in the literature that truthful bidding is a dominant strategy, and therefore the VCG mechanism leads to a fully efficient allocation. However, the VCG mechanism has been criticized for numerous drawbacks, leading to a very rare use in practice. For example, Ausubel and Milgrom (2006) point out that in the face of complementarities, the total procurement cost under the VCG mechanism can be prohibitively high. Although the priority is generally given to allocative efficiency in a public procurement project, the procurement cost is also an important performance measure and can be a decisive factor to deny the use of the VCG mechanism. Therefore it is on itself interesting to see how the VCG mechanism would work in a real world application.

Having estimated the bidders' supplying costs which would be their reported bids if VCG had been used, we were able to compute the procurement cost under this counterfactual scenario. Interestingly, and contrary to the theoretical predictions in the literature, the results show that the VCG mechanism performs well in that particular application despite the significant cost synergies – the VCG procurement costs are very close to those of the first-price CAs in both 2003 and 2005. We believe this result is driven by the significant amount of competition introduced by the large number of package bids submitted by firms.

Motivated by this important observation using real-world data, Chapter 4 focuses on the revenue properties of the VCG mechanism, addressing such apparent paradox between the theory and our empirical application. More specifically, we examine the impact of competition on the revenue performance of the VCG mechanism using an asymptotic analysis. The main insight that our analysis provides is that the first order impact would be the competition measured by the amount of bids placed rather than measured simply by the number of bidders; our results emphasizes that the VCG mechanism is expected to work better when the bidders' interests are not limited to a small subset of units, resulting in high unit-wise competition, and when they place ample combination bids that contain such units in which they are interested. In our empirical application, it is relatively straightforward for bidders to estimate their supplying costs on many different packages, and therefore they were able to place a large number of bids over most of the units in the auction

– this scenario is precisely one in which our results suggest that VCG should work well. Although the analysis we provide in Chapter 4 is yet preliminary, we believe that the findings in this chapter adds useful insights for the practical usage of the VCG mechanism.

Literature review

This work adds to the stream of research in the literature that examines CAs. Most notably, Cramton et al. (2006) provide a comprehensive overview of CAs. In addition, de Vries and Vohra (2003) present a useful survey on CAs and Pekec and Rothkopf (2003) discuss the design issues in CAs.

CAs have been studied in various fields including economics, computer science, operations research, and operations management. In operations research and computer science, there is an active body of works that focus on the complexity of the auction mechanism. For example, Rothkopf et al. (1998) and Sandholm (2006) analyze the winner determination problem of CAs, identifying structures that lead to computational efficiency. By allowing combination bidding, CAs may result in excessive amount of possible number of packages on which bidder can place bids, and this can also cause challenges to bidders. Day and Raghavan (2009) propose a CA format where bidders submit several matrix bids, effectively making the bidding expression compact. Also, An et al. (2005) design a simple and efficient model for bidders' valuation and efficient bidding strategies, and evaluate the impact of the proposed strategies on the revenue of the CA.

In the economics literature, researchers have been particularly interested in the performance analysis of package bidding through CAs. There are several papers analytically examine small size CAs to study the implications of allowing package bidding. Baranov (2010) analyzes exposure and threshold problems in CA settings and examines their impact on auction performance in terms of both revenue and efficiency. Levin (1997) considers the optimal selling mechanism for complementary items. By analyzing a small CA, Maréchal and Morand (2009) also studies free-riding problem in a first-price sealed-bid CA. Krishna and Rosenthal (1996) consider the case where multiple units are auctioned simultaneously in the presence of synergies. There is also a growing body of research

that empirically study CAs. For example, Ausubel et al. (1997) and Moreton and Spiller (1998) take reduced-form approaches to analyze the extent of synergies among wireless licenses using data from Personal Communication Service (PCS) spectrum auctions. There is also an increasing number of papers that conduct structural estimation on auctions in multiple unit settings – we will discuss some of these papers when we provide review on the structural estimation papers on auctions.

Practical applications of using CAs for procurement decisions are frequently discussed in the literature. One major field of application is industrial sourcing decisions especially in B2B settings. For example, Bichler et al. (2006) and Hohner et al. (2003) study the case of using CAs by Mars Inc. for its important sourcing decisions. Similarly, Metty et al. (2005) examine Motorola’s world-wide procurement operations that use CAs. Another field where CAs have been actively discussed and implemented is the procurement of transportation services (Sheffi (2004) provides an overview of using CAs in the transportation procurement). Caplice and Sheffi (2006) examine implementation of a CA for truckload transportation and Elmaghraby and Keskinocak (2004) describe the CAs used for transportation services by retail firms such as The Home Depot Inc. and Wal-Mart Stores Inc. Cantillon and Pesendorfer (2006a) provide detailed description of the first-price CA used to procure public bus services to private providers in the city of London.

The application that we study in Chapter 3 is the CA used for procuring public school meal services in Chile. Details of this particular application including the background, history as well as the specific auction rules are presented in Epstein et al. (2002). Epstein et al. (2004) develop and discuss the techniques to solve the winner determination problem that was used in this particular application. Similarly, Catalán et al. (2009) also focus on the computational aspects, reporting the advances in the techniques of solving the winner determination problems in this application. Chapter 3 is very closely related to Olivares et al. (2012), in which the authors conduct an empirical investigation on this application. By taking a reduced-form approach, they found suggestive evidence of cost synergies. Our analysis in Chapter 3 complement their results – we provide direct estimates of the bidders’ costs and synergies as well as the efficiency of the allocation.

Our work is related to other structural estimation papers on auctions (see Athey and Haile

(2006), Hendricks and Porter (2007), and Paarsch and Hong (2006) for good surveys). Most notably, Reguant (2011) uses a first-order-conditions approach to structurally estimate a model of the day-ahead wholesale electricity market in Spain, where “complex bids” allow companies to express cost complementarities of operating across different hours in the same day. In addition, Fox and Bajari (2013) use an estimator based on a pairwise stability condition to estimate complementarities in an FCC spectrum auction, which is run in an ascending auction format without package bidding. There has also been an important structural estimation literature studying multi-unit auctions of homogeneous goods (see, e.g., Hortaçsu and McAdams (2010), Kastl (2011), and Chapman et al. (2005)). Our method developed in Chapter 2 is built upon the structural estimation method on CAs by Cantillon and Pesendorfer (2006b) which was inspired by the seminal work of Guerre et al. (2000) for single-unit auctions. Cantillon and Pesendorfer (2006b) applied their method to the London bus routes CAs with two or three units, but a direct application of their method is not possible to large-scale CAs. To the best of our knowledge, this is the first work that does structural estimation on large-scale CAs.

This work is also related to the growing literature in operations management that uses structural estimation. Olivares et al. (2008) develop a structural approach to impute the cost of overage and underage of a newsvendor, which is applied to the reservation of operating room time by an hospital. Allon et al. (2011) conduct a structural estimation to measure the implicit waiting cost of customers in the fast food industry. Similarly, Aksin-Karaesmen et al. (2013) estimate customer waiting costs but develop a dynamic structural model to explain customer abandonments in a bank’s call center. Li et al. (2011) also model consumer’s forward looking behavior through a dynamic structural model, using data from the airline industry. Lu et al. (2012) study queues in the context of retail stores and examine how customers’ waiting affects their purchasing behavior. Musalem et al. (2010) uses store-level sales data to study the effect of out-of-stocks on customer choice. We add to this stream of research by applying structural estimation in a service procurement setting, an important area in operations and supply chain management.

Finally, Chapter 4 of this work is related to the literature that examine the pitfalls of the VCG

mechanism and propose variants of the VCG mechanism that overcome such deficits. Most notably, Ausubel and Milgrom (2006) provide a detailed analysis on the drawbacks of the VCG mechanism, including low revenue problem and the possibility of collusion and shill bidding. Moreover, they show that these problems are closely related to the core of the transferable utility cooperative game played among the bidders and the auctioneer. Related to this, there is a growing body of work that develops the so called core-selecting auctions intended to overcome such drawbacks that the VCG mechanism faces. For example, Day and Milgrom (2008) discuss properties of the core-selecting auctions and relates them to the stable mating mechanisms. Also Ausubel and Milgrom (2002) explore ascending proxy auctions with package bidding, where the outcome is a point in the core. Day and Raghavan (2007) develop a bidder-Pareto-optimal core-selecting auction, addressing the computational difficulties in the payment calculation, and Day and Cramton (2012) develop further a method to compute a (unique) point in the core that minimizes the distance from the VCG payments in a given norm, using quadratic programming. One limitation in these papers is that they only perform the analysis in complete information settings. In particular, Goeree and Lien (2009) discover in a simple setting that the core-selecting auctions may result in inefficient allocation with slightly worse revenue performance than the VCG auction in an incomplete information setting. However, later work by Ausubel and Baranov (2010) identify the cases where the core-selecting proxy auction outperforms the VCG mechanism in terms of the auction revenue also through a full equilibrium analysis in a simple setting with incomplete information. We believe that our analysis in Chapter 4 contribute to this line of research, taking a different direction – we identify the environment where the VCG mechanism is expected to perform well in terms of the procurement cost (or revenue), to enhance the practicality of the mechanism.

Organization of the Dissertation

The remainder of this dissertation is organized as follows. In Chapter 2 we propose a structural estimation approach for large-scale first-price sealed-bid CAs. After describing a structural estima-

tion framework to estimate the primitives of first-price sealed-bid CAs in Section 2.2, we explain the limitations in large-scale CAs. Sections 2.3 and 2.4 develop our proposed characteristic-based markup approach and describe its further details. We then provide detailed description of the estimation procedure using our approach in Section 2.5.

We apply our structural estimation method to the Chilean school meals auction in Chapter 2. In Section 3.2, after providing a detailed description of the Chilean school meals auction and the data, we discuss how the structural model assumptions fit into this application. We then present the estimates for the distribution of the competitors' bids in Section 3.3. Using these estimates, 3.4 provides cost and markup estimation results in both small and large-scale auctions. Finally, we conduct performance analysis of the large-scale CAs using the cost and markup estimates in Section 3.5.

Chapter 4 studies the revenue properties of the VCG mechanism. Section 4.2 describes the rules of the VCG mechanism and relates its outcome to the core of coalitional games. Using the estimates from the Chilean school meals auction in Chapter 3 we perform a counterfactual analysis to see how the VCG mechanism would work in this particular application in Section 4.3. Then we provide our main analysis on the asymptotic revenue properties of the VCG mechanism in Section 4.4.

Chapter 2

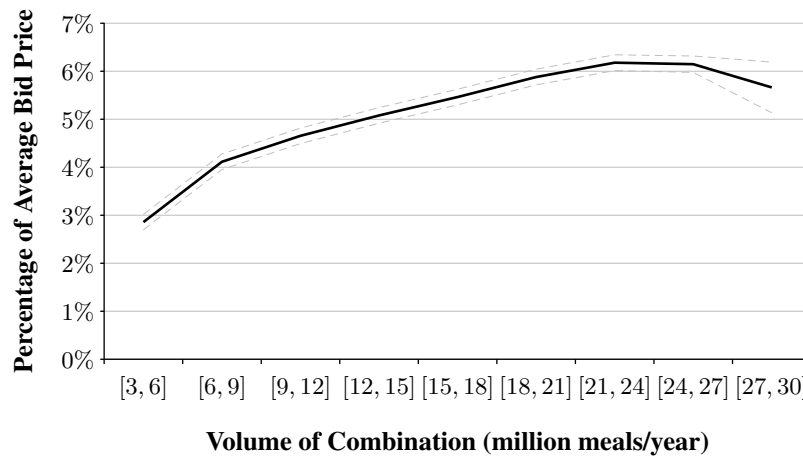
Structural Estimation Approach for Large-Scale First-Price Sealed-Bid Combinatorial Auctions

2.1 Introduction

A reduced-form analysis of the bid data can be used to provide a direct measurement of the package discounts relative to single-unit bids observed in a CA. For example, in the context of the application we study in Chapter 3, previous work by Olivares et al. (2012) provides evidence of significant package discounts (see Figure 2.1). However, the presence of package discounts is not conclusive about the performance of the auction. While bid discounts may reflect cost synergies, they could also reflect the types of strategic behavior as described in Chapter 1; bidders could inflate their single-unit bids relative to package bids to increase the probability of winning larger packages with relatively high margins, even in the absence of cost synergies. Such strategic inflation of single-unit bids also results in package discounts and a reduced-form analysis of the bid data cannot directly distinguish between this and a cost synergy-based explanation. This is limiting when evaluating the efficiency of the auction, since we expect a CA to perform well only if package bid discounts are mostly

explained by cost synergies. Moreover, since a reduced-form analysis does not identify bidders' cost information from the observed bids, it cannot be used to evaluate alternative mechanism designs.

Figure 2.1 – Average scale discounts in per-meal bid prices.



Average scale discounts in per-meal bid prices placed during 1999 - 2005. Dashed lines indicate 95% confidence interval of the estimates. The discount is measured by the decrease in the per-meal bid price when individual units are combined into a multi-unit package in the corresponding volume range. For example, the price of a meal in a package with 7 units (about 20 million meals) is cheaper by 6% of the average bid price compared to the bid prices of a meal in individual units. All the bid prices are normalized to 1999 values using consumer price index.

As an alternative to this reduced-form approach, we propose a structural estimation approach which directly identifies the bidders' costs using actual bid data. In particular, our structural method disentangles whether the discounts observed in bids are driven by cost synergies or strategic markup adjustments. Our method is based on the seminal work of Guerre et al. (2000) for single-unit auctions that was later extended by Cantillon and Pesendorfer (2006b) to a CA setting with a small number of units. The main idea behind this structural approach is to use the first-order conditions from the bidder's profit maximization problem to find the imputed costs that would rationalize the bids observed in the data. More specifically, Cantillon and Pesendorfer (2006b) conduct the structural estimation of first-price CAs in two steps. In the first step, a statistical distribution of the competitors' bids is estimated from bidding data. In the second step, the first-order conditions from

the bidder's profit maximization problem are used to find the imputed costs that would rationalize the bids observed in the data. These first-order conditions involve beliefs about the competitors' bidding behavior, and the distribution estimated in the first step is used to sample competitors' bids and form these beliefs.

In a large-scale CA such as the one we analyze in Chapter 3 – where each bidder submits in the order of hundreds or thousands of bids – a direct application of the Cantillon and Pesendorfer (2006b) method may not be possible due to the large number of decision variables in the bidder's profit maximization problem. We develop a novel approach to overcome this issue, assuming a “simplified” version of the bidder's problem where the markups charged on the package bids are chosen based on a reduced set of package characteristics. With this simplification, the bidder's problem becomes computationally and econometrically tractable so that the structural approach can be effectively applied to large-scale CAs. Recall, however, that the main objective of the structural approach is to identify the cost structure – which is a primitive in the structural model – separately from the markups, which is chosen strategically by the bidders. Therefore, we impose reasonable restrictions to the structure of the markups that reduce the complexity of the bidders problem but still provide sufficient flexibility to capture strategic behavior that can hurt the performance of a CA. In addition, we introduce a parsimonious, yet flexible parametric description of the distributions of competitors' bids for CAs that involve heterogeneous units, and scale and density discounts. This specification makes the estimation of the distributions of competitors' bids tractable. Overall, these two simplifications make the use of the structural approach feasible in large-scale CAs.

To the best of our knowledge, our work is the first structural estimation approach on large-scale CAs. We expect that our approach, based on pricing package characteristics, can be used in several real-world large-scale auctions. Notably, CAs have been actively used in outsourcing transportation services by, for example, The Home Depot (Elmaghraby and Keskinocak, 2004) and Proctor & Gamble (Sandholm, 2006). As Caplice and Sheffi (2006) point out, in many cases these auctions are done in large-scale – with a few hundreds of lanes. In addition, CAs also have been used in many industrial procurements. For example, Mars Inc. used them to source packaging materials and

raw materials (Hohner et al., 2003). In such settings, CAs are also used in large-scale, sometimes including hundreds of items in one auction (Bichler et al., 2006). The FCC has also recently allowed package bidding in its spectrum auctions, in which they can sell hundreds of licenses simultaneously (McDuff, 2003). Even though our method may need modification to accommodate different pricing and auction rules in these different settings, we believe that our idea of using characteristic-based pricing can be a useful starting point to reduce the complexity of econometric analysis in these large-scale settings.

The rest of the chapter is structured as follows. Section 2.2 develops a structural estimation framework to estimate the primitives of first-price sealed-bid CAs and explain its limitations in a large-scale setting. Section 2.3 introduces the characteristic-based markup approach for large-scale CAs and Section 2.4 provides details on how to specify the markup restrictions. We present details of the two step estimation procedure in Section 2.5, which includes the description of estimating the distribution of competitors' bids as well as estimating markups and costs based on the estimated distribution. Section 2.6 provides the main conclusions of this chapter.

2.2 Structural Estimation Approach for Combinatorial Auctions

This section develops a structural estimation framework to estimate the primitives of first-price single-round sealed-bid CAs. First, Section 2.2.1 describes the structural model and its primitives. Section 2.2.2 describes a two-step estimation approach to estimate the primitives of the structural model. Section 2.2.3 contains an informal discussion on identification. The standard structural approach to estimate auctions was pioneered by Guerre, Perrigne, and Vuong (2000) (GPV) for single-unit auctions. Cantillon and Pesendorfer (2006b) (CP) extended this approach to CAs and applied it to the London bus route auctions with three or fewer units. The structural approach introduced below closely follows the approach in CP, with some differences that we specify. Finally, Section 2.2.4 discusses the limitations of applying this approach to large-scale CAs, due to the high dimensionality of the bid vectors.

2.2.1 A Structural Model for First-Price Sealed-Bid Combinatorial Auctions

First, we describe the basic setting of a CA. Let U denote the set of N units to be procured by an auctioneer. There is a set F of supplier firms, referred to as bidders and indexed by f . A package or combination, indexed by a , is a non-empty subset of units in U . We let \mathcal{A} denote the set of all possible packages and $A = |\mathcal{A}| (= 2^N - 1)$ be the total number of them. Let b_{af} denote the bid price asked by bidder f to supply package a , and $b_f = \{b_{af}\}_{a \in \mathcal{A}}$ the *bid vector* containing all bids from that bidder.

The following assumption describes the auction format.

Assumption 2.1 (Auction Format). *The auction has a first-price single-round sealed-bid format, so that bidders submit their bids simultaneously and winning bidders are paid their submitted bid prices for the packages awarded to them. The auction mechanism determines the winning bids by solving the following mathematical integer program, referred to as the winner determination problem:*

$$\begin{aligned} & \text{minimize} && \sum_{a \in \mathcal{A}, f \in F} b_{af} x_{af} && (2.1) \\ & \text{subject to} && x \in X, \quad x_{af} = \{0, 1\}, \forall a \in \mathcal{A}, f \in F, \end{aligned}$$

where x_{af} is a binary decision variable that is equal to one if and only if package a is assigned to bidder f , and $x = \{x_{af}\}_{a \in \mathcal{A}, f \in F}$. We denote by X the set of feasible allocations; the set imposes that each unit is allocated to exactly one bidder, that each bidder can win at most one package, and potentially some additional allocative constraints.

The winner determination problem (2.1) minimizes the total procurement cost of the auctioneer, given the submitted bids. We note that the additional constraints in the set of feasible allocations could impose, for example, market share constraints that limit the maximum package size that a single bidder can be awarded, which may be used to keep a diversified supplier base. In Section 3.2, we provide more details on the winner determination problem and its integer program formulation in the context of our specific empirical application.

The structural estimation approach is based on an auction model with private information and requires assumptions on the bidders' information structure and bidding behavior in order to identify their supplying costs.

Assumption 2.2 (Bidders' Costs). *Bidders have independent private costs. In particular, given an auction, each bidder $f \in F$ gets an independent random draw of a cost vector $c_f = \{c_{af}\}_{a \in \mathcal{A}}$, in which c_{af} is the cost of supplying package a for bidder f .*

Before submitting its bid, each bidder observes its own vector of costs, but does not observe the costs' realizations of its competitors. Moreover, because costs are private, a bidder's costs only depend on its own private signal and it is not a function of the costs' realizations of other bidders. Based on this information structure, we make the following assumption on the bidders' bidding strategies.

Assumption 2.3 (Strategies). *Bidders are risk-neutral and play pure bidding strategies. In particular, for a given auction, a bidder's strategy is a function $b_f : \mathbb{R}_+^{\mathcal{A}} \mapsto \mathbb{R}_+^{\mathcal{A}}$ that depends on its own costs c_f . Bidders place bids on all possible combinations of units.*

In our sealed-bid format, bidders submit their bids in a game of incomplete information without directly observing the bids nor the cost realizations of their competitors. Therefore, bidders face uncertainty on whether they will win any given package. For each bidder, we capture this uncertainty with the vector $G_f(b_f) = \{G_{af}(b_f)\}_{a \in \mathcal{A}}$, where $G_{af}(b_f)$ is the probability that bidder f wins package a with bid vector b_f . Using vector notation, we can then write a bidder's expected profit maximization problem as:

$$\underset{b \in \mathbb{R}_+^{\mathcal{A}}}{\text{maximize}} \quad (b - c)^T G(b), \tag{2.2}$$

where v^T denotes the transpose of a vector v . Note that each bidder has its own optimization problem with its own cost and winning probability vectors. To simplify the notation, we omit the subscript f whenever clear from the context.

To formulate the optimization problem above, a bidder needs to form expectations about the bidding behaviors of its competitors, so that it can evaluate the vector of winning probabilities $G(b)$, for a given value of b . Note that if bidder f anticipates that bidder f' uses a bidding strategy $b_{f'}(\cdot)$, bidder f' 's bids are random from bidder f 's perspective; they correspond to the composition $b_{f'}(c_{f'})$, where $c_{f'}$ is the cost vector for bidder f' . Note that $c_{f'}$ is random from bidder f 's perspective, because it is private information. Assumption 2.4, described next, formalizes this. Assumptions 2.1, 2.2, 2.3, and 2.4 are kept throughout the chapter.

Assumption 2.4 (Bid Distributions). **a)** Consider a given auction and any bidder $f \in F$. From the perspective of other bidders, the bid vector of firm f , $b_f = b_f(c_f)$, is random and is given by the composition of the strategy used by firm f in the auction and its random cost vector c_f (see Assumptions 2.2 and 2.3). Accordingly, denote by $H_f(\cdot|Z)$ the distribution of b_f , where Z is a vector of observable bidders and auction characteristics. This distribution is common knowledge among bidders.

b) For all bidders $f \in F$, the competitors' random bid vectors $\{b_{f'}\}_{f' \neq f}$ are mutually independent conditional on Z .

c) For all bidders $f \in F$, the distributions of competitors' bids $\{H_{f'}(\cdot|Z)\}_{f' \neq f}$ and the winner determination problem (2.1) induce the beliefs on the winning probabilities $G_f(b_f)$, for any given b_f .

d) For all bidders $f \in F$, $H_f(\cdot|Z)$ has a continuous density everywhere.

We note that while Assumptions 2.1, 2.2, and 2.3 (or similar variations of them) are commonly made in the literature, Assumption 2.4 departs from the standard structural approach followed by CP and GPV in the following sense. The standard approach assumes that the primitives of the model such as the number of bidders, the probability distribution of costs, and the utility functions are common knowledge and that bidders play a Bayes Nash equilibrium (BNE) of the game induced by the auction. In many settings, such as the first-price single-unit auction studied in GPV, this is well justified because under mild conditions a unique symmetric pure strategy BNE always exists.

However, there is no theoretical result available that guarantee existence of a *pure strategy* BNE in a CA. As we describe next, Assumption 2.4 is weaker than assuming BNE play, but still lends itself to using the two-step estimation approach in CP. In fact, an alternative would be to assume that bidders play a *mixed strategy* BNE in the CA; this is guaranteed to exist. However, we believe that formulating a structural model in terms of pure strategies is more transparent, has a clearer interpretation, and yields simpler identification arguments.

More specifically, note that assuming pure strategy BNE play imposes two conditions for each bidder: (i) the bidder correctly anticipates the strategies of its competitors, and therefore correctly estimates the vector of winning probabilities given its own bids; and (ii) the bidder selects a bid vector that maximizes its expected profit, given its costs and the winning probabilities function. While conditions (a)-(c) in Assumption 2.4 are weaker than condition (i) in a private cost setting, they impose the same restriction over bidders' beliefs that we use in our structural estimation approach: bidders in the auction can correctly anticipate their winning probabilities. This follows because Assumption 2.4 imposes that bidders' beliefs on winning probabilities are induced by the distributions $\{H_f(\cdot|Z)\}_{f \in F}$, which are constructed with the correct strategies used by the competitors in the auction together with their actual costs' distributions. We also make a weaker assumption relative to the aforementioned condition (ii) imposed by BNE: we will only assume that each bidder selects a bid vector that satisfies the necessary first-order conditions of the expected profit maximization problem (2.2). Despite these differences in the formulation of the structural model, the first-order conditions introduced below in Section 2.2.2 are the same as the ones used by CP to identify bidders' costs.

Condition (d) in Assumption 2.4 guarantees the differentiability of the winning probability vector $G(\cdot)$ that is needed to use the first-order conditions for estimation. Note that this assumption is over the bids' distributions, that are endogenously determined in the auction game. Although we would prefer to make assumptions over model primitives that imply the assumptions on behavior, the lack of theoretical results regarding the existence and characterization of pure strategy equilibria in CAs does not allow us to follow this approach. We formalize the differentiability of $G(\cdot)$ in the

following proposition. The proof of this proposition as well as all other proofs are provided in the appendix of this chapter (see Section 2.7).

Proposition 2.1. *In a given auction, the winning probability vector $G_f(b)$ is continuous and differentiable at all b , for all bidders $f \in F$.*

2.2.2 A Two-Step Structural Estimation Method

For a given bidder, the necessary first-order conditions of the optimization problem (2.2) are given by the matrix equation:

$$c = b + \{[\mathcal{D}_b G(b)]^T\}^{-1} G(b), \quad (2.3)$$

where \mathcal{D}_b refers to the Jacobian matrix operator with respect to the variable vector b so that the ij^{th} element is $[\mathcal{D}_b G(b)]_{ij} = \frac{\partial}{\partial b_j} G_i(b)$. Note that the Jacobian $\mathcal{D}_b G(b)$ is a square matrix which can have non-zero off-diagonal elements because packages of the same bidder compete against each other. Now, for a given auction, there is one first-order-condition matrix equation per bidder. The standard structural approach assumes that the observed bid vector of each bidder satisfies equation (2.3). An important difference between first-price single-unit auctions and CAs is that in the former this first-order condition is necessary and sufficient for optimality, while in the latter it is only necessary. However, in principle it is possible to test computationally whether the observed bid vector that satisfies (2.3) is locally or globally optimal for optimization problem (2.2). We provide more details in the context of our application.

The first-order conditions (2.3), evaluated at the observed bid vector in the data, are the basis to point identify that bidder's cost vector, because the right-hand side only depends on the observed bid vector b , the winning probabilities $G(b)$ and their derivatives. Note that Assumption 2.4 implies that bidders have the correct expectations about the vector of winning probabilities $G(b)$. Hence, these winning probabilities must be consistent with the actual auction play, and therefore can be potentially estimated using bidding data from all bidders. For example, in the first-price single-unit auction analyzed by GPV, the winning probability distribution – which in this case corresponds to

the tail distribution of the competitors' minimum bid – and its derivative can be estimated non-parametrically. GPV replace these estimates in the first-order conditions to obtain point estimates of bidders' costs.

In a CA setting, $G_f(\cdot)$ is a vector of probabilities determined by the bid distributions of competitors $\{H_{f'}(\cdot|Z)\}_{f' \neq f}$ and the winner determination problem (2.1) which has no analytical solution, thereby complicating its estimation. CP uses a simulation-based two-step method to estimate $G_f(\cdot)$ and to then use the first-order conditions to obtain point estimates of the bidders' costs. This procedure can be summarized as follows:

Step 1. Use bid data to estimate the distribution of bids, $H_{f'}(\cdot|Z)$, for all firms $f' \in F$.

Step 2. To obtain the cost vector of firm f , c_f , estimate via simulation the vector of winning probabilities $G_f(b)|_{b=b_f}$ and its Jacobian matrix $D_b G_f(b)|_{b=b_f}$ evaluated at the observed bid vector submitted by firm f , b_f . Replace these on equation (2.3) to obtain a point estimate of c_f .

In step 2, winning probabilities are estimated via simulation, where each simulation run $r = 1, \dots, R$, consists of the following:

- Fix the bid vector by firm f , b_f , and for each competitor $f' \neq f$, independently sample the competitor's bids from the distribution $H_{f'}(\cdot|Z)$ estimated in the first step. Let $\{b_{f'}^r\}_{f' \neq f}$ be the bids sampled for each competitor firm in simulation run r .
- Solve the winner determination problem with bid vectors $(b_f, \{b_{f'}^r\}_{f' \neq f})$. Record the packages won by firm f with indicators $\mathbf{1}[win_{af}^r] = 1$ if and only if firm f wins package a in run r .

The winning probabilities are estimated with the empirical frequency of wins over all runs in the simulation, that is, $\hat{G}_{af}(b_f) = 1/R \sum_{r=1}^R \mathbf{1}[win_{af}^r]$, where R is an appropriately chosen large number. The Jacobian matrix of $G_f(b_f)$ is computed numerically using a similar simulation technique together with a finite-difference method. For example, to calculate the a^{th} row and s^{th} column

element of the Jacobian, one can use a central finite-difference method described by the following equation:

$$[\mathcal{D}_b G_f(b)|_{b=b_f}]_{as} = \frac{\partial}{\partial b_s} G_{af}(b)|_{b=b_f} \approx \frac{G_{af}(b_f + he_s) - G_{af}(b_f - he_s)}{2h},$$

where e_s is the s^{th} canonical vector – the s^{th} component of e_s is its only non-zero element with size one. The step-size h is an appropriately chosen small value (see, for example, Glynn (1989)). The estimations of $G_{af}(b_f + he_s)$ and $G_{af}(b_f - he_s)$ can be obtained via simulation as above. Following these two steps, once the winning probability vector $G_f(b_f)$ and its Jacobian matrix $[\mathcal{D}_b G_f(b)|_{b=b_f}]$ are estimated, the cost vector is obtained by plugging the two quantities into the first-order condition equation (2.3).

2.2.3 Identification

It is helpful to understand what patterns in the bid data drive the identification of the cost estimates from equation (2.3). This first-order condition equation implies that the bid vector is equal to a cost plus a markup vector, where the markup vector for bidder f , $-\{[\mathcal{D}_b G(b)]^T\}^{-1} G(b)$, depends implicitly on the competitors' bid distributions estimated in the first step of the structural method, $\{H_{f'}\}_{f' \neq f}$, from the bid data. It is therefore useful to analyze how these distributions affect the estimated markups and the cost estimates. For this purpose, we conducted two types of numerical experiments in small-scale instances. First, for given bids of a firm, we used the second step of the structural method to study how that firm's cost estimates changed for different distributions of competitors' bids. Second, we performed the opposite exercise and computed, for given costs, the firm's optimal bid vector for different distributions of competitors' bids. In both cases, we experimented with various model parameters; we present the results of a typical instance.

We assume a parametric model for the distribution of competitors' bids. In particular, for the base case, stand-alone unit bids from each bidder follow a bivariate normal distribution, so that the bid by competitor f' for unit $i = 1, 2$, is given by $b_{if'} \sim N(100, 15^2)$ with correlation ρ between them. We model the package bids to be the sum of the realized stand-alone bids minus a

deterministic discount of 10: $b_{12f'} = b_{1f'} + b_{2f'} - 10$.

In the first experiment, we fix a bidder f and her bids $b_f = (b_{1f}, b_{2f}, b_{12,f}) = (90, 90, 170)$; there is a discount of 10 in the package bid b_{12} . Given the model of competitors' distribution, we perform the second step estimation to compute the winning probability vector and its Jacobian matrix, and replace them in the first-order condition. This will give the cost estimates for this particular bidder f , denoted by $\hat{c} = (\hat{c}_1, \hat{c}_2, \hat{c}_{12})$, which then enable us to disentangle the cost synergies from the discount. By varying the distributional parameters, we seek to examine the impact of the distribution on the estimated cost synergies and markup adjustments. Table 2.1 reports, for each scenario, the estimated cost synergy $\hat{c}_1 + \hat{c}_2 - \hat{c}_{12}$ and the markup adjustment $\hat{m}_1 + \hat{m}_2 - \hat{m}_{12}$, where $\hat{m}_a := b_a - \hat{c}_a$.

In the second experiment, we again fix a bidder f and its costs $c_f = (c_{1f}, c_{2f}, c_{12f}) = (80, 80, 152)$; there is a cost synergy of 8 in the package cost c_{12f} . Given the model of competitors' distribution, we computed the optimal bid vector for bidder f , using a simulation-based optimization routine over a grid of bid prices. The optimal bid vector selected is the one that maximizes the expected profit among all grid points. Of course, this bid vector satisfies the first-order conditions (modulus the numerical error introduced by using a discrete grid). Table 2.2 reports, for each scenario, the optimal bid vector $b^* = (b_1^*, b_2^*, b_{12}^*)$ and the markup adjustment, $m_1^* + m_2^* - m_{12}^*$, where $m_a^* := b_a^* - c_a$ is the markup of package a .

To examine the impact of the competitors' average unit prices, we perform the experiments for different average prices on unit 1 while keeping the distribution of unit 2 unchanged. We assume independence of the unit prices ($\rho = 0$) in this case. The top panel of Table 2.1 reports the estimated cost synergy and markup adjustment in each scenario from the first experiment. The results show that as the competitors' average bid prices increase, the estimated markup adjustments are larger. Similarly, the results from the second experiment (top panel of Table 2.2) reveal that the markup adjustments become larger as the competitors' bid increase on unit 1. Note that the costs of the firm have not changed, so this markup adjustment is purely strategic. The intuition is that as competitors' bids on unit 1 increase, the firm becomes more competitive in unit 1 as well

Table 2.1 – Distributional parameters and the cost estimation.

Mean of b_1	92	96	100	104	108
Cost Synergy	6.21	5.46	4.28	2.38	-0.79
Markup Adjustment	3.79	4.54	5.72	7.62	10.79
Correlation	-0.8	-0.4	0	0.4	0.8
Cost Synergy	-0.43	2.49	4.28	6.06	8.55
Markup Adjustment	10.43	7.51	5.72	3.94	1.45
Bid Discount	6	8	10	$N(10, 2^2)$	$N(10, 4^2)$
Cost Synergy	3.33	3.83	4.28	4.26	4.20
Markup Adjustment	6.67	6.17	5.72	5.74	5.80

Impact of parameters distribution of competitors' bids on the estimated cost synergies and markup adjustments for a given firm with bid vector (90, 90, 170). The cost synergies and the markup adjustments denote the differences between the individual units and the package in the cost and in the markup, respectively.

Table 2.2 – Distributional parameters and the optimal bids.

Mean of b_1	92	96	100	104	108
Optimal Bids	(88, 89.5, 165)	(89.5, 90, 166.5)	(91, 91, 168)	(92.5, 92, 170)	(94.5, 93, 172.5)
Markup Adjustment	4.5	5	6	6.5	7
Correlation	-0.8	-0.4	0	0.4	0.8
Optimal Bids	(93, 93, 169)	(91.5, 91.5, 168)	(91, 91, 168)	(90.5, 90.5, 169)	(90, 90, 169.5)
Markup Adjustment	9	7	6	4	2.5
Bid Discount	6	8	10	$N(10, 2^2)$	$N(10, 4^2)$
Optimal Bids	(91.5, 91.5, 169.5)	(91, 91, 169)	(91, 91, 168)	(91, 91, 168.5)	(91.5, 91.5, 169)
Markup Adjustment	5.5	5	6	5.5	6

Impact of parameters distribution of competitors' bids on markup adjustments in the optimal bid vectors for a given firm with cost structure (80, 80, 152). The optimal bids are given in the form of (b_1^*, b_2^*, b_{12}^*) . The markup adjustments denote the difference in the markup between the individual units and the package.

as in the bundle. Therefore, the firm can more easily exert market power and take advantage of the additional flexibility provided by package bidding and submit a discounted package bid with strategic motivations.

To analyze the impact of correlation among single-unit prices, we use the parameters of the base case and repeated the numerical experiment with different correlation coefficients. The results, reported on the middle panels of Table 2.1, show that as the correlation becomes more negative, the estimated markup adjustments are larger and the estimated cost synergies are smaller. We observe the same pattern in the second experiment; the results show that the estimated markup adjustments are larger as the correlation becomes more negative (the middle panel of Table 2.2). This effect is related to bundling motives in the multi-product monopolist literature. Cantillon and Pesendorfer (2006b) and Olivares et al. (2012) discuss in more detail the relation between strategic markup adjustments incentives for a bidder in a CA and bundling incentives for a multi-product monopolist.

Finally, to explore the impact of competitors' discounts in the package bid b_{12} , we conducted the numerical experiment using different levels of deterministic discount as well as normally distributed random discount with different variances. We chose the ranges of the level and the variance of discounts to be close to what we observe in our data. The results from the first experiment, reported in the bottom panel of Table 2.1, suggest that the magnitude and variance of the discounts do not affect much the estimated markup adjustments and cost synergies if kept within these ranges. The results from the second experiment shown in the bottom panel of Table 2.2 also suggest that the optimal bid prices and markup adjustments are not affected much in these parameter ranges.

The results in the numerical examples provide useful insights about identification of costs. Recall that observed package discounts are the sum of cost synergies plus markup adjustments. Hence, the results suggest that as the means of the estimated distribution of competitors' bids become larger, and as the correlation among individual prices become smaller, the estimated markup adjustments should increase, and therefore, the fraction of the package discounts explained by cost synergies should decrease. In summary, when estimating the distribution of competitors' bids, correctly capturing the correlation and the heterogeneity of unit prices plays an important role in the

estimation of the bidders' cost structures.

2.2.4 Limitations in Large-Scale CAs

CP were able to effectively use the previous approach in auctions of at most three units. However, there are two significant limitations in using the standard approach in large-scale CAs with more units.

First, in large-scale CAs that are typically found in practice (including our empirical application), firms may submit hundreds or even thousands of bids. In that case, the bid vectors $\{b_f\}_{f \in F}$, and therefore the distributions $\{H_f(\cdot|Z)\}_{f \in F}$ that need to be estimated in the first step, are high dimensional. For this reason, parametric restrictions need to be imposed to make the estimation tractable. However, it is important to allow for sufficient flexibility in these restrictions. In their application, CP developed a reasonable parametric model that balances flexibility with feasibility in the estimation. We extend their approach to large-scale CAs. In particular, in Section 2.5.1 we provide more details about a parsimonious, yet flexible parametric description of the distributions of competitors' bids for CAs that involve geographically dispersed and heterogeneous units as well as scale and density discounts. These distributions are then taken as an input for the second step.

Second, there is a limitation in the second step of the CP approach when applied to large-scale CAs: the high dimensionality of the first-order conditions (2.3). The dimension of this matrix equation is determined by the number of bids, which increases exponentially with the number of units in the CA. As the number of bids submitted by a bidder gets large, the winning probability of each bid is likely to become very small and the simulation errors in estimating these rare-event probabilities become large. Moreover, equation (2.3) requires taking derivatives over a large number of variables; simulation error for these quantities may be even larger. These problems may not be resolved by simply increasing the length of the simulation runs, because in the course of a simulation run one needs to repeatedly solve the winner determination problem which is known to be NP-hard, and solving these problems gets increasingly expensive computationally as the scale increases. For example, using state-of-the-art solvers for integer programming, it takes in the order of seconds to

solve a single instance of the winner determination problem in our empirical application that we examine in Chapter 3.¹ Hence, computation of $G(b)$ and $\mathcal{D}_b G(b)$ via simulation becomes quickly intractable as the number of units auctioned increases. The difficulties in estimating $G(b)$ make it also unreasonable to assume that bidders would be able to solve (2.2) optimally.

An important methodological contribution of this chapter is to address the second problem – the high dimensionality of the first-order conditions (2.3). Our approach imposes reasonable restrictions in the structure of the markups which allow us to reduce the dimensionality of the problem. We describe this approach in detail in the next section.

2.3 The Characteristic-Based Markup Approach for Large-Scale CAs

Our model is based on the approach described in Section 2.2. As mentioned above, a significant complication of using this model in large-scale CAs is that the dimensionality of the first-order conditions is too large. We develop an approach to reduce the dimensionality of the problem by imposing additional assumptions on the bidders' bidding behavior that have behavioral appeal and make the estimation approach econometrically and computationally feasible in large-scale CAs.

Notice that in the first-order condition (2.3), the markup term $-\{[\mathcal{D}_b G(b)]^T\}^{-1} G(b)$ provides the flexibility to the bidder to assign a different and separate markup to each package. Hence, we refer to this model as the *full-dimension* model. In contrast, we propose that the markup of each bid is specified through a reduced set of package characteristics. Specifically, let w_a be a row vector of characteristics describing package a , with dimension $\dim(w_a) = d$ that could be potentially much smaller than A . The markup for package a is given by the linear function $w_a \theta$, where θ is a (column) vector of dimension d specifying the markup components associated with each package characteristic. Instead of choosing the markup for each package, the bidder now chooses θ . Let

¹In the application described in Chapter 3, typically we have about 30 units in an auction. To solve the winner determination problems, we used CPLEX V12.1 called by a C routine and ran on Columbia Business school's shared cluster, where each machine has eight 2.4 GHz CPUs.

$W \in \Re^{A \times d}$ be a matrix containing the characteristics of all packages, so that the a^{th} row of W is w_a . The following assumption, kept throughout the chapter, formalizes this simplification to the bidders' bidding behavior.

Assumption 2.5 (Characteristic-Based Markups). *Consider a given bidder in a particular auction. Its bid vector is determined by $b = c + W\theta$, where W is a fixed $(A \times d)$ -dimensional matrix of package characteristics and θ is a d -dimensional vector of decision variables chosen by the bidder.*

It is worth noting that our approach allows the specification of W to vary across bidders. Like other quantities, we omit the firm index f in W for notational simplicity. Now under this assumption, the bidder's optimization problem becomes:

$$\underset{\theta \in \Re^d}{\text{maximize}} \quad (W\theta)^T G(W\theta + c), \quad (2.4)$$

whose first-order conditions yield:

$$[\mathcal{D}_\theta W^T G(W\theta + c)]^T \theta = -W^T G(W\theta + c). \quad (2.5)$$

Here again the ij^{th} element of the Jacobian matrix above is $[\mathcal{D}_\theta W^T G(W\theta + c)]_{ij} = \frac{\partial}{\partial \theta_j} [W^T G(W\theta + c)]_i = \frac{\partial}{\partial \theta_j} W_i^T G(W\theta + c)$, where W_i is the i^{th} column of matrix W . Re-arranging and replacing terms, we can solve for the decision vector θ as follows:

$$\theta = - \{ [\mathcal{D}_\theta W^T G(b)]^T \}^{-1} W^T G(b). \quad (2.6)$$

As in GPV and CP, this first-order-condition equation constitutes the basis of identification in our structural model. Again, note that in each auction there is one first-order-condition matrix equation per bidder and different bidders may have different characteristic matrices W . For each bidder, under Assumption 2.5, the cost is given by $c = b - W\theta$. Hence, costs are uniquely determined by θ , and moreover, if the matrix $\mathcal{D}_\theta W^T G(b)$ is invertible, equation (2.6) uniquely identifies the markup vector θ . Therefore, equation (2.6) provides an alternative to (2.3) to estimate costs. We provide conditions for the invertibility of this matrix in Section 2.4.3. In what follows, we formalize this discussion with the following assumption that is kept throughout the chapter.

Assumption 2.6 (First-Order Conditions). *The observed bid vector of a given bidder in the auction satisfies the necessary first-order conditions of the characteristic-based markup model given by equation (2.5).*

As with equation (2.3), the right-hand side of equation (2.6) can be estimated purely from observed bidding data. In fact, our approach using equation (2.6) closely follows the two-step method described in Section 2.2.2. However, the reduced dimensionality of equation (2.6) significantly simplifies the computation burden in the second step, making it feasible in large-scale applications.

To see this, note that (2.6) is similar to (2.3), with the winning probability vector $G(b)$ and its Jacobian matrix $\mathcal{D}_b G(b)$ replaced by the vector $W^T G(b)$ and its Jacobian matrix $\mathcal{D}_\theta W^T G(b)$, which is now with respect to the markup vector θ . The first simplification is that the derivatives are now taken with respect to $d \ll A$ variables, effectively reducing the dimension of the problem. Second, in the specifications we propose later, we will see that each element of the vector $W^T G(b)$ is a (weighted) sum of winning probabilities over many packages. These aggregated probabilities are larger than the winning probabilities of each individual package, and therefore easier to estimate via simulation. Besides, there are fewer probabilities to be estimated; altogether these make the second step computationally tractable.

One apparent limitation of Assumption 2.5 is that the markup is additive as oppose to multiplicative to costs, which may be more appropriate in some applications. A multiplicative markup, however, would lead to different first-order conditions from which it is mathematically intractable to identify bidders' costs using bid data. A relatively simple way to make the additive assumption less restrictive is to include package characteristics in W which are related to costs, so that the markup can be scaled based on these cost-related characteristics. This approach is effective when the cost heterogeneity across packages can be captured, at least partially, by a reduced set of known variables. We come back to this point in the sequel.

Note that the characteristic-based markup model is very general and flexible in the specification of markup structures. For example, if we specify the package-characteristic matrix W as the

identity matrix, each package has its own markup and we are back to the full-dimension problem (2.2). On the opposite extreme, one could choose $d = 1$ so that the markups of all packages are determined by a single decision variable; while this specification significantly reduces the dimension of the problem, this may be too restrictive. Between these two extremes there are many possible specifications for W . Different specifications may be chosen depending on the details of the large-scale application at hand, for example, those mentioned in Section 2.1. In particular, the next section describes an approach to specify W that is sufficiently flexible to capture strategic markup adjustments that arise in package bidding, but that at the same time is parsimonious and maintains computational tractability.

2.4 Specifying Markup Restrictions

Recall from our discussion in the introduction that an important objective of our structural estimation approach is to measure what portion of the observed package discounts can be attributed to cost synergies versus strategic markup reductions. In fact, economic theory and the numerical examples described in Section 2.2.3 show that firms may have incentives to make strategic markup adjustments. Therefore, package discounts cannot be fully accrued to cost synergies a-priori. Moreover, previous literature suggests that scale is likely to be the main driver of these strategic markup adjustments (see the references in the introduction, in particular Olivares et al. (2012), for a more detailed discussion). For this reason, we focus on developing a specification that allows for markups to vary on the size of the package. Doing so helps separating what portion of the volume discounts observed in the bid data arise from markup adjustments vis-à-vis cost synergies. In Section 2.4.1 we show how to specify W using size as a package characteristic. Section 2.4.2 extends this base model by providing additional flexibility to provide a better approximation of the estimates we would obtain with the full-dimension model. Finally, Section 2.4.3 describes some additional requirements on the package characteristic matrix W that ensure identification of the firms' costs.

2.4.1 Group-Based and Size-Based Markup Models

The key idea in our estimation approach is to impose restrictions on the markup structure to reduce the dimensionality of the bidders' problem. A special case of the characteristic-based markup approach is to create a partition of the set of all packages, and allow each group (or set) of the partition to have its own separate markup parameter. This group markup parameter then determines the markup of all the packages in the corresponding group. This approach, referred to as the *group-based markup model*, is defined formally as follows.

Definition. A markup specification follows a *group-based markup model* if each row of the package-characteristic matrix W is composed by zeroes except for one and only one positive component.

Consider the following group-based markup model. Let $\{\mathcal{A}_s\}_{s=1}^S$ be a partition that covers all possible packages. From this partition, a potential candidate for the package-characteristic matrix $W \in \mathbb{R}^{A \times S}$ can be generated using indicator variables $W_{as} = \mathbf{1}[\text{package } a \text{ belongs to set } \mathcal{A}_s]$. With this specification, the term $W^T G(b)$ in equation (2.6) has the following form:

$$W^T G(b) = \begin{bmatrix} W_1^T G(b) \\ W_2^T G(b) \\ \vdots \\ W_S^T G(b) \end{bmatrix} = \begin{bmatrix} \text{Probability of winning any package in } \mathcal{A}_1 \\ \text{Probability of winning any package in } \mathcal{A}_2 \\ \vdots \\ \text{Probability of winning any package in } \mathcal{A}_S \end{bmatrix}.$$

As seen above, the group-based markup model could significantly reduce the dimensionality of the problem; if $S \ll A$, there are much less probabilities to estimate, as well as derivatives to take in the Jacobian matrix. Moreover, while the winning probability of any given package a is typically small and hard to estimate via simulation, the winning probability of a *group* of packages aggregates these individual probabilities over a potentially large set of packages and, therefore, is often much larger. For this reason, we require fewer simulation runs to obtain precise estimates of these aggregated probabilities. All this makes the computation of the right-hand side of the first-

order condition (2.6) tractable.

A special case of the group-based model is when the packages are grouped by their sizes. For some defined measure of package size (e.g. the number of units in the package), let \mathcal{A}_s be the set of all packages of size s . The markup parameter θ_s represents the *common* markup charged to all packages of size s ; the bidder chooses S different markups, one for each possible size. This is referred to as the *pure size-based markup model*. Recall that we want to disentangle what portion of the observed bid discounts are explained by markup adjustments when bidders submit larger packages. The pure size-based markup provides the minimum level of flexibility to capture such strategic markup adjustments, and therefore, we believe it is a reasonable starting point to impose markup restrictions in our approach.

Having defined the pure size-based markup model, we seek to understand whether it provides a good approximation to the estimates of the full-dimension model. To do so, we provide an analytical comparison of the markups estimated by the full-dimension model with those estimated via the group-based markup approach.

Proposition 2.2. *Consider a bidder submitting a bid vector b in a CA. Assume that all bids in b have strictly positive probabilities of winning.*

a) *Suppose the CA has A packages. Let θ_a , $a = 1, \dots, A$, be the estimated markup for package a by the full-dimension model (2.3), and θ_u be the common markup estimated by the group-based markup model (2.6) when the A packages form a single group, that is, $b_a = c_a + \theta_u$, $a = 1, \dots, A$. Then, $\theta_u = \sum_{a=1}^A \beta_a \theta_a$, for appropriately defined weights $\beta_a \geq 0$, $a = 1, \dots, K$, that satisfy $\sum_{a=1}^A \beta_a = 1$.*

b) *Suppose the CA has 2 units. Let $(\theta_1, \theta_2, \theta_{12})$ be the estimated markup vector by the full-dimension model (2.3) and let (θ_u, θ_v) be the estimated markup vector by the size-based model (2.6), where θ_u is the common markup for single units and θ_v is the markup for the package. Then, $\theta_u = \beta\theta_1 + (1 - \beta)\theta_2$ and $\theta_v = \theta_{12} + \gamma(\theta_1 - \theta_2)$, for appropriately defined constants $\beta \geq 0$ and γ .*

We provide the proof of the above proposition in the appendix of this chapter (see Section 2.7)

as well as the detailed expressions for the constants β 's and γ , which are functions of the partial derivatives of the winning probabilities evaluated at the observed bid vector b (we omit them here for brevity). The results from Proposition 2.2 provide important insights regarding the implications on the group-based markup estimates. First, from the result in part (b), we observe that grouping the units affect the estimated markup of the package, where the impact depends on the coefficient γ and the difference of the individual unit markups. If the unit markups are very close to each other, the effect of grouping on the package markup will be negligible. Moreover, note that if γ is small, the effect of having a common markup for the units has a negligible effect on the estimated markup for the package. In fact, extensive numerical experiments have shown that in our application, grouping a set of packages so that they share a common markup merely affects the markups of other packages that are not in that particular group. In Section 3.4.1, we provide some numerical validations supporting these observations in the context of our empirical application.

The previous discussion together with part (a) in the proposition can be summarized as follows. Consider a situation in which the packages in the set \mathcal{A}_s are grouped together, and let θ_s be the common markup estimated by the group-based markup model. Let θ_a be the individual markup for package $a \in \mathcal{A}_s$ estimated by the full-dimension model. Then, the previous discussion basically suggests that:

$$\theta_s \approx \sum_{a \in \mathcal{A}_s} \beta_a \theta_a,$$

where $\beta_a \geq 0, \forall a \in \mathcal{A}_s$, are appropriately defined weights that satisfy $\sum_{a \in \mathcal{A}_s} \beta_a = 1$. The result is useful because it suggests that the estimated common markup is a convex combination of the individual markups we would obtain from the full-dimension model (if we were able to estimate them). Moreover,

$$|\theta_a - \theta_s| \approx \left| \theta_a - \sum_{a' \in \mathcal{A}_s} \beta_{a'} \theta_{a'} \right| \leq \sum_{a' \in \mathcal{A}_s, a' \neq a} \beta_{a'} |\theta_a - \theta_{a'}|, \forall a \in \mathcal{A}_s.$$

Therefore, the estimated common markup would be a good approximation to the individual markup estimates from the full-dimension model if the latter markups are close to each other. Of course, checking this condition is computationally intractable, because we would need to solve the full-

dimension model. The next section describes a computationally tractable heuristic that aims at providing more flexibility in the markup restrictions without increasing much the computational burden of the method.

2.4.2 A Refinement of the Size-Based Markup Model

As suggested above, the main issue with the size-based markup model would be whether the packages in the same size group are significantly heterogenous or not. For example, if one package has a significantly different markup to the rest of the group in the full-dimension model (if we were able to estimate it), it ideally should not be a part of the group.

Recall that one of the difficulties in the full-dimension model arises in the computation of a large number of small winning probabilities via simulation. However, for a given firm, there still may be a small number of packages with reasonably large winning probabilities that can actually be computed with precision – we refer to these as “special packages”. It is then possible to assign and estimate a separate markup for these special packages, without forcing them into a group and therefore alleviating the potential biases previously discussed. Given their high winning probabilities, special packages are also more likely to be part of the winning CA allocation, so it is useful to obtain more precise estimates for their markups. Finally, in Section 3.4.1 we provide empirical evidence in the context of our application that high winning probability packages tend to have larger estimated absolute per-meal markups in the full-dimension model relative to the rest of their group. Hence, removing them from the group and estimating a separate markup for each of them is likely to reduce the bias associated with grouping in a significant way.

Another extreme alternative would be to estimate the model with special packages only, ignoring the rest of the packages. Although the cost information of bids with small winning probabilities may be less important for the estimation of the performance measures (since they are less likely to be part of the winning CA allocation), they cannot be entirely eliminated in the estimation procedure. One reason is that these packages may have significant winning probability in aggregate, and therefore ignoring them in the first-order condition (2.6) can result in an inaccurate estimation

for the costs of the large winning probability bids. In fact, we have estimated models with and without the packages with small winning probabilities and found that the markup estimates of the special packages changed substantially.

In addition, in some applications (including the one analyzed in this work) the units of the same group can be heterogeneous and this could lead to differences in markups, even after separating the special packages. In our application, units differ in their volume and so packages with the same number but different composition of units could have different markups. To account for this heterogeneity, let v_i be the volume of unit i and define $v_a = \sum_{i \in a} v_i$ as the total volume of the package. The package-characteristic matrix W can be specified as $W_{as} = v_a \cdot \mathbf{1}[\text{package } a \text{ has } s \text{ units}]$, which is also in the class of group-based markups. With this specification, the packages in the same size group will share the same markup parameter θ_s which is the per unit of volume markup, so that the markup of package a is $v_a \theta_s$. Also, the s^{th} element of the vector $W^T G$ is equal to the expected volume of winning packages of size s . Overall, this specification makes the additive nature of Assumption 2.5 less restrictive.

Based on the previous insights, we propose the following heuristic to build the package-characteristic matrix W for a given firm:

1. Group packages according to their sizes and let $W_{as} = v_a \cdot \mathbf{1}[\text{package } a \text{ has } s \text{ units}]$, so that initially all packages with the same number of units share the same markup per unit of volume.
2. Run a simulation to obtain rough estimates of the winning probabilities of each package; this simulation is quicker to run than solving for the first-order conditions. For each size group, identify bids that have high winning probabilities relative to the rest. Each of these packages is associated with a separate individual markup parameter.
3. For each size, further divide the rest of the bids into two groups: medium and low winning probability groups. This step is motivated by the observation discussed in Section 3.4.1 that winning probability is related to the magnitude of markups. In that section, we further justify this step with empirical evidence in the context of our application.

Through this heuristic procedure, we construct the corresponding package-characteristic matrix W for each firm allowing for separate markups for each of the specified groups (including the groups with a single special package). We refer to this approach to define the package-characteristic matrix as the *extended size-based markup model*, which is a particular case of a group-based markup model. For each firm, we use this specification within a two-step method similar to the one described in Section 2.2.2. In the first step, we parametrically estimate the distribution of competitors' bids. This procedure requires a separate treatment and is described in detail in Section 2.5.1. In the second step, we use the specification of W given by the previously described heuristic in the first-order condition (2.6) to obtain a point estimate of θ , and therefore of c .

Our heuristic based on the extended size-based markup model aims to improve the approximation to the full-dimension model starting from the pure size-based markup model. However, it is important to provide some empirical validation of this claim. For this purpose, we collected data from two exceptionally small CAs in our application. The full-dimension approach was feasible to implement in these smaller auctions and was compared with the results provided by the extended size-based model using our heuristic method described above. Notably, the results of this analysis presented in Section 3.4.1, suggest that the markups estimated with the two approaches are very similar, providing support for our method. Finally, we note that in the related context of multi-product monopolist pricing, Chu et al. (2011) provides computational and empirical evidence of the effectiveness of size-based pricing in some settings. They also show examples where this restricted pricing strategy is used in practice.

2.4.3 Further Requirements on the Package-Characteristics Matrix

We finish this section discussing issues related to identification that are important for the specification of W . In particular, we provide conditions for which $\mathcal{D}_\theta W^T G(b)$ is invertible in equation (2.6), and therefore, the first-order conditions uniquely identify the markup vector θ , and hence the costs.

In some applications, bidders may not submit bids on all packages. For example, in our

empirical application analyzed in Chapter 3, firms do not place bids on all possible combinations because of two reasons: (1) firms have limits on the maximum number of units that can be included in a package (these limits depend on the firm's financial capacity); and (2) the number of possible combinations is too large. However, this case can still be handled with our proposed approach by treating the missing packages as bids with very high prices that have no chances of winning. We refer to the bids which never win as *irrelevant bids*. In addition, some bids that are actually submitted may also be irrelevant, in the sense that they have zero probabilities of winning. For example, this could arise as a strategic decision in order not to win a specific package when the auction rules require submission of bid prices on all packages. In contrast, a *relevant bid* has a strictly positive probability of winning.

CP shows that in the full-dimension model irrelevant bids do not play a role in the first-order conditions, and one can identify the markups for relevant bids after eliminating irrelevant bids from the estimation. We extend this discussion to our group-based markup model, where we can still identify markup variables as long as each group has at least one relevant bid. To see that, recall that each column of package characteristics in W is associated with a markup variable in the bidder's decision θ . We say that a package a is associated with the markup variable θ_i if $W_{ai} \neq 0$, that is, the bid price of a depends on the value of θ_i . The following lemma is useful to characterize the conditions needed for identification. It is also used for the proof of Theorem 2.1. We provide the proof of this lemma in Section 2.7.

Lemma 2.4.1. *Consider a given bidder and auction. For any package $a \in \mathcal{A}$, $G_a(b) = 0$ implies $\frac{\partial}{\partial \theta_i} G_a(W\theta + c) = 0$, for all $i = 1, \dots, d$.*

The lemma implies that if all the bids associated with a markup variable θ_i are irrelevant, then the i^{th} row of the Jacobian matrix $\mathcal{D}_\theta W^T G(b)$ will be all zero, and the matrix will not be invertible. In this case, the markup vector of that bidder will not be identified, because (2.6) requires invertibility of the Jacobian matrix. CP shows that, in the full-dimension model, this problem can be resolved by eliminating irrelevant bids from the estimation, and by doing so, one can still identify markups

for the relevant bids. We extend the discussion to our characteristic-based model, where we can still identify markup variables as long as each of them has at least some relevant bids that are associated with it. In what follows we examine this issue in more detail.

Consider a given firm. Without loss of generality, we assume packages are ordered such that all the relevant bid packages (superscripted by R) are followed by the group of irrelevant bid packages (superscripted by I), so that:

$$W = \begin{bmatrix} W^R \\ \dots \\ W^I \end{bmatrix}, \quad c = \begin{bmatrix} c^R \\ \dots \\ c^I \end{bmatrix}, \quad b = \begin{bmatrix} b^R \\ \dots \\ b^I \end{bmatrix}, \quad \text{and} \quad G(b) = \begin{bmatrix} G^R(b) \\ \dots \\ G^I(b) \end{bmatrix}.$$

By replacing these terms in equation (2.6), we obtain:

$$\begin{aligned} \theta &= - \left\{ [\mathcal{D}_\theta ((W^R)^T G^R(b) + (W^I)^T G^I(b))]^T \right\}^{-1} ((W^R)^T G^R(b) + (W^I)^T G^I(b)) \\ &= - \left\{ [\mathcal{D}_\theta (W^R)^T G^R(b)]^T \right\}^{-1} (W^R)^T G^R(b) \\ &= - \left\{ [\mathcal{D}_\theta (W^R)^T G^R(b^R)]^T \right\}^{-1} (W^R)^T G^R(b^R), \end{aligned} \tag{2.7}$$

where the second to last equation follows from $G^I(b) = 0$ and Lemma 2.4.1. In the last equation, it is implicitly assumed that the bidder only submit relevant bids. Because irrelevant bids never win and moreover small changes in the markup vector will not turn them into relevant bids by Lemma 2.4.1, it is the same as if the bidder would not have submitted them (recall that non submitted bids are also irrelevant). Therefore, the right-hand sides of equations (2.6) and (2.7) are equivalent. Consequently, the elimination of irrelevant bids will not affect the identification of the markup vector θ as long as the Jacobian in equation (2.7) is invertible. Finally, the following theorem provides necessary and sufficient conditions to ensure the invertibility of the Jacobian $D_\theta W^T G$ for the class of group-based markup models.

Theorem 2.1. *Consider a given bidder and auction. Assume that the package-characteristic matrix W is the specification of a group-based markup model. If the Jacobian matrix $\mathcal{D}_\theta W^T G(b)$ evaluated at the observed bid vector b is invertible, then every group contains at least one relevant bid.*

The latter condition becomes sufficient for the invertibility of the Jacobian matrix if the following additional conditions hold: (i) the observed bid vector is such that $b - c \geq 0$; and (ii) all the elements of W are non-negative ($W \geq 0$). In this case, the markup vector θ is uniquely identified by equation (2.6).

Note that the assumption $b - c \geq 0$ is a mild rationality assumption on bidders' behavior that guarantees bidders make positive profits on each package conditional on winning that package. Also note that under the previous assumption, assuming $W \geq 0$ is essentially done without loss of generality, in the sense that for a W matrix with negative entries, one can find another W matrix with non-negative entries that produces the same markup estimates. A practical implication of the theorem is that when implementing the heuristic described in Section 2.4 one needs to make sure that each group of packages must include at least one relevant bid. After we imposed this, we were always able to invert the Jacobian matrix computationally.

Another implication of the previous discussion is that the proposed method can only identify the cost structure of packages associated with relevant bids, that is, $c^R = b^R - W^R\theta$, because irrelevant bids provide no information to the first-order conditions (c^R is the cost vector including only components associated with relevant bid packages). Although it is not possible to point identify the costs of irrelevant bids, CP showed that bounds on the costs of such irrelevant bid packages can be obtained. However, computing these bounds is computationally expensive and becomes infeasible in large-scale CAs. Instead, we infer the costs of those irrelevant bids using extrapolation. We will come back to this point in Section 3.5.1 in the context of our application.

Finally, an important assumption needed for our approach is that bidders can win at most one package. This is a frequent requirement in many real-world CAs, specially in settings with rich and expressive package bidding. Without this requirement, it may not be possible to point identify costs. For example, consider a CA with 2 units and suppose a bidder only submits bids for the individual units. Suppose the bidder has a positive chance of winning both individual bids simultaneously, which is equivalent to winning the two-unit package. Then, we have three unknowns to estimate

(the cost for each individual unit and the cost for the package), but only two equations (the two first-order conditions with respect to the individual bid prices).

2.5 Estimation

In this section, we propose an estimation method for large-scale CAs using the characteristic-based markup model that we developed in the previous section. We adopt the two-step approach introduced in Section 2.2.2. In the first step, we need to estimate the distributions of competitors' bids, $\{H_f(\cdot|Z)\}_{f \in F}$, which are then used in the simulation-based routine in the second step to sample competitors' bids and estimate the terms in the first-order conditions given by equation (2.6). Sections 2.3 and 2.4 addressed the complexity introduced in the second step due to the large-scale nature of the auction. In particular, we simplified the first-order conditions of bidders by imposing some structure in their markups. It was important that the structure was flexible enough to allow for strategic markup adjustments.

As discussed in 2.2.4, for the estimation of large-scale CAs, another important challenge arises in the first step. The complication in the first step is that in large-scale CAs, firms may submit hundreds or even thousands of bids. Therefore, the bid vectors $\{b_f\}_{f \in F}$ are high dimensional precluding the use of a non-parametric approach like GPV to estimate the distribution of competitors' bids; CP faced a similar challenge even for a three-unit CA. We address this challenge by proposing a parametric approach to model the bid distribution which can be used in CAs that involve geographically dispersed and heterogeneous units that are subject to discounts due to scale and density, like in our application. Section 2.5.1 describes our parametric approach. In Section 2.5.2, we then explain the simulation-based estimation procedure in the second step under the characteristic-based markup model. Finally, Section 2.5.3 describes an algorithm to enhance the computational efficiency in the second step estimation.

2.5.1 Estimating the Distribution of Competitors' Bids

It is important to emphasize that the simplifications in the two steps have different objectives. In the first step, the objective is to introduce a parametric model that fits the competitive bidding landscape data well. In the second step, the objective is to simplify the bidders' decision space in the first-order conditions. We note that the parametric model of the *bid* data in the first step will be more flexible than the model for *markups* in the second step, because it will allow for scale and density discounts both of which could be observed in the data. On the other hand, our extended size-based markup model explicitly considers strategic discounts associated with scale, but not to density. The reason is that, as mentioned in Section 2.4, economic theory suggests that scale (and not density) is likely to be the main driver of the strategic markup adjustments we are trying to identify. In Section 3.4.1 we provide some empirical evidence of this claim in the context of our application.

The parametric approach we follow has an important difference with CP in that in our estimation method the identification of the distribution of competitors' bids is based on variation across package bids and firms in a *single auction*, and hence exploits the large number of package bids, which is a key characteristic of large-scale CAs. In the standard structural approach to auctions (including CP and GPV), the estimation of the bid distribution uses variation in a *cross-section of auctions*, implicitly assuming that the same equilibrium is being played across these auctions. Hence, our identification strategy can be more robust when there is unobserved heterogeneity across auctions – changes in the auction characteristics and firm characteristics from auction to auction that are observed by bidders but unobserved by the econometrician (see Krasnokutskaya (2011) for a more detailed discussion on this issue).

Imposing parametric restrictions to the multivariate bid distribution needs to balance flexibility with estimation feasibility. There are three key aspects typical in applications of CAs that are important to account for: (i) units are heterogeneous; (ii) the correlation structure among the bids from the same bidder; and (iii) package discounts. We discuss each of these three in what follows.

First, in many CAs, the bid prices are heterogenous among units and among firms. In ap-

plications that involve logistics and transportation across dispersed geographic units (as the one we study), heterogeneity among units arises primarily from the costs of serving different territories. For example, units located in isolated rural areas tend to be more expensive than units in urban areas. There is also heterogeneity across firms: some firms may have national presence, are vertically integrated, and may have well functioning and efficient supply chains; other firms may be more rustic local firms.

Second, package bids of the same bidder may be correlated. In CAs, there are two main factors that can generate correlation between bids. First, a bidder that has a high cost in a given unit is likely to submit higher prices for all packages containing that unit. Second, if there are local advantages, a supplier charging a low price for a unit may also charge lower prices for nearby units. Hence, the unit composition of the package bids together with the spatial distribution of the territorial units provides a natural way to parameterize the covariance structure among package bids. As described in Section 2.2.3, the correlation structure of the competitors' bids has direct implications on the incentives to engage in strategic markup adjustments, so it is important to allow for a flexible covariance structure that incorporates these effects.

Third, CAs exhibit package discounts in the bids; the price per unit may decrease as the size of the package increases. In applications where economies of density matter, the *geographic location* can be another factor that determines the magnitude of the discounts; for example, combining two units located nearby could lead to larger discounts (relative to a package with two distant units).

Accordingly, we develop the following econometric model for package bids that captures heterogeneity among units, correlation, and discounts. In particular, from the perspective of all other firms, firm f 's bids are specified by the following parametric model:

$$b_{af} = -g^{scale}(v_a, \beta_{k(f)}^{scale}) - \sum_{c \in Cl(a)} g^{density}(v_c, \beta_{k(f)}^{density}) \cdot \frac{v_c}{v_a} + \sum_{i \in a} \tilde{\delta}_{if} \frac{v_i}{v_a} + \tilde{\varepsilon}_{af}. \quad (2.8)$$

We note that the structure in this equation that separates individual prices with discounts is motivated by Olivares et al. (2012). As defined earlier, v_i denotes the volume of unit i and $v_a = \sum_{i \in a} v_i$ is

the total volume of package a . With some abuse of notation, the dependent variable, b_{af} , denotes the price per unit of volume submitted by firm f for package a ; that is, the actual bid price divided by the total volume of the package, v_a . The four terms in the right-hand side of equation (2.8) capture: (i) the effect of discounts due to size or scale (g^{scale}); (ii) the effect of discounts due to density ($g^{density}$); (iii) the effect of the specific units contained in the package (the sum over units i in package a), where $\tilde{\delta}_{if}$ can be viewed as an *average implicit price* that bidder f is charging for unit i among all the packages submitted, net of any scale and density discounts; and (iv) a Gaussian error term $\tilde{\varepsilon}_{af}$ capturing other factors affecting the bid price. It is important to emphasize that the discount functions g^{scale} and $g^{density}$ should not be interpreted directly as cost synergies because part of the discounts could also arise from strategic behavior.

This parametric specification also assumes that the bids across bidders are independent and that the bid distribution of a bidder depends only on its own characteristics, $H_f(\cdot|Z) = H_f(\cdot|Z_f)$. Assuming that the bid distribution of a firm depends only on its own characteristics is not restrictive when the distribution is estimated separately for each auction, because the characteristics of the competitors are held fixed within the auction. In addition, to avoid making strong assumptions on how firms choose which combinations to submit, we use the same package composition observed in the data. That is, when generating competitors' bid prices, we fix the packages on which the bids were actually submitted by a particular bidder and simulate new prices for these packages.

The competitors' bid distribution captures the relevant uncertainty faced by a bidder due to asymmetric information in the auction game. Hence, it is important to distinguish which elements of equation (2.8) are known by all other firms at the time of bidding and which are private information to firm f submitting the bid vector. We use tilde (e.g. $\tilde{\delta}_{if}$) to denote factors that are private information to firm f and therefore treated as random parameters from the perspective of all other bidders. As a consequence, the bid distribution $H_f(\cdot|Z_f)$ is characterized by the deterministic parameters $\{\beta_k^{scale}, \beta_k^{density}\}_{f \in F}$ (to be defined shortly) and the *distribution* of the random parameters $\{\tilde{\delta}_{if}\}_{i \in U, f \in F}$ and $\{\tilde{\varepsilon}_{af}\}_{a \in \mathcal{A}, f \in F}$. This distinction between deterministic and random parameters in equation (2.8) is important for simulating winning probabilities. Next, we provide more details on

how these different components are specified and estimated.

Model Specification and Estimation Method. First, consider the terms capturing scale and density discounts, $(\beta_{k(f)}^{scale}, \beta_{k(f)}^{density})$. The model allows for some observed heterogeneity of these discounts across firms, with $k(f)$ indicating the type of firm f . For example, firms could be categorized based on their business size, because larger firms may operate at a different cost scale and therefore their synergies could be different. Moreover, larger firms may also be able to bid on larger packages, so their markup adjustments could also be different. We assume that the heterogeneity in the discount curves across firms is considered common knowledge and that all the uncertainty associated with the magnitude of the discounts is provided by the error terms $\tilde{\epsilon}_{af}$. In fact, as discussed in Section 2.2.3, through small experiments we also found that modeling these discount parameters as random variables does not affect the cost estimates by much, given the magnitude of discounts observed in the data in our application.

To measure scale discounts, g^{scale} is specified as a step function of the package volume v_a . Because density discounts depend on the proximity of the units in the package, $g^{density}$ depends on the volume of *clusters* of units in a package, where a cluster is a subset of the units in package a which are located in close proximity. In equation (2.8), $Cl(a)$ denotes the set of clusters in the package and c indicates a given cluster in this set, with size v_c . This approach follows directly from the work of Olivares et al. (2012), and further details on a specific way of computing clusters used in our application is described in the appendix of that article.

Consider now the term $\sum_{i \in a} \frac{v_i}{v_a} \tilde{\delta}_{if}$, a weighted average of firm-unit specific random parameters that capture the effects of the individual units contained in package a . The $\tilde{\delta}_{if}$'s are average implicit prices that bidder f charges for each unit among all the packages submitted, net of any discounts. These implicit prices capture heterogeneity in the unit characteristics (e.g. urban vs. rural units) and local advantages of a firm in that unit, among other factors. Part of the heterogeneity of this implicit prices is considered to be private information. Accordingly, we let the vector of average implicit prices $\tilde{\delta}_f$ follow a multi-variate normal distribution with mean and covariance matrix (μ, Σ) . More

specifically, let:

$$\tilde{\delta}_{if} = \bar{\delta}_i + \beta^Z Z_{if} + \tilde{\psi}_{r(i),f} + \tilde{\nu}_{if}, \quad (2.9)$$

so that $\mu_i = E(\tilde{\delta}_{if}) = \bar{\delta}_i + \beta^Z Z_{if}$ is specified by a unit fixed effect and the firm characteristics Z_{if} . Firm characteristics depend on the specific application, but may include an indicator on whether the firm was awarded the unit in the previous auction and other covariates that capture local advantages of the firm. The error terms $(\tilde{\psi}_{r(i),f}, \tilde{\nu}_{if})$ impose restrictions on the covariance matrix Σ based on the spatial location of units. Let \mathcal{R} be a set of regions that cover all the units in U and $r(i)$ denote the region that contains unit i ; the number of regions, R , is smaller than the number of units. Each firm is associated with a realization of the random vector $\tilde{\psi}_f = (\tilde{\psi}_{1f}, \dots, \tilde{\psi}_{Rf})$ from a multivariate normal distribution with zero mean and covariance matrix Ω . The error term $\tilde{\nu}_{if}$ follows an independent, heteroscedastic, zero-mean normal distribution with variance σ_i^2 .

Under the specification (2.9), the covariance structure of any two average implicit prices $\tilde{\delta}_{if}$ and $\tilde{\delta}_{jf}$ is given by $\text{Cov}(\tilde{\delta}_{if}, \tilde{\delta}_{jf}) = \Omega_{r(i),r(j)} + \sigma_i \sigma_j \mathbf{1}[i = j]$. Thus, under this model, two unit prices will be more positively correlated if the regional effects of the corresponding regions are more positively correlated. Note that this specification imposes positive correlation among unit prices in the same region; this restriction can be validated with data from the specific application. The model is flexible in allowing positive or negative correlation among units in different regions. Because R may be much smaller than the number of units, this specification provides a substantial dimensionality reduction over the fully flexible covariance matrix Σ .

In summary, the competitor's bid distribution $H_f(\cdot | Z_f; \phi)$, where ϕ is the set of parameters that fully describe the distribution, is a mixture defined by equations (2.8) and (2.9), $\tilde{\psi}_f \sim MVN(0, \Omega)$, $\tilde{\nu}_{if} \sim N(0, \sigma_i)$, and the error $\tilde{\varepsilon}_{af}$ which is assumed to have a zero-mean normal distribution with variance dependent on the package size, $\sigma_{\varepsilon,|a|}^2$. We seek to estimate the vector parameters β^{scale} , $\beta^{density}$, $\bar{\delta} = (\bar{\delta}_1, \dots, \bar{\delta}_N)$, β^Z , $\sigma^2 = (\sigma_1^2, \dots, \sigma_N^2)$, the covariance matrix Ω , and $\{\sigma_{\varepsilon,|a|}^2\}$ for different package sizes. The following two-step method is used to estimate these parameters:

- First step: Estimate (2.8) via a Generalized Least Squares (GLS) regression to obtain es-

imates of β^{scale} , $\beta^{density}$, $\{\sigma_{\varepsilon,|a|}^2\}$, and point estimates of the realizations of the average implicit prices $\tilde{\delta}_{if}$'s.

- Second step: Replace the estimated $\tilde{\delta}_{if}$'s into equation (2.9), and estimate the parameters characterizing its multivariate-normal distribution through maximum likelihood.

Identification of the parameters is based on variation across units and firms within a single auction. More specifically, the estimation of the scale and density discounts uses variation across different combinations submitted by the same firm over the same set of units. Given consistent estimates of the realized implicit average unit prices $\tilde{\delta}_{if}$, the second step provides consistent estimates of $\{\bar{\delta}_i, \sigma_i\}_{i \in U}$, β^Z , and Ω as long as Z_{if} is orthogonal to the error components $\tilde{\psi}_{r(i),f}$ and $\tilde{\nu}_{if}$. The consistency of our two-step method is a special case of the 2-step M-estimators described in Wooldridge (2002).

2.5.2 Markup and Cost Estimation

Similar to the approach discussed in 2.2.2, under the characteristic-based markup model developed in Section 2.3, markups are estimated using equation (2.6) in two steps. In the first step, the distribution of the competitors' bids are estimated, which are described in the previous section. Then in the second step, given the estimated distribution of competitors' bids, the aggregated winning probabilities $W^T G(b)$ and its Jacobian $\mathcal{D}_\theta W^T G(b)$ are computed. In this section we describe a simulation-based approach to estimate these quantities under the characteristic-based markup model.

Specifically, given the point estimates $\hat{\phi}$ for the distribution parameters estimated in the first step, each simulation run r consists of the following:

1. For each competitor f' (different from bidder f), draw independently a bid vector $b_{f'}^{(r)}$ (containing all submitted packages by that firm) from the estimated bid distribution $H(\cdot | Z_{f'}; \hat{\phi})$.
2. Using the observed bid b_f for bidder f and the sampled competitors' bids $\{b_{f'}^{(r)}\}_{f' \neq f}$, solve the winner determination problem.

3. Let $\iota^{(r)}$ be a vector of A binary variables indicating the packages awarded to the bidder f .

Store in memory the vector $W^T \iota^{(r)}$.

At the end of the simulation after R replications, the aggregated winning probability vector can be estimated by:

$$W^T G(b) \approx \frac{1}{R} \sum_{r=1}^R W^T \iota^{(r)}.$$

Note that if the distribution of competitors' bids is estimated consistently, then the previous equation provides consistent estimates of the aggregated winning probabilities as R becomes large. For the computation of the Jacobian matrix $\mathcal{D}_\theta W^T G(b)$, a central finite difference method can be used, which requires calculating the change in the winning probabilities from a small change in each markup variable of θ . Because the bid is linear in θ (by Assumption 2.5), this is equivalent to considering a small change in the observed bid vector b in the direction of each markup variable. Specifically, consider a change in the j^{th} component of the markup vector, and let W_i be the i^{th} column of the package-characteristic matrix W . Then the following central finite difference equation can be used to calculate the i^{th} row and j^{th} column element of the Jacobian:

$$[\mathcal{D}_\theta W^T G(b)]_{ij} = \frac{\partial W_i^T G(b)}{\partial \theta_j} \approx \frac{W_i^T G(b + hW_j) - W_i^T G(b - hW_j)}{2h}.$$

Notice that since there are d number of markup parameters, we need to take d upward perturbations ($b + hW_j$, for $j = 1, \dots, d$) as well as d downward perturbations ($b - hW_j$, for $j = 1, \dots, d$). Then the computation of $W_i^T G(b + hW_j)$ and $W_i^T G(b - hW_j)$ is done via simulation as before: in each simulation run, we solve the winner determination problems with each of the perturbed bid vectors and keep track of the winning bids. Note that this means we need to solve the winner determination problem $(2d + 1)$ times in each simulation run (d with upward perturbations, d with downward perturbations, and one with unperturbed bids). This could cause significant computational challenge in large-scale CAs – as the size increases solving one winner determination problem becomes computationally more expensive, and the number of required perturbations also increase. We will discuss this issue in more detail in the following section and propose an algorithm to achieve a more efficient computation than this brute forth approach.

Once the aggregated winning probability vector $W^T G(b)$ and its Jacobian matrix $\mathcal{D}_\theta W^T G(b)$ are estimated, the markup vector θ for this bidder is obtained through the identification equation (2.6): $\theta = - \{[\mathcal{D}_\theta W^T G(b)]^T\}^{-1} W^T G(b)$. The winning records used in the central finite difference method also enable to estimate the direct second-order derivatives of the bidders' expected profits with respect to each of their markup variables. We obtained that these estimates are negative for all firms, which is consistent with the local optimality of the estimated markups.

2.5.3 Improving the Computational Efficiency

In the proceeding section, we described simulation-based approach for the second step estimation under the characteristic-based markup model. As noted earlier, in each simulation run one needs to solve the winner determination problems $(2d + 1)$ times in total, while the sampled bid prices of the competitors are kept fixed. This is because using common random numbers is known to improve the estimation quality in this type of settings (see, for example, Glynn (1989)). As the size of the auction increases, however, this task can become computationally intense. In this section, we propose an algorithm that can alleviate the computational burden in this step. Specifically, we examine whether we can reduce the number of winner determination problems one needs to solve in a given simulation run.

Once the competitors' bid prices are sampled at the start of a simulation run, one may begin with solving the unperturbed problem and saving its solutions. Then the rest $2d$ of perturbed problems are considered. In each perturbed problems, we set the initial solution to be the optimal solution of the unperturbed problem. We note that when the perturbation size is small, the perturbed bid vector may not affect the optimality of the initial solution. In our proposed algorithm, we check the optimality of the initial solution first, then solve the winner determination problem only if the initial solution is declared not to be optimal. Otherwise, we can skip solving the winner determination problem, using the initial solution as the optimal solution for the particular perturbed problem. Since solving the winner determination problem usually is the bottleneck of the computational process in the second step, if we don't need to solve the problem for many of the perturbed bids, that will save

computation time in great deal.

We focus on the group markup model where any bid is associated with only one mark-up parameter. Recall that in our structural model described in Section 2.2.1, we have an allocative constraint in the winner determination problem that ensures that any winner can win at most one bid. Therefore, the focal bidder f has at most one winning bid in the optimal solution. We will call the mark-up parameter that is associated with the focal bidder's winning bid to be a *winning markup* parameter and call the rest *non-winning markup* parameters. We now examine the winner determination problems in the following two types of perturbations.

First, consider the upward perturbations – making some of my bids more expensive. Note that the perturbation will change the value of feasible allocations that includes allocating one of the perturbed bids to the focal bidder. Therefore an upward perturbation will result in increased total value of the feasible solutions that include the bids associated with the perturbed markup parameter. For this reason, the optimal solution may change by an upward perturbation only if we perturb the winning mark-up parameter. Moreover, increasing bids that is not a part of current optimal solution (and therefore the winning mark-up) will never change the optimal solution after perturbation. Since there is at most one winning bid in any given solution, during the d number of upward perturbation, at least $(d - 1)$ upward perturbations will not affect optimality of the current optimal solution, and therefore we can skip solving for the winner determination problem.

We now turn our attention to the case of downward perturbations. A downward perturbation of a non-winning markup may result in solution changes, since it decreases the total values of the feasible allocations that involve the bids associated with the perturbed markup parameter. Also a downward perturbation of a winning markup parameter may change the solution as well because the decrease in each of the associated bids by this downward markup perturbation depends on the volume of the package. Hence for downward perturbations, we conduct a sensitivity analysis to examine the possibility of reducing the number of solving winner determination problems.

The main idea of the sensitivity analysis is to find a second-best solution at extra cost of computation – which is might be compensated later – and use that information to detect the cases

where solving the winner determination problem is unnecessary. The second-best solution is an optimal solution with following two additional constraints: (1) one of the focal bidder's bids must win; and (2) current optimal solution is not feasible. The second-best solution found in this way is the best possible solution where the focal bidder wins a bid other than the current optimal solution. As will be clear below, we are interested in the case where the downward perturbation makes the focal bidder a winner from a non-winner or the case where the focal bidder wins a bid other than her current winning bid after the perturbation. In both cases, the new optimal solution is the second-best solution that we have already computed. Therefore, the difference in the objective values between the actual optimal solution and second-best solution will be used as a threshold that a perturbation should overcome in order to affect the optimality of the initial solution. We explain further details in what follows.

To proceed, we will use the following notations. First, we let P_i be the set of package indexes corresponding to the bids associated with i_{th} markup parameter. Denoted by v_p is the volume of package p . We will also let p^* is the package index of the focal bidder's winning bid in the initial solution, which is the optimal solution of the unperturbed problem, and let γ to denote the aforementioned threshold defined by "the objective value of the optimal solution minus the objective value of the second-best solution. Finally, define $V_i := \max_{p \in P_i} v_p$, the maximum volume of the packages in P_i . Note that given a per-meal perturbation size h , the decrease in a perturbed bid is obtained by the volume of the package multiplied by h . Therefore, when the i_{th} markup parameter is perturbed, the decreases in the associated bid prices are at most $h \cdot V_i$. Now we consider the following four cases.

Case 1: a downward perturbation of non-winning markup i In this case, the perturbation can decrease the objective value of the feasible allocations that include perturbed bids, which are not part of the winning bids in the optimal solution. Therefore larger than the threshold the perturbation may make such an allocation an optimal solution. In other words, the perturbation will not change the optimal solution if the maximum possible decrease in the perturbed bid prices is less than the

threshold. That is, one can skip the winner determination problem under the condition, $h \cdot V_i < \gamma$.

Case 2: a downward perturbation of winning markup i Now the bid prices associated with the winning markup will be all decreased by the perturbation. However, the magnitude may differ depending on the volume of the packages. If the relative price drop of the currently winning bid is sufficiently smaller than the price decrease of the other bids that are associated with the winning markup, the initial solution may not remain optimal. Otherwise, the initial solution will remain optimal, and one can skip the winner determination problem. Since the price drop in the winning package p^* will be $h \cdot v_{p^*}$, the maximum possible relative price drop of bids in P_i is $h \cdot (V_i - v_{p^*})$. Therefore solving the winner determination problem is unnecessary if the relative price change is less than the threshold, that is $h \cdot (V_i - v_{p^*}) < \gamma$.

In summary, for an upward perturbation, we can skip solving the winner determination problem if the perturbation is with respect to a non-winning markup parameter. For a downward perturbation, we can apply the above inequalities to determine the winner determination problem is necessary. Algorithm 1 describes the procedure in the l_{th} simulation run. We let $S_0^{(l)}$ be the vector of the optimal solution in the l_{th} simulation run for the unperturbed problem. Similarly, we let $S_{i+}^{(l)}$ and $S_{i-}^{(l)}$ be the optimal solution vectors for the i_{th} upward and downward perturbed problems, respectively, obtained in the l_{th} simulation run.

In Chapter 3 we apply our method described in this Chapter to an empirical application, where we have about 30 units in a CA. In that setting, it takes in the order of seconds to solve an instance of the winner determination problem, which is a significant bottleneck of the computational procedure. The approach developed in this section effectively reduced the computational burden – on average it took about 40% of computational time compared to the brute forth method.

Algorithm 1 Simulation Run l

- 1: For each bidder $f' \neq f$, sample $b_{f'}^{(l)}$ from $H(\cdot | Z_{f'}; \hat{\phi})$;
 - 2: Solve the winner determination problem (WDP) with $\{b_{f'}^{(l)}\}_{f' \neq f}$ and b_f ;
 - 3: Store the solution in the vector $S_0^{(l)}$, and store winning markup index w^* ;
 - 4: **for** $i = 1$ to d (*upward perturbation cycle*) **do**
 - 5: **if** i is a winning markup **then**
 - 6: Set $S_0^{(l)}$ as an initial solution and solve the WDP with $\{b_{f'}^{(l)}\}_{f' \neq f}$ and $(b_f + hW_i)$;
 - 7: Store the solution in the vector $S_{i+}^{(l)}$;
 - 8: **else**
 - 9: Set $S_{i+}^{(l)} \leftarrow S_0^{(l)}$;
 - 10: **end if**
 - 11: **end for**
 - 12: Find the second-best solution with $\{b_{f'}^{(l)}\}_{f' \neq f}$ and b_f , and set γ ;
 - 13: **for** $i = 1$ to d (*downward perturbation cycle*) **do**
 - 14: **if** $(i = w^*$ and $h \cdot (V_i - v_{p^*}) < \gamma$) or $(i \neq w^*$ and $h \cdot V_i < \gamma)$ **then**
 - 15: Set $S_{i-}^{(l)} \leftarrow S_0^{(l)}$;
 - 16: **else**
 - 17: Set $S_0^{(l)}$ as an initial solution and solve the WDP with $\{b_{f'}^{(l)}\}_{f' \neq f}$ and $(b_f - hW_i)$;
 - 18: Store the solution in the vector $S_{i-}^{(l)}$;
 - 19: **end if**
 - 20: **end for**
-

2.6 Conclusion

In this chapter, we develop a structural estimation approach for large-scale first-price CAs. An important methodological contribution of our work is to introduce a restricted markup model in which bidders are assumed to determine their markups based on a reduced set of package characteristics. The main advantage of this approach is that it reduces the computational burden of the structural

approach so that it can be applied to large-scale CAs. We establish markup restrictions that are parsimonious yet sufficiently flexible to capture strategic markup adjustments that can undermine the performance of CAs. We expect that our approach, based on pricing package characteristics, can be a useful starting point to reduce the complexity of econometric analysis in other real-world large-scale auctions.

Moreover, we also propose a parametric approach to model the competitors' bid distribution that needs to be estimated in the first step, addressing the complexity of estimation in large-scale CAs. We believe that our model provides a parsimonious, yet flexible parametric description of the distributions for the competitors' bids which can be especially useful in CAs that involve geographically dispersed and heterogeneous units that are subject to discounts due to scale and density.

2.7 Appendix for Chapter 2 - Proofs

2.7.1 Notation

We begin by defining notation that is frequently used in this section. First, we consider a focal bidder f , whose observed bid vector is denoted by b . All of the analysis is focused on this particular bidder, and as before we omit the firm index f whenever it is clear from the context. Recall that from the perspective of this focal bidder, competitors' bid prices are random. All such random quantities are defined over a probability space $(\Omega, \mathcal{F}, \mathbf{P})$. Note that \mathbf{P} measures the probability of each of the events characterized by the allocation of units to bidders in the CA. Hence, it defines the vector of winning probabilities $G(\cdot)$. In addition, we define $\Omega^* \subseteq \Omega$ to be the sample space where ties never happen in the winner determination problem. By Assumption 2.4, the distribution of competitors' bids is absolutely continuous, and hence we can find such a sample space so that $\mathbf{P}(\Omega^*) = 1$. In words, this means that the winner determination problem has a *unique* solution for any realization $\omega \in \Omega^*$. Accordingly, in our analysis we do not consider any issues related to tie-breaking in determining the auction allocation.

We let b' be the vector of competitors' bid prices. That is, given a realization of $\omega \in \Omega^*$, $b'(\omega) = \{b'_{f'}(\omega)\}_{f' \neq f}$, where $b'_{f'}(\omega)$ is a vector of bids for competing firm f' . Furthermore, we let $x = \{x_{af}\}_{a \in \mathcal{A}, f \in F}$ be a $A \times |F|$ dimensional vector such that x_{af} takes 1 if bidder f wins package a and 0 otherwise. A vector x uniquely determines an allocation outcome. We denote by X , the set of all feasible allocation outcomes that satisfy all the allocative constraints in the CA including the one that each bidder can win at most one package (see Assumption 2.1). All the proofs in this section are valid under any additional allocative constraints in the CA as long as they do not depend on bid prices (so the constraints in our empirical application described in Sections 3.2.2 and 3.2.3 are all valid). In addition, we let $X_a \subset X$ be the set of allocations such that bidder f wins package a . We adopt the *null* package, indexed by 0, and accordingly, we use X_0 to denote the set of allocations in which bidder f wins no package. We also let $G_0(b)$ be the probability that bidder f wins no package given her bid vector b . Note that because bidders can win at most one package, X_a and X_s

are disjoint for any packages $a \neq s$ in $\mathcal{A}_0 = \mathcal{A} \cup \{0\}$, and we have $\bigcup_{a \in \mathcal{A}_0} X_a = X$.

Without loss of generality, we assume x is ordered in a way that the vector of bidder f 's allocation decisions, denoted by x_f , is followed by the vector of competitors' allocation decisions, denoted by x' , so that $x = (x_f, x')$. Additionally, we define a cost function: $p_a(\omega) := \min_{x \in X_a} (b, b'(\omega))^T x$, for each $a \in \mathcal{A}_0$. This is the minimum total procurement cost out of all the allocations where bidder f wins package a given a realization $\omega \in \Omega^*$. It is important to note that because each bidder can win at most one package, for any $a \in \mathcal{A}$, $p_a(\omega)$ only depends on the value of b_a among bidder f 's bids in b .

Finally, for notational simplicity, we use $G_{a,s}(b)$ to denote the partial derivative of the winning probability $G_a(b)$ with respect to the bid price b_s . Similarly, when dealing with a characteristic-based markup model, we let $\mathcal{A}_i \subseteq \mathcal{A}$ to denote the set of packages associated with the i^{th} markup variable θ_i and let $G_{a,\theta_i}(b)$ to denote the partial derivative of the winning probability $G_a(b)$ with respect to the markup variable θ_i .

2.7.2 Proofs of Main Results

We will use some side-lemmas for the proofs of the main results. The proofs of these side-lemmas are provided in Section 2.7.3, if not from a reference. We start with the following lemma that is useful for the proof of Proposition 2.1.

Lemma 2.7.1. *Define a function $F : \mathbb{R}^n \mapsto \mathbb{R}$ such that:*

$$F(y) = \int_{D(y)} f(x) dx,$$

where $f : \mathbb{R}^m \mapsto \mathbb{R}$ is continuous and integrable in \mathbb{R}^m . Assume that the domain of integration $D(y)$ is a polyhedron formed by a given matrix $A \in \mathbb{R}^{k \times m}$ and a vector function $b(y) \in \mathbb{R}^k$ with $k \in \mathbb{N}$ such that $D(y) := \{x \in \mathbb{R}^m : Ax \leq b(y)\}$. If $b(y)$ is differentiable with respect to y , then F is continuous and differentiable everywhere in \mathbb{R}^n .

Proof of Proposition 2.1. To prove the differentiability of the winning probability vector $G(b)$ with respect to b , first fix an arbitrary package $a \in \mathcal{A}$ and look at the winning probability that bidder f

wins package a , $G_a(b)$. Notice that bidder f wins package a if one of the allocations in X_a achieves the minimum procurement cost among all possible allocations in X . We let $K := |X_a|$, the number of distinct allocations in X_a , and index them by $k = 1, 2, \dots, K$. Now we specifically consider the event that bidder f wins package a as a result of allocation $x_k \in X_a$. Accordingly, we let $G_a(b; x_k)$ denote the probability that $x_k \in X_a$ becomes the final allocation (hence the minimizer of the total procurement cost). Because the probability of ties is zero, the winning probability $G_a(b)$ can be expressed as $G_a(b) = \sum_{k=1}^K G_a(b; x_k)$. Therefore it suffices to show that $G_a(b; x_k)$ is continuous and differentiable for any given allocation x_k .

Now given an arbitrary allocation $x_k \in X_a$, we show the differentiability of $G_a(b; x_k)$ using Lemma 2.7.1. By letting $h(b')$ denote the joint probability density function of competitors' bids b' , $G_a(b; x_k)$ can be written as $G_a(b; x_k) = \int_{D_k(b)} h(b') db'$, where $D_k(b)$ is the set of b' 's for which x_k is the optimal allocation given b . Observe that $D_k(b)$ can be expressed by the following set of inequalities:

$$x_k^T(b, b') \leq y^T(b, b'), \quad \forall y \in X \quad (\Rightarrow) \quad (x'_k - y')^T b' \leq (y_f - x_{kf})^T b, \quad \forall y \in X.$$

The inequalities ensure that, given the placed bids (b, b') , the total procurement cost incurred by allocation x_k is cheaper than those of any other feasible allocations if we do not consider ties. Therefore, we get: $D_k(b) = \{b' \in \mathfrak{R}^{A \times (|F|-1)} : (x'_k - y')^T b' \leq (y_f - x_{kf})^T b, \forall y \in X\}$. If we let $J = |X|$ and index the feasible allocations by j , then $D_k(b)$ is a polyhedron in $\mathfrak{R}^{A \times (|F|-1)}$ defined by $Mb' \leq q(b)$, where the j^{th} row of M is $(x'_k - y'_j)^T$ and the j^{th} element of vector $q(b)$ is $(y_{jf} - x_{kf})^T b$, for $j = 1, \dots, J$.

By Assumption 2.4, the density $H_{f'}(\cdot|Z)$ for each competitor f' is continuous everywhere and independent across bidders, and hence, the joint density $h(b')$ is continuous on $\mathfrak{R}^{A \times (F-1)}$. The integrability of $h(b')$ is readily obtained as it is a probability density function. Finally, the function $q(b)$ is a linear function of b , and hence differentiable with respect to b . Therefore, by Lemma 2.7.1, $G_a(b; x_k)$ is continuous and differentiable with respect to the bid vector b . Since the choice of package $a \in \mathcal{A}$ and allocation $x_k \in X_a$ was arbitrary, the proof is complete. ■

It is useful to examine some of the properties of the Jacobian matrixes $\mathcal{D}_b G(b)$ and $\mathcal{D}_\theta W^T G(b)$ for the proof of Proposition 2.2 and Theorem 2.1. The following lemma investigates those properties.

Lemma 2.7.2. *For any given bidder and her bid vector b , we have the following properties for the winning probability vector $G(b)$.*

1. *The Jacobian matrix $\mathcal{D}_b G(b)$ is symmetric.*
2. *For any package a , we have i) $G_{a,a}(b) \leq 0$; ii) $G_{s,a}(b) \geq 0$ for any $s \neq a$; and iii) $\sum_{s \in A} G_{s,a}(b) \leq 0$.*
3. *Consider a group-based markup model specified by a package-characteristic matrix W whose elements are all non-negative. Let the markup vector θ and $D := \mathcal{D}_\theta W^T G(b)$. Then $D_{ij} \geq 0$ for any $i \neq j$.*

Proof of Proposition 2.2. (Part a): In the full-dimension markup model, we have $b_a = c_a + \theta_a$ for $a = 1, \dots, A$, and the first-order conditions, (2.3) yields:

$$[\mathcal{D}_b G(b)]^T \theta = -G(b), \text{ where } \theta := [\theta_1, \dots, \theta_A]^T. \quad (2.10)$$

Similarly, for the group-based markup model, we have $b_a = c_a + \theta_u$ for all $a = 1, \dots, A$. Note that the package-characteristic matrix $W \in \mathfrak{R}^A$ is then $W = [1, 1, \dots, 1]^T$. By letting $\alpha := [\alpha_1, \dots, \alpha_A]^T$ where $\alpha_a := G_{a,\theta_u}(b)$, we have $\mathcal{D}_{\theta_u} W^T G(b) = W^T \mathcal{D}_{\theta_u} G(b) = W^T \alpha$. Then the first-order condition of this characteristic-based markup model, (2.5) now becomes:

$$[\mathcal{D}_{\theta_u} W^T G(b)]^T \theta_u = -W^T G(b) \quad (\Rightarrow) \quad \alpha^T W \theta_u = -W^T G(b). \quad (2.11)$$

Observe that by definition, $\frac{\partial b_s}{\partial \theta_u} = 1$ for all $s = 1, \dots, A$. Therefore by the chain rule, we get:

$$\alpha_a = G_{a,\theta_u}(b) = \sum_{s=1}^A G_{a,s}(b) \quad (\Rightarrow) \quad W^T [\mathcal{D}_b G(b)]^T = \alpha^T.$$

Using this, left-multiplying by W^T on both sides of equation (2.10) and then equating the right-hand sides of equations (2.10) and (2.11) yields:

$$\sum_{a=1}^A \alpha_a \theta_a = \left(\sum_{a=1}^A \alpha_a \right) \theta_u \quad (\Rightarrow) \quad \theta_u = \frac{1}{\sum_{a=1}^A \alpha_a} \sum_{a=1}^A \alpha_a \theta_a. \quad (2.12)$$

Note that by symmetry of the Jacobian matrix $\mathcal{D}_b G(b)$, shown in part 1 of Lemma 2.7.2, we have $\alpha_a = \sum_{s=1}^A G_{a,s} = \sum_{s=1}^A G_{s,a}$. Then part 2 of the same lemma implies $\alpha_a \leq 0$ for all $a = 1, \dots, A$. Next we show that at least one $\alpha_a < 0$, which then implies that $\sum_{a=1}^A \alpha_a < 0$. Assume for the purpose of contradiction that α_a 's are all zero. This implies that the sum of all the column vectors in the Jacobian matrix $\mathcal{D}_b G(b)$ is a zero vector and therefore they are not linearly independent. However, since all the bids have strictly positive winning probabilities, the Jacobian matrix $\mathcal{D}_b G(b)$ is invertible as shown in Theorem 2.1, hence a contradiction. Therefore, we have at least one α_a that is strictly negative, and so does $\sum_{a=1}^A \alpha_a < 0$. By defining $\beta_a := \alpha_a / (\sum_{s=1}^A \alpha_s)$, we get $\beta_a \geq 0$ and $\sum_{a=1}^A \beta_a = 1$. Finally plugging them into equation (2.12), we get $\theta_u = \sum_{a=1}^A \beta_a \theta_a$, which completes the proof.

(Part b): The first-order conditions of the full-dimension model, (2.3) gives:

$$\begin{bmatrix} G_{1,1}(b) & G_{2,1}(b) & G_{12,1}(b) \\ G_{1,2}(b) & G_{2,2}(b) & G_{12,2}(b) \\ G_{1,12}(b) & G_{2,12}(b) & G_{12,12}(b) \end{bmatrix} \begin{bmatrix} \theta_1 \\ \theta_2 \\ \theta_{12} \end{bmatrix} = - \begin{bmatrix} G_1(b) \\ G_2(b) \\ G_{12}(b) \end{bmatrix} \quad (2.13)$$

Now consider the case where we use common markup θ_u for single unit bids and markup θ_v for the bundle of the two, so that the package-characteristic matrix W is formed as follows:

$$W = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{array}{l} \rightarrow \text{Unit 1: apply unit markup } \theta_u, \\ \rightarrow \text{Unit 2: apply unit markup } \theta_u, \\ \rightarrow \text{Package 12: apply package markup } \theta_v. \end{array}$$

Note that by the chain rule, $G_{a,\theta_u}(b) = G_{a,1}(b) + G_{a,2}(b)$, for $a = 1, 2, 12$. Hence, the first-order conditions (2.5) of this characteristic-based model yields:

$$\begin{bmatrix} G_{1,1}(b) + G_{1,2}(b) + G_{2,1}(b) + G_{2,2}(b) & G_{12,1}(b) + G_{12,2}(b) \\ G_{1,12}(b) + G_{2,12}(b) & G_{12,12}(b) \end{bmatrix} \begin{bmatrix} \theta_u \\ \theta_v \end{bmatrix} = - \begin{bmatrix} G_1(b) + G_2(b) \\ G_{12}(b) \end{bmatrix} \quad (2.14)$$

Left-multiplying by W^T on both sides of (2.13) and then equating the right-hand sides of equations

(2.13) and (2.14) give:

$$\theta_u = \beta\theta_1 + (1 - \beta)\theta_2,$$

$$\theta_v = \theta_{12} + \gamma(\theta_1 - \theta_2),$$

$$\text{where } \beta := \det^{-1} \{G_{12,12} (G_{1,1} + G_{1,2}) - (G_{12,1} + G_{12,2}) G_{1,12}\},$$

$$\gamma := \det^{-1} \{(G_{2,1} + G_{2,2}) G_{1,12} - (G_{1,1} + G_{1,2}) G_{2,12}\},$$

$$\det := (G_{1,1} + G_{2,1} + G_{1,2} + G_{2,2}) G_{12,12} - (G_{12,1} + G_{12,2}) (G_{1,12} + G_{2,12}).$$

Note that by Theorem 2.1, the Jacobian matrix in (2.14) is invertible and therefore its determinant, denoted by \det , is not zero. To show $\beta \geq 0$, first observe that \det is strictly positive, since $-(G_{1,1} + G_{2,1} + G_{1,2} + G_{2,2}) \geq (G_{1,12} + G_{2,12}) \geq 0$ and $-G_{12,12} \geq (G_{12,1} + G_{12,2}) \geq 0$ by Lemma 2.7.2. Similarly, the same lemma also implies $-G_{12,12} \geq (G_{12,1} + G_{12,2}) \geq 0$ and $-(G_{1,1} + G_{1,2}) \geq G_{1,12} \geq 0$. Therefore we get $\beta \geq 0$, which completes the proof. \blacksquare

Proof of Lemma 2.4.1. Fix a package $a \in \mathcal{A}$. Note that by the chain rule and Assumption 2.5, we have $G_{a,\theta_i}(b) = \sum_{s \in \mathcal{A}} \frac{\partial b_s}{\partial \theta_i} G_{a,s}(b) = \sum_{s \in \mathcal{A}} W_{si} G_{a,s}(b)$. Therefore, it suffices to show that $G_{a,s}(b) = 0$ for all $s \in \mathcal{A}$.

First, we let $\underline{p}(\omega) := \min_{t \in \mathcal{A}_0} p_t(\omega)$, the minimum procurement cost given $\omega \in \Omega^*$. Note that $G_a(b) = 0$ implies $p_a(\omega) > \underline{p}(\omega)$ in a set of $\Omega'_a \subseteq \Omega^*$, such that $\mathbf{P}(\Omega'_a) = 1$. Also, we let $e_a \in \mathfrak{R}^A$ be the a^{th} canonical vector, whose a^{th} component is equal to one while all others are equal to zero.

We now show that $G_{a,s}(b) = 0$ for all $s \in \mathcal{A} \setminus \{a\}$. First, take any package $s \neq a$ and consider a perturbation of decreasing b_s by $\epsilon > 0$. Recall that bidder f can win at most one package and therefore $p_a(\omega)$ does not depend on the value of b_s . Therefore, decreasing b_s will not change the value of $p_a(\omega)$. However, depending on whether b_s is part of the current optimal allocation or not, the value of the current optimal allocation may decrease by ϵ or stay the same ($\underline{p}(\omega)$) after the perturbation. Thus, after such a perturbation the value of the current allocation will still be lower than $p_a(\omega)$. This implies that bidder f remains not winning package a for all $\omega \in \Omega'_a$. Hence, we

obtain $G_a(b) - G_a(b - \epsilon e_s) = 0$ for all $\epsilon > 0$. Then the differentiability of $G(b)$ established in Proposition 2.1 implies $G_{a,s}(b) = 0$.

Similarly, to show that $G_{a,a}(b) = 0$, consider a perturbation of increasing b_a by $\epsilon > 0$. Then again for all $\omega \in \Omega'_a$, after such a perturbation, $p_a(\omega)$ only increases (to be $p_a(\omega) + \epsilon$) and remains being larger than the optimal value $\underline{p}(\omega)$. Hence bidder f can never win package a after the perturbation, which implies $G_a(b + \epsilon e_a) - G_a(b) = 0$ for all $\epsilon > 0$. Again by Proposition 2.1, we obtain $G_{a,a}(b) = 0$.

By combining these results, we get $G_{a,\theta_i}(b) = \sum_{s \in \mathcal{A}} W_{si} G_{a,s}(b) = 0$, which completes the proof. ■

The following Lemma provides invertibility conditions of a matrix, which is used to prove Theorem 2.1.

Lemma 2.7.3 (Theorem 6.1.10 in Horn and Johnson (1985)). *A matrix $D \in \Re^{n \times n}$ is said to be strictly diagonally dominant, if it satisfies:*

$$|D_{ii}| > \sum_{j \neq i} |D_{ij}|, \quad \forall i = 1, 2, \dots, n.$$

If D is strictly diagonally dominant, then D is invertible.

Proof of Theorem 2.1. (Necessity): We first show that if the Jacobian matrix $\mathcal{D}_\theta W^T G(b)$ is invertible, it must be that every markup variable has at least one relevant bid associated with it. For this, assume there exists a markup variable, say θ_i , whose associated bids are all irrelevant. Now note that $[\mathcal{D}_\theta W^T G(b)]_{ij} = \sum_{a \in \mathcal{A}_i} W_{ai} G_{a,\theta_j}(b)$. But then Lemma 2.4.1 implies that $G_{a,\theta_j}(b) = 0, \forall a \in \mathcal{A}_i$, leading to $[\mathcal{D}_\theta W^T G(b)]_{ij} = 0$. Since this is true for any $j = 1, 2, \dots, d$, the i^{th} row of Jacobian matrix $\mathcal{D}_\theta W^T G(b)$ will be a zero vector. Having a row of zeros implies that the matrix is not invertible. This completes the proof of necessity.

(Sufficiency): We now show that if every markup variable has at least one relevant bid associated with it and the additional conditions in the statement of the theorem hold, then the Jacobian matrix

$\mathcal{D}_\theta W^T G(b)$ evaluated at the observed bid vector b is invertible, and therefore the markup vector θ is uniquely determined by equation (2.6). For notational simplicity, we let $D := \mathcal{D}_\theta W^T G(b)$.

First, recall that in a group-based markup specification, for any package a , there is only one markup variable that is associated with it, say markup variable θ_i . Then the profit that bidder f makes from winning package a is $W_{ai}\theta_i$. By assumption, $W_{ai}\theta_i \geq 0$ and $W_{ai} \geq 0$, for all packages a . Therefore, $\theta_i \geq 0$, for all i . We now proceed to show that θ_i is indeed strictly positive for all $i = 1, 2, \dots, d$. By Assumption 2.6, θ satisfies equation (2.5): $D^T \theta = -W^T G(b)$. For the purpose of contradiction, we fix i and assume that θ_i is zero. We examine the i^{th} equation in (2.5):

$$D_{ii}\theta_i + \sum_{j \neq i} D_{ji}\theta_j = -W_i^T G(b). \quad (2.15)$$

The first term on the left-hand side is zero by assumption. The second term is non-negative since we know that (i) $\theta_j \geq 0, \forall j$; and (ii) $D_{ji} \geq 0$ by part 3 of Lemma 2.7.2. However, the right-hand side is strictly negative because there is at least one relevant bid, say b_a , that is associated with markup variable θ_i , so that $W_i^T G(b) \geq W_{ai}G_a(b) > 0$. Therefore it is impossible for θ to satisfy equation (2.5), which contradicts Assumption 2.6. Hence, $\theta_i > 0$, for all i .

Now, we construct a diagonal matrix Θ so that $\Theta_{ii} = \theta_i$ for all $i = 1, 2, \dots, d$. Because $\theta_i > 0, \forall i$, it is clear that Θ is invertible. We now show that equation (2.5) implies that the matrix $D^T \Theta$ is strictly diagonally dominant, and therefore invertible by Lemma 2.7.3. To see this, take any $i \in \{1, 2, \dots, d\}$, and consider the i^{th} equation in (2.5) (see (2.15)), for which we know that its right-hand side is strictly negative. Therefore, using $[D^T \Theta]_{ij} = D_{ji}\Theta_{jj} = D_{ji}\theta_j$, we reach the following inequality:

$$[D^T \Theta]_{ii} + \sum_{j \neq i} [D^T \Theta]_{ij} = -W_i^T G(b) < 0 \quad (\Rightarrow) \quad \sum_{j \neq i} [D^T \Theta]_{ij} < -[D^T \Theta]_{ii}.$$

Recall that when $i \neq j$, we have $[D^T \Theta]_{ij} = D_{ji}\theta_j \geq 0$, and this implies $\sum_{j \neq i} |[D^T \Theta]_{ij}| < |[D^T \Theta]_{ii}|$. Since this is true for any $i = 1, 2, \dots, d$, we conclude that $D^T \Theta$ is strictly diagonally dominant and hence invertible by Lemma 2.7.3. Since Θ is also invertible, the invertibility of D follows with $D^{-1} = (\Theta^T D)^{-1} \Theta^T$, and the proof for sufficiency is now complete. \blacksquare

2.7.3 Proofs of Side Lemmas

Proof of Lemma 2.7.1. We will prove differentiability of F by induction. Consider the basis case of $m = 1$. Since $D(y)$ is a polyhedron, it is a closed interval in \mathfrak{R} . Moreover, by assumption, each boundary is either a differentiable function of y or infinity. We let such boundary functions to be $\underline{h}(y), \bar{h}(y)$ for lower and upper bounds (possibly infinity), respectively. Then F can be rewritten as follows:

$$F(y) = \int_{\underline{h}(y)}^{\bar{h}(y)} f(x_1) dx_1.$$

Then, for any $j = 1, 2, \dots, n$, we have:

$$\frac{\partial F}{\partial y_j} = \frac{\partial \bar{h}(y)}{\partial y_j} f(\bar{h}(y)) - \frac{\partial \underline{h}(y)}{\partial y_j} f(\underline{h}(y)), \text{ by Leibniz Integral Rule.}$$

Note that each term will vanish in case of infinite boundary by the integrability of $f(\cdot)$. Thus, by the differentiability of the boundary function h and continuity of f , F is differentiable. Now set the induction hypothesis that it is true for $m - 1$. For the case with m , without loss of generality pick the first element in x and define F_1 as follows:

$$F_1(y, x_1) = \int_{D(y; x_1)} f(x_1, x_{-1}) dx_{-1},$$

where $x_{-1} := (x_2, \dots, x_m)$, A_i is the i^{th} column vector in A , $A_{-1} := [A_2, \dots, A_m]$, and $D(y; x_1) := \{\tilde{x}_{-1} \in \mathfrak{R}^{m-1} : A\tilde{x} \leq b(y), \tilde{x}_1 = x_1\} = \{\tilde{x}_{-1} \in \mathfrak{R}^m : A_{-1}\tilde{x}_{-1} \leq b(y) - A_1x_1\}$. It is clear that $D(y; x_1)$ is also a polyhedron in \mathfrak{R}^{m-1} defined by a given matrix A_{-1} and a differentiable vector $b(y) - A_{-1}x_1$. Hence, $F_1(y, x_1)$ is differentiable with respect to y for any given x_1 , by induction hypothesis. Now using F_1 , F can be expressed as:

$$F(y) = \int_{D_1(y)} F_1(y, x_1) dx_1$$

. Since $D_1(y)$ is a projection of $D(y)$ onto x_1 axis, it is also a polyhedron in \mathfrak{R} and thus a closed interval. Hence after letting these boundaries to be $\underline{h}(y), \bar{h}(y)$, and again by Leibniz Integral Rule, its partial derivative with respect to y_j becomes:

$$\frac{\partial F}{\partial y_j} = \frac{\partial \bar{h}(y)}{\partial y_j} f(\bar{h}(y)) - \frac{\partial \underline{h}(y)}{\partial y_j} f(\underline{h}(y)) + \int_{\underline{h}(y)}^{\bar{h}(y)} \frac{\partial}{\partial y_j} F_1(y, x_1) dx_1,$$

Therefore, $\frac{\partial F}{\partial y_j}$ is continuous for the case of m . Hence we conclude that F is differentiable for any dimension $m \in \mathbb{N}$ and the proof is complete. \blacksquare

Proof of Lemma 2.7.2. (Part 1): We fix two arbitrary but distinct packages a and s , and we first show that $G_{s,a}(b) \leq G_{a,s}(b)$. We then establish the reversed inequality by exchanging the two packages and using a symmetric argument. The arbitrary choice of the two packages a and s then provides the completion of the proof.

Accordingly, take any two distinct packages $a, s \in \mathcal{A}$ and an arbitrary scalar $\epsilon > 0$. We begin by defining the following events:

$$\begin{aligned}\Omega_a &:= \{\omega \in \Omega^* : p_a(\omega) = \min_{t \in \mathcal{A}_0} p_t(\omega)\}, \\ \Omega_{a,s} &:= \{\omega \in \Omega_a : p_s(\omega) = \min_{t \in \mathcal{A}_0 \setminus \{a\}} p_t(\omega)\}, \\ \Omega_{a,s}^\epsilon &:= \{\omega \in \Omega_a : p_s(\omega) < p_a(\omega) + \epsilon\}.\end{aligned}$$

By definition, Ω_a denotes the event where bidder f wins package a , and $\Omega_{a,s} \subset \Omega_a$ denotes the event where the minimum allocation without bidder f winning package a is the one with her winning package s . Also, $\Omega_{a,s}^\epsilon \subset \Omega_a$ is the event where the minimum allocation with bidder f winning package s is less than ϵ above from the current optimal value, $p_a(\omega)$. Finally, we let $\Omega_s \subset \Omega^*$ to be the event where bidder f wins package s . Note that Ω_a and Ω_s are disjoint.

We use the following random variables: $Y_s^{a \pm \epsilon}(\omega) := \mathbf{1}[p_s(\omega) = \min(\min_{t \in \mathcal{A}_0 \setminus \{a\}} p_t(\omega), p_a(\omega) \pm \epsilon)]$. The random variables $Y_s^{a \pm \epsilon}(\omega)$ indicate bidder f 's winning of package s when her bid b_a changes by $+\epsilon$ and $-\epsilon$, respectively. Similarly, we define $Y_s^0(\omega) := \mathbf{1}[p_s(\omega) = \min_{t \in \mathcal{A}_0} p_t(\omega)]$, that is, the indicator that the bidder wins package s given her bid price b at the realization of ω . Now we divide the event set Ω^* into the following four disjoint subsets and examine the values of the random variables $Y_s^{a+\epsilon}(\omega)$ and $Y_s^0(\omega)$.

1. $\forall \omega \in \Omega^* \setminus (\Omega_a \cup \Omega_s)$: Bidder f is winning neither a nor s , so $Y_s^0(\omega) = 0$. Moreover, increasing her bid b_a by ϵ will not let her win s , hence, $Y_s^{a+\epsilon}(\omega) = 0$.

2. $\forall \omega \in \Omega_s$: Bidder f is winning package s and increasing her bid on non-winning package a will not change her winning s . Thus, $Y_s^0(\omega) = Y_s^{a+\epsilon}(\omega) = 1$.
3. $\forall \omega \in \Omega_{a,s} \cap \Omega_{a,s}^\epsilon$: Bidder f is winning package a , so $Y_s^0(\omega) = 0$. Since $\omega \in \Omega_{a,s}^\epsilon$, after increasing b_a by ϵ , the value of the current optimal allocation $p_a(\omega) + \epsilon$ becomes larger than $p_s(\omega)$. But then, $\omega \in \Omega_{a,s}$ implies $p_s(\omega)$ becomes the lowest procurement cost after such a perturbation. Hence, $Y_s^{a+\epsilon}(\omega) = 1$.
4. $\forall \omega \in \Omega_a \setminus (\Omega_{a,s} \cap \Omega_{a,s}^\epsilon)$: Bidder f is winning package a , so $Y_s^0(\omega) = 0$. If $\omega \notin \Omega_{a,s}$, after increasing b_a by ϵ , $p_s(\omega)$ is not the lowest procurement cost. If $\omega \notin \Omega_{a,s}^\epsilon$, $p_s(\omega)$ is still larger than the value of the current allocation, $p_a(\omega) + \epsilon$, even after the perturbation. Hence, $Y_s^{a+\epsilon}(\omega) = 0$.

In words, $(\Omega_{a,s} \cap \Omega_{a,s}^\epsilon)$ is the only event in which bidder f 's winning status of package s changes by an ϵ increase in her bid b_a . Therefore, we obtain:

$$\frac{G_s(b + \epsilon e_a) - G_s(b)}{\epsilon} = \frac{1}{\epsilon} \mathbf{E}[Y_s^{a+\epsilon} - Y_s^0] = \frac{1}{\epsilon} \mathbf{P}(\Omega_{a,s} \cap \Omega_{a,s}^\epsilon), \quad (2.16)$$

where e_a is the a^{th} canonical vector whose a^{th} component is the only non-zero element and is equal to one.

Now we look at the effect of decreasing b_s by ϵ to the winning of package a . Similarly, we divide the event set Ω^* into the following three disjoint subsets and examine the values of random variables $Y_a^{s-\epsilon}(\omega)$ and $Y_a^0(\omega)$.

1. $\forall \omega \in \Omega^* \setminus (\Omega_a)$: Since bidder f is not winning package a , $Y_a^0(\omega) = 0$. Moreover, decreasing her bid b_s by ϵ will never let her win package a , hence, $Y_a^{s-\epsilon}(\omega) = 0$.
2. $\forall \omega \in \Omega_{a,s}^\epsilon$: Bidder f is winning package a , so $Y_a^0(\omega) = 1$. Since $\omega \in \Omega_{a,s}^\epsilon$, after decreasing b_s by ϵ , $p_s(\omega) - \epsilon$ becomes lower than the current optimal value, $p_a(\omega)$, so bidder f will win package s instead of a . Hence, $Y_a^{s-\epsilon}(\omega) = 0$.

3. $\forall \omega \in (\Omega_a \setminus \Omega_{a,s}^\epsilon)$: Bidder f is winning package a , so $Y_a^0(\omega) = 1$. Since $\omega \notin \Omega_{a,s}^\epsilon$, decreasing b_s by ϵ cannot make the value $p_s(\omega) - \epsilon$ cheaper than the current optimal value, $p_a(\omega)$. Hence, the previous optimal allocation will remain optimal and $Y_a^{s-\epsilon}(\omega) = 1$.

This time, $\Omega_{a,s}^\epsilon$ is the only case that bidder f 's winning status of package a is affected by an ϵ decrease in her bid b_s . Therefore, we get:

$$\frac{G_a(b) - G_a(b - \epsilon e_s)}{\epsilon} = \frac{1}{\epsilon} \mathbf{E}[Y_a^0 - Y_a^{s-\epsilon}] = \frac{1}{\epsilon} \mathbf{P}(\Omega_{a,s}^\epsilon). \quad (2.17)$$

Since $(\Omega_{a,s} \cap \Omega_{a,s}^\epsilon) \subseteq \Omega_{a,s}^\epsilon$, from (2.16) and (2.17) we get the following inequality:

$$\frac{G_s(b + \epsilon e_a) - G_s(b)}{\epsilon} = \frac{1}{\epsilon} \mathbf{P}(\Omega_{a,s} \cap \Omega_{a,s}^\epsilon) \leq \frac{1}{\epsilon} \mathbf{P}(\Omega_{a,s}^\epsilon) = \frac{G_a(b) - G_a(b - \epsilon e_s)}{\epsilon}.$$

Recall that ϵ is an arbitrary positive scalar and Proposition 2.1 ensures the differentiability of $G(b)$ with respect to b . Thus, by letting ϵ vanish, we get $G_{s,a} \leq G_{a,s}$.

In the previous argument, the only condition for the packages a and s is that they are distinct. Hence, a symmetric argument also holds true and we get $G_{s,a} \geq G_{a,s}$, and therefore $G_{s,a} = G_{a,s}$. Since the choice of a and s was arbitrary, we conclude that the Jacobian matrix is symmetric and this completes the proof of part 1.

(Part 2): To show $G_{a,a}(b) \leq 0$, fix a realization of $\omega \in \Omega^*$ and consider a perturbation of increasing bidder f 's bid price b_a by $\epsilon > 0$. If she currently wins package a , she may or may not win package a after the perturbation. However, if she currently does not win package a , i.e., $p_a(\omega)$ is not the lowest procurement cost, she cannot win package a after the perturbation since $p_a(\omega) + \epsilon$ remains being larger than the current optimal value. Since these are true for any $\omega \in \Omega^*$, increasing bid price b_a will never increase her chances of winning package a . Hence we get $G_a(b + \epsilon e_a) \leq G_a(b)$, for all $\epsilon > 0$. Then the differentiability of $G(b)$, shown in Proposition 2.1, implies $G_{a,a}(b) \leq 0$.

Similarly, for the proof of $G_{s,a}(b) \geq 0$ for any $s \neq a$, consider a perturbation of decreasing b_a by an arbitrary $\epsilon > 0$. Given a realization of $\omega \in \Omega^*$, if she currently wins package s (possibly the null package), she can either win package a instead of s or still win package s after the perturbation.

However, if she currently wins package a , she will win package a for sure after the perturbation. Therefore, decreasing her bid b_a only possibly decrease her chances of winning package s , and we get $G_s(b) \geq G_s(b - \epsilon e_a)$, for all $\epsilon > 0$. Again the differentiability of $G(b)$ then implies $G_{s,a}(b) \geq 0$.

Finally, since $\sum_{s \in \mathcal{A}} G_s(b) = 1 - G_0(b)$, so we get $\sum_{s \in \mathcal{A}} G_{s,a}(b) = -G_{0,a}(b) \leq 0$, where the last inequality follows because $G_{0,a}(b) \geq 0$ by a similar argument as above. This completes the proof of part 2.

(Part 3): Note that by Assumption 2.5, $b = W\theta + c$ and by the chain rule, we have $D := \mathcal{D}_\theta W^T G(b) = W^T \mathcal{D}_b G(b) W$. Then for any $i \neq j$ we get:

$$D_{ij} = \sum_{a,s \in \mathcal{A}} W_{ai} W_{sj} G_{a,s}(b) = \sum_{a \in \mathcal{A}_i, s \in \mathcal{A}_j} W_{ai} W_{sj} G_{a,s}(b),$$

where the second equality comes from the fact that $W_{ai} = 0$ if $a \notin \mathcal{A}_i$ by its definition. In addition, recall that in a group-based markup model, \mathcal{A}_i and \mathcal{A}_j are disjoint if $i \neq j$. Therefore by part 2 of this lemma shown above, $G_{a,s}(b) \geq 0$ for all $a \in \mathcal{A}_i$ and $s \in \mathcal{A}_j$. The non-negativity of the elements in W then ensures that $D_{ij} \geq 0$ for all $i \neq j$, which completes the proof of part 3. ■

Chapter 3

Application to the Chilean School-Meal Auction

3.1 Introduction

We effectively apply our structural estimation method developed in Chapter 2 to the Chilean school meals CA (see Epstein et al. (2002) for a detailed description of the auction). This application fits well within the class of large-scale CAs: each auction has about 30 units and firms submit hundreds or even thousands of bids. This CA has a single-round sealed-bid first-price format. The auction is used by the Chilean government to allocate contracts among private catering firms to provide breakfast and lunch for 2.5 million children daily in primary and secondary public schools during the school year. In a developing country where about 14 percent of children under the age of 18 live below the poverty line, many students depend on these free meals as a key source of nutrition. The CA, one of the largest state-run auctions in Chile, was used for the first time in 1999 and has been used every year since its inception awarding more than \$3 billion of contracts. Although this application has been praised for bringing transparency and lowering the procurement costs of a high social impact public service, a detailed performance analysis of the CA format has not been conducted. Moreover, government officials running this auction have considered revising

its format, in part worried by the potential negative effects that can arise due to strategic bundling (Olivares et al., 2012). Thus motivated, we use our structural approach to inform this auction design question.

Our results show that for the Chilean auction, cost synergies are significant, amounting up to 6% of the cost. Roughly 75% of the discounts observed in the bid data arise from cost synergies (the rest is due to strategic markup adjustments). In part due to this large cost synergies, the CA achieves a strikingly high efficiency, with an actual cost allocation only 1.5% higher than the minimum-cost allocation. The results also show that while economies of scale (mostly generated by volume discounts in input purchases) are larger than economies of density (arising from common logistics infrastructure used to supply nearby units), they are both important in the firms' operational cost synergies. Finally, the estimated markups are on average around 5%, suggesting that the CA induces a reasonable amount of competition among the suppliers. The level of markups coincides with anecdotal evidence provided by the Chilean government.

Once we estimate the cost structure we can also perform other useful counterfactuals. One important consideration the government has when running these auctions, which arises frequently in other settings with synergies, is how to promote diversification and competition among bidders. On one hand, if cost synergies are significant, it may be efficient and optimal in the short-run to allocate all units to one or few firms. On the other hand, this could depress competition in the bidders' market for future auctions, as inactive firms may find it hard to compete head-to-head with incumbents, increasing expected payments in the long-run. In the Chilean school meals auction, the government imposes market share restrictions for bidders in the CA to promote long-run competition. The cost estimates provided by the structural estimation can be used to evaluate the efficiency loss due to these constraints. We find that the efficiency loss is very small, around 1%. The main reason for this result is that cost synergies get practically exhausted at the point where the market share constraints become binding.

Our results on the Chilean school meals auction complement those in Olivares et al. (2012) that provide further suggestive evidence of cost synergies in this application. However, their approach

does not provide direct estimates of the suppliers' costs, which are needed to evaluate the performance of the auction. In addition, their approach does not allow us to perform counterfactuals. Other reduced-form approaches suffer from the same limitations (see Ausubel et al. (1997), Gandal (1997), and Moreton and Spiller (1998)) that our structural estimation approach overcome. Going back to our initial motivating auction design question, the structural approach provides a direct measure of the good performance of the CA in this setting, which cannot be done through a reduced-form approach. As far as we know, this is the first work that provides empirical evidence that a CA performs well in a real-world application. In summary, our results suggest that package bidding and running a CA seems appropriate in this application.

The rest of this chapter is structured as follows. In Section 3.2 we provide a detailed description of the Chilean school meals auction and the data. We also discuss how the structural model assumptions fit into this application. Section 3.3 presents the estimates for the distribution of the competitors' bids. Using these estimates, we then provide cost and markup estimation results in Section 3.4 in both small and large-scale auctions. We conduct performance analysis of the large-scale CAs using the cost and markup estimates in Section 3.5 and Section 3.6 provides the main conclusions for this chapter.

3.2 The Chilean Auction for School Meals

The application we study in this chapter is the Chilean auction for school meals. In this section, we provide a detailed description of the auction, the data collected. We also provide the justification of the assumptions of the structural model developed in Chapter 2 in the context of our application.

3.2.1 Brief History

Junta Nacional de Auxilio Escolar y Becas (JUNAEB) is a government agency in Chile that provides breakfast and lunch for 2.5 million children daily in primary and secondary public schools during the school year. This is one of the largest and most important social programs run by the Chilean

government. In fact, in a developing country where about 14 percent of children under the age of 18 live below the poverty line, many students depend on these free meals as a key source of nutrition.

Since 1999, JUNAEB assigns its school meal service contracts through a single-round, sealed-bid, first-price CA, that was fully implemented for the first time that year. The CA has been used every year since its inception awarding more than US\$ 3 billion of contracts, being one of the largest state auctions in Chile (in recent years, each auction awards contracts for about half a billion dollars). The auction process begins with the registration of potential suppliers followed by an evaluation conducted by the agency, which considers managerial, technical and financial performance metrics. Some companies may be excluded from the auction if they do not pass this evaluation. Meal plans and service quality are standardized, so that qualified suppliers compete on price.

For the purpose of the auction, Chile is divided into approximately 100 school districts or territorial units (TUs) in 13 geographic regions. Each year, JUNAEB holds an auction for one-third of the country (around 30 - 35 TUs), awarding three-year contracts. Typically about 20 firms participate in each auction and they are allowed to submit package bids that cover any combinations of TUs and specify the price to serve them. The maximum number of TUs that a firm is allowed to include in any given package (ranging from one to eight TUs) depends on the firm's financial evaluation. Vendors can submit many bids and each package bid is either fully accepted or rejected (i.e. the mechanism does not allocate a fraction of a bid); most firms submit hundreds or even thousands of bids. Two potential sources of cost synergies motivate the use of CAs in this context: (i) economies of scale, generated by volume discounts in the input purchases; and (ii) economies of density that arise from common logistics infrastructure used to supply nearby units.

3.2.2 Auction Process

The auction process begins when JUNAEB invites and registers potential vendors. The agency then evaluates the companies from a managerial, technical and financial point of view, and eliminates those that do not meet minimum reliability standards. Qualifying vendors are classified according to two characteristics: their financial capacity (based on data from the firms' balance sheets), and their

managerial competence. Usually, firms below a minimum level of managerial competence are not allowed to participate in the auction. Meal plans are standardized and service quality requirements are presented in detail. With that, firms compete on price basis. Potential vendors submit their bids simultaneously and in a single-round through an online system. Contract winners are paid the amount of their winning bids and are responsible for managing the entire supply chain associated with all meal services in the awarded TUs. This includes from sourcing food inputs going all the way to cooking and serving the meals in the schools.

Bidding language. A bid can cover any combination from one to eight TUs and specifies the price for which the firm would serve all meals included in the TUs in the combination. Vendors can submit many bids and each package bid is either fully accepted or rejected (i.e. the mechanism does not allocate a fraction of a bid); most firms submit hundreds or even thousands of bids.

Winner determination. The allocation is chosen by selecting the combination of bids that supply all of the TUs at a minimum cost. The problem is formulated as an integer program (IP) that incorporates other considerations and side constraints. There are four types of constraints implemented in the auction and the details of those constraints are as follows. We then provide the mathematical formulation of the IP in the following section.

1. **Cover all TUs:** the final allocation should cover all the TUs auctioned.
2. **Maximum Number of TUs:** There is a maximum number of TUs that each firm can be allocated in any given auction. This maximum is based on the financial evaluation conducted by JUNAEB every year and therefore can be different across firms and auctions, ranging from 1 to 8 TUs.
3. **Global Market Share Constraints:** To avoid excessive concentration and encourage diversification, at any point in time, the total standing contracts of any firm cannot exceed 16% of the total number of meals included in all TUs in the entire country. Hence, depending on

the volume of standing contracts, the maximum volume can be also different across firms and auctions.

4. **Local Constraints:** To facilitate supervision and control of the firms, there are constraints on the maximum number of firms serving in each geographical region. On the other hand, to actively respond to contingencies such as bankruptcies, there are also constraints on the minimum number of firms serving in each geographical region. Geographical regions in low population areas contain less than five TUs while regions with higher population typically contain between 10 and 20 TUs.
5. **Global Competition Constraint:** For similar reasons as the global market share constraints, there is a constraint in the minimum number of firms winning contracts in each auction (this number can vary across auctions, usually being around 10).

3.2.3 Winner Determination Problem Formulation

We now provide the details of the integer programming (IP) formulation of the winner determination problem (WDP). We begin by introducing notation that is not defined previously and we then formulate the IP.

Index Sets. We let \mathcal{R} denote the set of geographical regions indexed by r (recall that each geographical region contains several TUs). We let \mathcal{A}_f be the set of packages on which firm f places bids. They are to distinguish from \mathcal{A} in case of missing bids (unobserved bids) by firm f . $\mathcal{A}_{r,f} \subseteq \mathcal{A}_f$ represents the set of packages in \mathcal{A}_f that contain at least one TU in region r . Finally, we let $|a|$ denote the number of TUs in package a , and we let A_f and $A_{r,f}$ denote the number of packages in the sets \mathcal{A}_f and $\mathcal{A}_{r,f}$, respectively.

Constraints and Their Parameters. As described above, we have five types of allocative constraints in the auction. We also have an additional constraint imposed in our structural model, namely, that each firm can win *at most one* package. We label those constraints as follows: (A)

Cover all TUs ensures that all the TUs be contracted. (B) *At most one package* constraint imposes that firms can win at most one package. (C) *Maximum number of TUs* bounds the number of TUs that each firm can win. We let MXU_f denote the maximum number of TUs that firm f can win. (D) *Global Market Share Constraints* limits the total volume of standing contracts of each firm in terms of the number of meals served. We let MXM_f denote the total number of meals that firm f can win in the auction being considered. (E) *Local constraints* bound the minimum and maximum number of firms serving in each region. We use MNF_r and MXF_r to denote these bounds for region r . (F) *Global competition constraint* sets the minimum number of firms being contracted in the auction being considered. We let MNF_g denote this minimum number.

Decision Variables. We let x_{af} be the firm-package allocation decision variable for package a and firm f . This variable takes the value of 1, if firm f wins package a , and 0 otherwise. These variables determine the allocation. The variable y_{rf} is a regional allocation variable for region r and firm f , taking the value of 1 if firm f wins a package that contains at least one TU in region r , and 0 otherwise. They are used to count the number of firms serving in each geographical region for the local constraints. The decision variable z_f relates to the winning status of firm f . It is equal to 1 if firm f wins a package and 0 otherwise. They count the number of winning firms to be used in the global competition constraint.

IP Formulation of the WDP. First, notice that constraints (C) and (D) are firm-wise limits, and for each firm any bids placed on packages that exceed the firm's limits can never win. Therefore, we eliminate such bids *a priori* from \mathcal{A}_f for each firm $f \in F$. That is for any given firm f and for all $a \in \mathcal{A}_f$, we have $|a| \leq MXU_f$ and $v_a \leq MXM_f$. Then, constraints (C) and (D) will be automatically satisfied as long as firms win at most one package imposed by (B). Hence, we omit (C) and (D) in our IP formulation. Recall that the objective is to minimize the total procurement cost. Now we present the IP formulation of the WDP. The constraints that are not labeled impose the correct values for the auxiliary variables y_{rf} and z_f , and the integrality constraints for all decision

variables.

$$\begin{aligned}
& \text{minimize} && \sum_{f \in F} \sum_{a \in \mathcal{A}_f} b_{af} x_{af} \\
\text{subject to} & \text{(A)} && \sum_{f \in F} \sum_{a \in \mathcal{A}_f: i \in a} x_{af} \geq 1, \quad \forall i \in U \\
& \text{(B)} && \sum_{a \in \mathcal{A}_f} x_{af} \leq 1, \quad \forall f \in F \\
& \text{(E)} && MNF_r \leq \sum_{f \in F} y_{rf} \leq MXF_r, \quad \forall r \in \mathcal{R} \\
& && \frac{1}{A_{rf}} \sum_{a \in \mathcal{A}_{rf}} x_{af} \leq y_{rf} \leq \sum_{a \in \mathcal{A}_{rf}} x_{af}, \quad \forall r \in \mathcal{R}, \forall f \in F \\
& \text{(F)} && \sum_{f \in F} z_f \geq MNF_g, \\
& && \frac{1}{A_f} \sum_{a \in \mathcal{A}_f} x_{af} \leq z_f \leq \sum_{a \in \mathcal{A}_f} x_{af}, \quad \forall f \in F \\
& && x_{sf}, y_{rf}, z_f \in \{0, 1\}.
\end{aligned}$$

3.2.4 Description of the Data

Data were collected and processed for all auctions between 2002 and 2005. The dataset contains all bids placed by all firms in each auction, the identity and characteristics of participating firms in each auction, and detailed information on the auction parameters, including all the parameters used to determine the side constraints of the winner determination problem. TU data includes its annual demand (number of meals to be served), referred to as the *volume* of the TU, as well as the geographic location of its schools. We also know the set of winning bids in each auction and therefore, at every point in time, we know the identity of the firms serving each TU. Additional details of the data can be found in Olivares et al. (2012).

We apply our method to the large-scale CA of 2003. In 2002, the auction faced some regularity issues, and a second subsequent auction was used to award the contracts. Hence, we conservatively

Table 3.1 – Summary statistics for the 2003 auction.

Bidder Characteristics by Business Size						
Business Size	Small (1-2)	Medium (3-4)	Med-Large (5-6)	Large (7-8)	Total	
Number of Bidders	5	5	2	8	20	
Average Number of Bids per Bidder	308	817	2540	3718	2022	
Unit (TU) Characteristics by Region						
Region	4	5	9	12	13	Total
Number of TUs	5	10	9	1	7	32
Average Volume	2.50 (0.62)	2.32 (0.47)	2.62 (0.67)	2.03 (-)	2.78 (0.70)	2.52 (0.60)
Number of Bidding Firms	17	19	18	14	19	20

Business Size is measured by the number of TUs allowed to win, which is specified next to the business size type. e.g., medium firms are those who can win up to 3 or 4 TUs. Unit characteristics are provided for each of the five geographical regions auctioned in 2003. TU volumes are shown in million meals per year and standard deviations are shown in parenthesis. Bidding firms for each region refers to the number of firms that place at least one bid on a package that includes at least one unit in that particular region.

decided to exclude this year from our analysis. In 2004, the government introduced an electronic bidding system to the auction process that resulted in a huge increase in the number of submitted bids. On average, firms placed four times as many bids as they did in 2003, imposing an onerous amount of computation time in the estimation. However, the estimation was more manageable for the 2005 auction as the number of units auctioned and the participating firms were smaller, so we estimated that year to cross-validate the results. Table 3.1 provides summary statistics of the 2003 auction.

Additionally, we also collected data from two exceptionally smaller-scale auctions that were run between 1999 and 2005. These auctions were used to replace contracts from a few firms that had some irregularities. The auctions had eight and six units, respectively, and about thirteen firms participated. Given their smaller scales, these auctions can be estimated with the full-dimension model described in Section 2.2. We used them to compare results between the full-dimension and the extended size-based markup models, thereby providing validation of the methods developed in this work. Section 3.4.1 reports this comparison.

3.2.5 Discussion on the Assumptions of the Structural Model

In this section, we discuss how the assumptions of the structural model fit into this application. First, Assumption 2.1 ensures that the auction allocates at most one package per firm. While this restriction is not explicitly imposed in our empirical application, firms actually win at most one package in practice. In fact, among 41 winning firms between 2002 and 2005, only in one occasion a firm won more than one package; this happened in 2002. Table 3.2 summarizes the winning outcome in those years. Moreover, the government closely monitors the firms that participate in the auction and keeps track of strict records regarding firms' ownerships. Firms that are divided are actually treated as a single firm in the entire auction process, and the government can prohibit firms to win multiple packages through different entities. Hence, imposing this assumption is reasonable in this setting.

Assumption 2.2 imposes independent private costs, which seems adequate in this application.

	YEAR	2002	2003	2004	2005
Total number of winning firms		10	9	11	11
Number of winning firms winning one bid		9	9	11	11

Table 3.2 – Fraction of winning firms that won only one package bid for different auctions.

Roughly, 75% of the cost structure of firms is associated with food inputs and labor. A significant amount of these costs are common to all firms. However, this common part is not subject to uncertainty; it is determined by food prices and wages that are common knowledge to all parties involved at the time of the auction (wages of the cooks in this industry are actually regulated by the government). There could still be some uncertainty about future food prices due to the three year extension of the contracts. However, if prices change too much, there are rules in the auction that allow *all* firms to adjust their bids accordingly, dramatically reducing this risk (these rules are based on variation of food price indexes). Therefore, we believe the cost uncertainty is basically driven by firm specific differences in (1) logistics and management abilities (constitute the other 25% of the costs); and (2) idiosyncratic cost advantages related to food inputs determined by, for example, better contracting terms with providers. We think (1) and (2) are well captured by an independent private cost model.

Finally, Assumption 2.4 imposes that firms have the correct expectations regarding the vector of winning probabilities given their bids and competitors' strategies. As previously discussed, this is similar to assuming equilibrium play as is usually done in the structural estimation literature. While this may generally be a strong rationality assumption given the complexity of the auction, we believe that in our application it may be less so. First, these auctions are repeatedly run every year and all past bidding history is publicly available information (including winning and non-winning bids). In this regard, we exclude auctions where the units were awarded for the first time (years 1999-2001), because bidders had less experience and history to rely on, and were less sophisticated,

so that our structural model assumptions may be harder to justify. On a related point, note also that we assume firms maximize expected profits in the current auction without incorporating the impact on future auctions. This assumption is fairly standard in the structural estimation auctions literature, and we believe it captures the first-order objective of firms in this market. Moreover, given that our model is already very complex, adding dynamics would make it even more challenging. In addition, note that under our characteristic-based markup model firms do not need to estimate the winning probability of each individual package, but instead aggregated probabilities over several packages, which can be more manageable.

Moreover, we know from anecdotal evidence that firms in our application are quite sophisticated when bidding. In fact, because stakes are so high, firms invest important amounts of money in business intelligence. Using the historical information together with current market intelligence, firms indeed try to estimate the competitive landscape they will face. A personal interview with a former CEO of one of the supplying companies (that also consulted for other firms) provided more details about the bidding process. He began by creating a large spreadsheet with the packages he was interested on and calculated detailed cost estimates for these. Then, he would choose a markup for the different packages. Historical bid prices were used to decide the average margins to be charged; the margins became smaller over time as the market became more competitive (e.g. when larger catering companies entered the market). In addition, he would typically adjust the markup depending on the number of units in the package, asking for a lower per-meal markup for larger packages. Finally, based on historical bid data, he would also adjust markups for a few packages in which he was “more competitive”.¹ This provides further support for Assumption 2.4 and is also consistent with the extended size-based markup model.

¹The interviewee asked for confidentiality of his identity; we are grateful to this anonymous contributor for the insights provided.

3.3 Estimates of the Bid Distribution Parameters

This section describes the estimates of the distribution of the competitors' bids, based on the model presented in Section 2.5.1. We provide the results for the 2003 auction. Results for the 2005 auction have similar pattern and magnitude, and are omitted for brevity.

The school meals auction exhibits significant differences in discounts across the largest firms and the rest of the firms, so we categorize the bidders into two types, $k(f) \in \{L, O\}$ (for Large and Other), to estimate discounts. Recall from Section 3.2.1 that based on the financial evaluation and business capability, each firm has a maximum number of TUs that it is allowed to win in a given auction. We use this to measure the size of a firm. On the covariates Z_{if} (see equation (2.9)) we include an indicator on whether the firm won the unit on the previous auction (other covariates were also tested but they did not exhibit explanatory power).

Table 3.3 reports estimates of β^{scale} and $\beta^{density}$ from the first step regression (equation (2.8)). The scale and density per-meal discount curves, $g^{scale}(v_a, \beta_{k(f)}^{scale})$ and $g^{density}(v_c, \beta_{k(f)}^{density})$, are specified as step functions with interval of three million meals per year in the package volume v_a and cluster size v_c , respectively. Each number indicates the average discount in per-meal price when units are combined to form a package that belongs to the corresponding volume level. For example, when units are combined into package a with volume $v_a \in [18, 21]$, then, on average, a large firm submits a bid that is Ch\$ 22.78 cheaper per meal than the weighted average bid price of those individual units in the package. If all these units are located nearby and form a cluster, there is an additional discount of Ch\$ 11.27 on average for a large firm. The results show that large firms were able to provide higher discounts which can be as large up to 8.5% of the average bid price (the average bid price in the 2003 auction is Ch\$ 423). All the coefficients are estimated with precision and are different from zero with statistical significance (0.01% significance level). The R -square of the regression corresponding to equation (8) is 0.98 (with $\tilde{\delta}_{if}$'s as fixed effects), which provides some support that the parametric model adopted provides a reasonable approximation to the bid data generating process.

Table 3.3 – The first step regression results for the 2003 auction.

Large Firms			Other Firms		
Volume	Scale	Density	Volume	Scale	Density
[3, 6]	8.33 (1.30)	6.46 (0.51)	[3, 6]	8.50 (0.62)	1.82 (0.14)
[6, 9]	15.21 (1.33)	7.81 (0.53)	[6, 9]	11.86 (0.64)	3.31 (0.19)
[9, 12]	17.82 (1.31)	8.10 (0.55)	[9, 12]	13.50 (0.65)	3.92 (0.24)
[12, 15]	19.10 (1.30)	8.57 (0.56)	[12, 15]	13.44 (0.67)	5.69 (0.28)
[15, 18]	20.76 (1.29)	9.13 (0.57)	[15, 18]	12.42 (0.69)	6.96 (0.36)
[18, 21]	22.78 (1.30)	11.27 (0.65)	[18, 21]	10.90 (0.72)	
[21, 24]	24.38 (1.30)				
[24, 27]	24.95 (1.35)				

Regression results by equation (2.8) for the 2003 auction. Robust standard errors are shown in parenthesis. Volume is measured in million meals per year and discounts in Ch\$.

The second step estimation (equation (2.9)) provides estimates for the distribution of the average implicit prices $\tilde{\delta}_{i,f}$'s, characterized by $\{\bar{\delta}_i, \sigma_i\}_{i \in U}$, the covariance matrix Ω and β^Z , the coefficients of the firm characteristics. Due to space limitations, we do not report the estimates of the $\bar{\delta}_i$ parameters, but these were estimated with precision - on average, the standard errors are 1.2% of the point estimates. The estimated coefficient for β^Z is -5.986 with a p-value of 0.012, suggesting that on average firms that were awarded the unit in the previous auction submit bids that are around 1.5% cheaper.

Table 3.4 shows the correlations between the region effects $\psi_{r(i),f}$ (which were calculated based on estimates of the variance/covariance matrix Ω). These estimates imply a significant positive correlation among units: on average, the correlation between the implicit prices of two units in the same region is 0.68, and 0.45 for units located in different regions. The last column of the table

shows the standard deviations of each region effect $\psi_{r(i),f}$ (which corresponds to $\sqrt{\Omega_{rr}}$). All the standard errors of the maximum likelihood estimates of equation (2.9) are computed via a parametric bootstrapping procedure.

We also tested some of the parametric assumptions in our model. First, our approach assumes that the implicit prices δ_{if} follow a normal distribution. For all the units, a Shapiro-Wilks test cannot reject this assumption at 5% significance level (p-values are in the range 0.073 to 0.92). To test the restrictions on the covariance structure of the implicit prices $\{\delta_{if}\}_{i \in U}$ imposed by equation (2.9), we compared this model against a more general model where the full covariance matrix across all units Σ is unrestricted. A likelihood ratio test cannot reject that the two models are equivalent (p-value > 0.6 for both 2003 and 2005 auctions). These tests confirm that our parametric assumptions make the estimation tractable, while being reasonably flexible.

3.4 Cost and Markup Estimation Results

We estimate markups and costs for the school meals auction application using the two-step method described in Section 2.2.2 and the competitors' distribution of bids estimated in the previous section. For the second step we use the extended size-based markup model heuristic described in Section 2.4.2. In addition, to validate our estimation approach, we also compared the estimation via the full-dimension model and the extended size-based markup model using the two small CAs. Before presenting these estimation results, we first discuss issues on specifying the package-characteristic matrix W in our specific application. Then the small auction results are provided followed by the results for the large-scale CAs.

First, in our empirical application, package volume – defined in terms of the annual demand of meals for the TUs contained in the package – has a first-order effect on the bid price and so prices may vary substantially even within each size group. For this reason, we assume that bids within each group have a common *per-meal* markup, instead of a fixed *absolute* markup. Defining b_a and c_a as the per-meal bid and per-meal cost of package a , and defining the non-zero components of W

Table 3.4 – Correlation structure among regions from the second step regression.

Region	Correlation Coefficients						Standard	
	4	5	9	12	13	Deviation	Deviation	
4	1.00 (0.00)	0.52 (0.21)	0.31 (0.27)	0.45 (0.24)	0.67 (0.17)	14.56 (3.20)		
5	-	1.00 (0.00)	0.65 (0.16)	0.69 (0.17)	0.69 (0.13)	14.52 (2.55)		
9	-	-	1.00 (0.00)	0.42 (0.22)	0.09 (0.27)	22.92 (4.02)		
12	-	-	-	1.00 (0.00)	0.48 (0.22)	46.48 (9.97)		
13	-	-	-	-	1.00 (0.00)	13.46 (2.29)		

This table shows the correlation structure among regions from the second step regression (equation (2.9)) for the 2003 auction. Parametric bootstrapping standard errors are shown in parenthesis. Standard deviations of regional effects are measured in Ch\$.

as $W_{as} = v_a$, the firm's decision variable θ can be interpreted as a per-meal markup vector. Related to this, it is worth noting that in Proposition 2.2 we assumed that $\{\theta_a\}_{a \in \{1, \dots, K\}}$ and θ_u are absolute markups. In the per-meal markup specification, we obtain absolute markups by multiplying them by the total meal volume in the package through the W matrix. In this case, the α_a 's are not necessarily all negative a priori. However, in our estimations they turned out to be negative for most of firms and auctions, leading to weighted average markups when aggregated.

Second, in Section 2.4.2 we discussed how to divide the markup groups to refine the size-based markup model. We try to have as many markup parameters as possible to the extent that computational tractability is maintained. In our actual estimation for large-scale CAs, we use the threshold probability of 10^{-3} to identify the high probability special package bids. Packages with winning probabilities above 10^{-4} are categorized in the medium probability groups, and the rest in the low probability groups. We selected these thresholds mainly for technical reasons. Recall that we use the first-order condition equation (2.6) for our estimation and that requires estimating the group winning probability and its Jacobian. Since we estimate these via simulation, to achieve accurate estimation in reasonable computation time, the group probability should be above certain magnitude. Given the simulation length that we use (around 100,000 runs), it seemed desirable to force the individual probability to be at least 10^{-3} . The two threshold probabilities were chosen to ensure this. With this grouping procedure, a typical firm has a markup vector with dimension $d = 20$.

3.4.1 Small Auctions Estimation

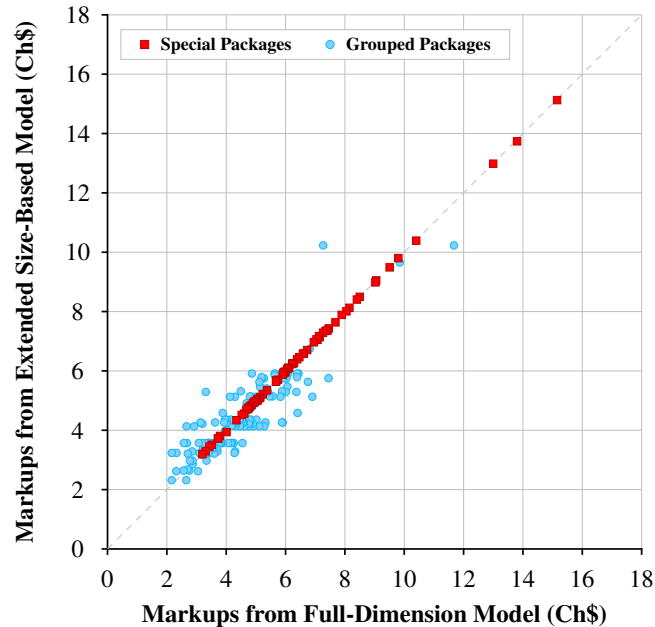
In this section, we estimate the markups and costs with the full-dimension model and the heuristic based on the extended size-based model in the small-scale CAs. As in the large-scale case, we follow the two-step approach described in Section 2.2.2; in the first step we parametrically estimate the distribution of competitors' bids using the model described in Section 2.5.1. Then, in the second step, we applied the two different markup models in the first-order conditions and compare the results. For the extended size-based markup model, bids with a winning probability above

10^{-3} were considered high winning probability (special) packages and were assigned a separate markup. On average, the extended size-based model uses only 35% of the markup variables of the full-dimension model. Note that because the relevance of each bid only depends on its winning probability and not on the markup specification, the set of relevant bids is the same for the two models, and therefore, they are comparable. Packages with winning probability below 10^{-4} were considered irrelevant. For each firm, the aggregated winning probabilities of these irrelevant bids are on average less than 1% of the firms' total winning probabilities. Hence, the effect of ignoring these irrelevant (with positive probability) bids was negligible.

Figure 3.1 shows a scatter plot of the estimated per-meal markups from the two methods. Overall, the markup estimates of the extended size-based markup model are similar to those obtained with the full-dimension model. The correlation between the markups is 0.982; their ratio is on average 1.003 with a standard deviation of 0.127. The figure illustrates that the estimates for the special packages are even closer to each other: the ratio is on average 0.998 with a standard deviation of 0.004 and the correlation is 0.999. This provides some support for the conjecture that grouping packages would have a negligible impact on the markups of the special packages, as discussed in Section 2.4.1.

We also note that this specification of the extended size-based markup model separates each size-group further into two sub-groups with medium and low winning probability packages as described in the heuristic in Section 2.4.2. We used a winning probability of 6×10^{-4} to divide the groups into medium and low winning probabilities. We observed that this additional refinement to the size-based markup model – which was used in the calculation of the markups shown in the figure – helped to improve the markup estimates (i.e. the estimates were closer to the estimates of the full-dimension model).

To provide the performance metrics of a CA (winning bidders' profit margins and efficiency loss), we need to estimate the total supplying costs of the winning firms in the CA and the efficient allocation. Notably, the estimated total supplying costs both in the winning CA and efficient allocations are very close between the two methods (differ by less than 0.1%).

Figure 3.1 – Markup estimates from the full-dimension and extended size-based model.

This is a scatter plot comparing the estimates from the full-dimension model with those from the extended size-based model. Circles denote markups of size-groups and squares show the special packages. Two special packages with large markups (around Ch\$ 25-28) are not shown to improve visualization. The differences in the estimated markups for these two special packages was negligible (less than 0.1%).

We also compared the estimates of the full-dimension model with those of the *pure size-based* markup model, which does not isolate the high winning probability packages. Here, the pure size-based markup model also specifies markups per meal, including the volume of packages in the W matrix as described in Section 2.4.2. In this case, the ratio of the estimated per-meal markups from the two methods (with the full-dimension markup in the denominator) is on average 1.138 with standard deviation of 0.365. As expected, the pure size-based markup model results in significant bias in the estimated markups relative to the extended size-based model. In particular, not separating the high winning probability packages leads to overestimating the group markups. A partial explanation is that high winning probability packages tend to have larger estimated markups relative to the rest of the group. To see that, let \bar{m}_s^h and \bar{m}_s^r be the *average* estimated per-meal

markups of special packages with high winning probabilities (h) and the rest of the packages (r), respectively, for a particular firm and package size s in the full dimension model. The ratio $\bar{m}_s^h / \bar{m}_s^r$ is on average 1.22 with a standard deviation of 0.37. Hence, following the discussion around Proposition 2.2, removing them from the group reduces the bias associated with grouping in an important way.

Finally, we also performed an experiment to examine the impact of package density on strategic markup adjustments. As briefly discussed in Section 2.4, economic theory predicts strategic markup adjustments mainly driven by scale, and we did not incorporate explicitly the density effects in our extended size-based markup model (even though the separation of special packages may correct for it to some extent). To further justify this, we enriched the extended size-based markup model with a markup variable associated with a per-meal density measure of the package. The measure ranges between 0 and 1 and becomes larger as the package has more co-located units. The density measure is motivated by the density discount function used in Section 2.5.1. In particular, we tested two different density measures for robustness of the results and they both gave very similar results. We provide the details of these density measures and the results in the appendix of this chapter (see Section 3.7.2). The estimates imply small markup adjustments associated with density; they are on average 0.11% of the average bid price; this is an order of magnitude smaller compared to the markup adjustments associated with the scale effect. Moreover, the extended size-based model with and without additional density parameter provide essentially the same markup estimates. The ratio of the two markup estimates is on average 0.999 with standard deviation of 0.063. This provides evidence that the density markup parameter do not play a significant role in the markup estimation.

Overall, this section provides evidence that our extended size-based model provides accurate approximations to the full-dimension model estimates, requiring significantly lower computational effort. In the small-scale CAs, the heuristic is an order of magnitude faster to run than the full-dimension model, producing similar estimates. In the large-scale CAs, the full-dimension model is computationally infeasible; we present the results using the extended size-based model heuristic in the next section.

3.4.2 Results for Large-scale CAs

The extended size-based markup model was used to estimate markups and costs for the package bids in the 2003 auction. We also performed estimation for the 2005 auction and report some of the results in the appendix of this chapter (see Section 3.7.1). In 2003, a total of 32 TUs in 5 regions were auctioned and 20 firms participated placing more than 2000 bids per bidder on average.

After estimating the markups and costs of these firms, we numerically checked if the estimated markup variables locally optimize the expected profit. Note that to fully evaluate the local optimality of the markups, we need to estimate the Hessian matrices of the bidders' expected profit. However, estimating the Hessian matrices is computationally very intense, requiring an order of magnitude more computational time than estimating the markups. Instead, we checked the second-order derivatives of the bidders' expected profits with respect to each of the markup variables and they were all negative, consistent with the local optimality of the estimates.

Given the firms' cost and markup estimates, two groups of firms were identified based on each firm's *total winning probability*, that is, the firm's aggregated winning probabilities over all packages in the auction. The "competitive" group consists of ten firms whose total winning probabilities are higher than 45%. Firms in the other group have very low winning probabilities (less than 2%) except for one with 16% of total winning probability. In addition, from the 20 participating firms there are two extreme firms with very competitive bids, for which the estimated markups are unreasonably high and lead to negative costs for some packages. Despite their competitive prices, these firms did not win any units and were disqualified from the allocation process, because of quality considerations. For these reasons, we omit them from our analysis hereafter. In terms of markups, the competitive firms have markup estimates ranging from 1.2% to 18% of the average bid price with an average markup of 4.4% of the average bid price (US\$ 0.88 per meal). The other firms have lower markups, resulting in an average markup over all firms of around 2.8% of the average bid price. Table 3.5 shows the average per-meal markup estimates for each package size (1 through 6 units) for representative firms in three different levels of total winning probabilities. The

estimates indicate that firms reduce their markups as the size of packages increases, showing that some portion of the discounts in package bids are due to markup adjustments.

Table 3.5 – Estimated markup levels and the winning probability levels.

Firm	Prob	Average Markup for Each Package Size						Overall
		1	2	3	4	5	6	Average
47	0.9193	22.64	15.07	12.14	7.98	7.54	7.19	9.88
36	0.6642	3.00	2.39	2.21	1.77	1.50	1.41	2.07
19	0.1578	0.81	0.82	0.84	0.79	0.72	0.71	0.79

Results from the markup estimation for representative firms of different winning probability levels for the 2003 auction. *Prob* refers to the probability that the firm wins any package. The rest are the average per-meal markups corresponding to each package size. The markups are shown as a percentage of the average bid price per meal (US\$ 0.88).

Firms submit hundreds to thousands of bids, and about 13% of them are relevant bids.² For the competitive firm group, the fraction of relevant bids is higher, 22%. With the estimated markup and cost information of relevant bids, we are able to compute the total cost and markup of the CA allocation. The total procurement cost for the government was US\$ 70.5 million per year and the supplying costs for firms was US\$ 67.2 million per year. This yields an average profit margin to winning firms of 4.8%. This level of profit margins is consistent with the Chilean government's estimate for this market. In addition, the Chilean government has their own estimates for the average TU costs, and these exhibit similar levels compared to our estimates. Both facts are re-assuring. Note that the government does not have estimates for TU costs for each firm, nor of the variance of these quantities. Moreover, they do not have reliable estimates for the level of cost synergies. Hence, the government's estimates are of course insufficient to evaluate the performance of the auction, which is the objective of this work.

²Packages with estimated winning probability below 10^{-5} were considered irrelevant.

Finally, to compare results, we also performed the estimation for the 2005 auction, where 16 firms participated for 23 units. The results are consistent with the 2003 auction, both in the shape and level of the estimated markups. The total procurement cost amounts to US\$ 53.4 million and the total supplying cost is US\$ 51.5, which give 3.5% of average profit margins to winning firms.

Next, we evaluate the cost synergies – cost savings from combining units together – implied by the estimates. Recall that our main objective is to determine what portion of the observed package discounts is due to cost synergies. Given the markup estimates, the per-meal cost of each package a submitted by firm f is given by $c_{af} = b_{af} - w_a \theta_f / v_a$, where θ_f is the markup vector estimated for that firm, b_{af} is the per-meal bid price placed by firm f for package a , and w_a is the a^{th} row of package-characteristic matrix W used for bidder f . A direct calculation of the per-meal cost synergy in this package, denoted by s_a , can be computed from $s_a = \sum_{i \in a} \frac{v_i}{v_a} c_i - c_a$, where c_i is the point estimate for the cost of unit i . Table 3.6 shows some summary statistics of the cost synergies. The cost synergies estimated directly from the cost estimates tend to increase as the size of packages grow, and suggest that there are significant cost synergies amounting up to around 4.5% of the average bid price.

Table 3.6 – Average cost synergies.

Package Size	2	3	4	5	6	7	8	Average
Cost Synergy (CH \$)	5.17	11.59	14.40	13.71	15.07	16.64	18.93	13.64
% of Average Bid Price	1.22	2.73	3.39	3.22	3.55	3.92	4.45	3.21
Number of Observations	280	85	121	49	126	169	205	

Average cost synergies computed directly from estimated costs of individual units and multi-unit packages for 2003 auction. Package size refers to the number of units in a package. The cost synergy measures the average per-meal cost savings when the units are combined to form a package of given size.

One disadvantage of estimating cost synergies in this direct way is that the synergies can only be computed for packages containing units whose single-unit bids are all relevant, which is not a

representative sample of the bid population. In order to use a larger portion of the packages to estimate cost synergies, we run a regression similar to (2.8) but replacing the dependent variable b_{af} by c_{af} :

$$c_{af} = \sum_{i \in a} \xi_{if} \frac{v_i}{v_a} - g^{scale}(v_a, \gamma_{k(f)}^{scale}) - \sum_{c \in Cl(a)} g^{density}(v_c, \gamma_{k(f)}^{density}) \cdot \frac{v_c}{v_a} + \varepsilon_{af}, \quad (3.1)$$

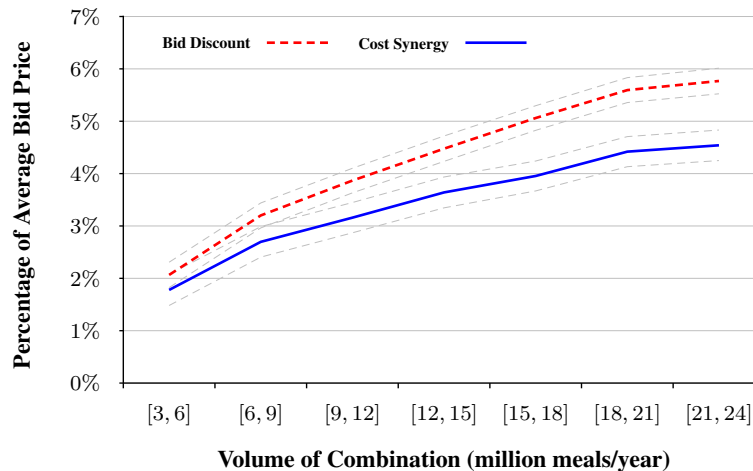
where again $k(f) \in \{L, O\}$ indicates one of the two firm types. This regressions projects the estimated costs on scale and density synergies, and provides estimates of the costs for every unit, ξ_{if} , including those for which the single-unit bid was irrelevant. For relevant single-unit bids, ξ_{if} and c_{if} are quite close: their correlation is 0.993, the absolute different is about 1% and their ratio averages 1.00 with a standard deviation of 0.011. Hence, equation (3.1) seems a reasonable approach to estimate cost synergies.

Figure 3.2 plots the total cost synergies and bid discounts as a function of package volume. We note that there are two small firms whose estimated cost synergies are significantly different from the rest firms, and they are not included in the figure. The results show that while there is some strategic markup adjustments, most bid discounts (at least 75%) are actually explained by cost synergies. These synergies are quite significant and can be as large as 4.5% of the bid price on average.

3.5 Efficiency and Counterfactual Analysis

The previous results suggest that in our application allowing package bidding may be appropriate: cost synergies are significant and account for most bid discounts vis-à-vis strategic markup adjustments. Moreover, the overall markups that firms gain do not seem too large, resulting in a reasonable total procurement cost. Overall, our results suggest that the advantages of using package bidding (allow bidders to express cost synergies) may be larger than its disadvantages (the additional flexibility that firms can use to strategize and game the mechanism).

In this section we use our estimates to provide sharper results concerning the efficiency and procurement cost of our CA. In particular, we study the allocative efficiency and procurement cost

Figure 3.2 – Overall bid discount and cost synergy curves.

Bid discount and cost synergy curves due to the combined effects of economies of scale and density for the 2003 auction. To capture the combined effects, the estimation is done using equations (2.8) and (3.1) without the density terms, $g^{density}$. 95% confidence intervals are shown in dashed lines. The number of observations is 4397.

of the first-price sealed-bid CA, and also examine the impact of the allocative constraints imposed in the auction on the efficiency of allocation. We provide results for the 2003 auction. The results for the 2005 auction are similar and consistent with the 2003 auction. We provide the counterfactual results for the 2005 auction in the appendix of this chapter (see Section 3.7.1).

3.5.1 Performance of the First-Price CA

We first study the allocative efficiency of the first-price CA. The winning bidders' costs under the first-price CA allocation can be directly computed using the cost estimates obtained in Section 3.4. If we had the cost estimates for *all* possible packages, we could also calculate the efficient allocation, that is, the combination of package bids among all firms that achieve the minimum possible total cost. Unfortunately, our structural estimation method only identifies the costs of relevant bids, and the efficient allocation over this subset of combinations could overestimate the cost of the true efficient allocation that considers all possible packages.

To address this issue, we propose estimating the cost of irrelevant bid packages through an out-of-sample extrapolation based on equation (3.1). However, the total number of feasible packages are in the order of millions and it is computationally infeasible to extrapolate to the entire set of (out-of-sample) packages. Instead, we choose the set of packages on which at least one bidder placed a bid, which is in the order of 30 thousand packages. We call this the *expanded package set*. Then, for each firm, we extrapolate costs to all packages in this expanded package set that are also in the set of feasible allocations. While this is a small subset of all possible packages, it provides a reasonable approach to extend the set of bids observed in the data.

This out-of-sample extrapolation approach implicitly assumes that the selection of the bids in the irrelevant bid sample is independent of the costs of these units. Recall that irrelevant bids include bids that were not submitted by the bidder. Hence, in our application, it could be possible that the sample selection of irrelevant bids is related to costs: for example, bidders are likely to bid on the subset of combinations where they are more competitive, so that higher-cost combinations are not submitted. If this is the case, then our cost extrapolation procedure could lead to a cost estimate of the efficient allocation which is lower than the true one, so that we could *overestimate the true efficiency-loss* of the first-price CA.

Recall that in 2003, the bidders' supplying costs given by the auction allocation were equal to US\$ 67.2 million per year. The efficient allocation that minimizes the total supplying cost among the feasible allocations over the set of relevant bids is equal to US\$ 66.7 million per year, implying an efficiency loss of 0.65%. When considering the expanded package set, the total supplying cost of the efficient allocation goes down to US\$ 66.2 million per year, with an efficiency loss of the first-price CA of 1.5%. It is worth noting that the first-price CA tends to identify the most cost-efficient firms in the different geographical regions. More specifically, there are nine firms in the CA allocation and ten firms in the efficient allocation; the majority of them –seven firms– are present in both cases. Two firms are allocated the exact same set of packages in both cases and other firms win packages that contain many overlapping units or units from the same geographical regions. The efficiency loss is arguably low. We believe this result is essentially driven by the high level of competition in

the auction; there is a reasonable number of firms and most firms compete in all units and submit many package bids. For this reason, firms do not seem to have enough market power to significantly harm efficiency by using the flexibility that package bidding allows with strategic motivations.

Although the efficiency loss is overall evidently low, a few firms indeed engage in strategic markup adjustments that are consistent with the economic arguments provided in Section 3.1. For example, there is one firm winning eight units in the CA allocation, that essentially leveraged its cost advantage in some units to win another unit in which it was not the cost-efficient firm. If this firm was forced to just win its cost-efficient bundle, the loss of 0.65% over relevant bids is significantly reduced. In summary, the high efficiency and relatively small profit margins for firms (around 5% as presented in Section 3.4.2) achieved by the school meals CA suggests that it is a reasonable mechanism for the procurement of this public service.

3.5.2 Supplier Diversification

The CA of our application imposes three types of constraints aimed at preserving a more diversified supplier base: (1) a single bidder cannot be awarded more than 16% of the total volume including outstanding contracts awarded in previous years (*market share* constraint); (2) a minimum number of winning firms on each auction (*global competition* constraint); and (3) a minimum number of winning firms on each of the 13 pre-specified geographic regions (*local* constraints). We now focus on measuring what is the efficiency loss imposed by these constraints.

To study efficiency of the first-price CA, we have already calculated the minimum-cost allocation that satisfies these constraints. We could compare this with the minimum-cost allocation obtained under the larger feasible set of allocations when the constraints are removed. However, this may not be a fair comparison because bids on packages that violate some of the constraints are not submitted by the bidders. In other words, in the counterfactual world without the constraints we should observe new package bids that are not observed under the current format with the constraints. To address this issue, we expand the set of submitted bids in the counterfactual without the constraints as we now explain.

First, consider the market share constraint that imposes a maximum volume of 16% of the total volume of the country, equivalent to about $K = 40$ million meals per year to each firm. Under this constraint, bids on packages with larger volume than K will never be observed in the data. It turns out that because of the 8 unit limit for the packages, the market share constraint is never binding for those firms which do not have any existing outstanding contracts, because the maximum volume that can be achieved with 8 units is less than K . So those firms do place bids on packages of any volume with at most 8 units. We call such firms whose bidding is not limited by the market share constraint the *unrestricted firms*.

To extrapolate costs to packages violating the market share constraint we do the following. Consider a large bidder f that has existing outstanding contracts for a total volume of X . This firm can only submit packages of volume less than or equal to $K - X$. Removing the constraint would allow this bidder to submit packages of any volume up to K , as long as they have 8 units or less. We denote by A_X the set of observed combinations that are infeasible for bidder f but feasible for the unrestricted bidders, hence contained in the expanded package set. We can use regression (3.1) to predict the costs of combinations in A_X for bidder f . Doing this for all bidders allows us to build a larger feasible set that contains bids that would not be feasible when the 16% market share constraint is included. Again the expanded package set which is in the order of 30 thousand packages – still less than the 20 million possible packages that could be submitted – and provides a reasonable set of bids to evaluate the effect of removing the market share constraints.

In contrast, packages that violate the local competition constraints are almost never observed. To illustrate why this is the case, consider region 13 which has seven units but the minimum number of firms required to win is four. Hence bids on any package containing five or more units in region 13 will violate this constraint and will never win. Note that unlike the market share constraint which is a firm-wise restriction, local constraints are applied to all firms and hence no such packages are found in the expanded package set. For this reason, we cannot analyze the effect of removing the local constraints using the same approach to expand the set of bids. Finally, we note that it is not a priori clear whether removing the global competition constraint would result in significantly

different package bids submitted, because in any event firms cannot submit packages larger than 8 units. Therefore, we do not include additional bids associated to removing that constraint, and we focus on the efficiency loss caused by the market share constraint and the global competition constraint.

To measure the efficiency loss due to the market share and the global competition constraints, we compare the minimum-cost allocations with all those constraints and without the two types of constraints. We find that this efficiency loss in 2003 is about 0.57%, which is relatively small. The final allocations in both cases look similar. Nine firms win in both cases and only one winner is replaced by another. Two firms win exactly the same packages, and six other firms have many of the winning units overlap in both cases or win units in the same region. The efficiency loss is mainly triggered by one large firm who won a package of two units with market share constraint and won a package of five units in the unconstrained case. The small impact of these constraints on efficiency can be partially explained by the structure of the cost synergies in the industry. As we saw in section 3.4, scale cost synergies get exhausted, so there are small cost reductions for combinations that lie beyond the volume range that is currently feasible in the auction. To further evaluate the inclusion of these constraints in the mechanism, it would be useful to measure the value that the constraints aimed at promoting supplier diversification bring in terms of increased competition. For example, Olivares et al. (2012) show that local competition, measured by the number of firms serving nearby units, has a significant effect in reducing prices in this application. This suggests that supplier diversification at a local level can lead to increased competition. We leave this analysis for future research.

3.6 Conclusions

In this chapter, we apply our structural estimation approach to the large-scale Chilean school meals CA. We find that cost synergies in this auction are significant and the current CA mechanism, which allows firms to express these synergies through package bidding, seems appropriate. In

particular, the current CA achieves high allocative efficiency and a reasonable procurement cost. We believe that this is the first empirical analysis documenting that a CA performs well in a real-world application.

More broadly, our results highlight the importance of the simultaneous consideration of the suppliers' operational cost structure and their strategic behavior for the successful design of a CA. In this way, we hope that this research agenda enhances the understanding of the performance of CAs and thereby provide insights to improve their design.

3.7 Appendix for Chapter 3

3.7.1 Results for the 2005 Auction

The results from the 2005 auction are similar to those of the 2003 auction: they give similar level of cost synergies and strategic markup adjustments as well as the winning firms' profit margins. For example, in 2003 the cost synergies ranges from 1.8% to 4.6% of the average bid price and in 2005 they were from 2.6% to 5.8% of the average bid price. The strategic markup adjustments were 75% or more of the discounts in 2003 and 70% or more in 2005. In addition, the winning firms' average profit margins were around 5% in 2003 and 3.5% in 2005. While the results provide roughly similar magnitude of these estimates, we still observe some differences between the two auctions. However, note that the units in these two auctions are different so the costs need not be the same (characteristic of the units and the meal plans are different). The number of bidders in the two auctions was also different which can lead to differences in the markups. In this section, we further provide the counterfactual results for the 2005 auction.

First, we find that the allocation is also highly efficient in 2005. Recall from Section 3.4.2 that the total annual supplying cost in the first-price CA is US\$ 51.53 million. The total annual supplying cost of the minimum-cost allocation is US\$ 51.49 million over the set of relevant bid packages and US\$ 50.70 million over the set of expanded package sets. This gives about 1.6% of efficiency loss in the allocation by the first-price CA.

Finally, in 2005, we have a bit larger but still small efficiency loss incurred by the allocative constraints. We consider the loss due to the market share constraints and global competition constraints. The efficiency loss in the constrained auction is 2.8% compared to the minimum-cost allocation without those constraints. In 2005 the impact of the global competition is higher; it imposes a minimum of 9 winners out of 16 bidders in 2005; in 2003 it also imposed the same minimum but out of 20 bidders.

3.7.2 Package Density Measures

In Section 3.4.1, we compared the estimated markups from the extended size-based model with and without an additional markup variable associated with the density of the packages. In this section, we provide two density measures we used in those comparisons.

To define measures of package density, we continue using the cluster volume described in Section 2.5.1. Recall that in equation (2.8), the density discount term is specified as a step function of cluster volume. Ideally, to test the effect of the density on the markup adjustments, we would also add multiple density markup variables for each level of cluster volume. However, this is computationally costly. Therefore, our objective in this section is to come up with a single-parameter measure of density that reasonably follows such description. Specifically, we consider the following two candidates; (i) one that assumes that markup adjustment due to density are “linear” in cluster volume; and (ii) one that allows non-linearity in cluster volume.

First, the linear density per-meal measure for package a , denoted by d_a^l is defined as follows:

$$d_a^l = \frac{1}{v_a} \sum_{c \in Cl(a)} v_c,$$

where $Cl(a)$ is the set of clusters in package a and v_c is the volume of cluster $c \in Cl(a)$. This measure takes the value between 0 and 1 (clusters are sets of co-located units so they always contain more than unit). The first term is to normalize the density by volume, and the second term implies that the cluster density is linear in its volume. Let θ_d be the per-meal markup variable associated with this density measure, then the total package markup adjustment from the density effect is given by $d_a^l \theta_d v_a = \sum_{c \in Cl(a)} v_c \theta_d$. One limitation of this linear density measure is that it cannot capture differences arising from different composition of clusters for packages of the same size. For example, consider two packages consisting of four units with identical volume. Suppose further that one package has a single cluster of four units while the other has two clusters of two units. Then, the linear density measure coincides for both packages. However, intuitively one would expect larger density effects for the package of four clustered units.

To accommodate this, we introduce a non-linear measure of density, defined as follows (per-

meal):

$$d_a^n = \frac{1}{v_a} \sum_{c \in Cl(a)} \frac{v_c}{v_a} v_c.$$

Note that this measure also takes the value between 0 and 1. The main difference is that now the cluster density has been weighted by its relative volume to the total package volume to capture the relative impact of each cluster to the package density. Under this measure, the density of the two packages in the above example will now be different as expected.

To capture the effect of package density on the potential markup adjustments, we added one more markup variable that is associated with a density measure. Then, we compared the markup estimates from this model with those of the extended size-based markup model. For robustness, we have tested using the two density measures. In both cases, the markup adjustments associated with the density was relatively very small; on average the magnitude of such markup adjustments was 0.096% or 0.134% of the average bid prices using non-linear and linear measures, respectively. In fact, this is an order of magnitude smaller than the markup adjustments associated with the scale effect. Given such a small effect, the markups with and without the density term are very similar with both measures. Using the non-linear measure, the ratio between the model with the density measure and the extended size-based markup is on average 0.999 with a standard deviation of 0.063. With the linear measure, they are 1.008 and 0.079, respectively. These results provide evidence that the strategic motivations associated with density seems much smaller, if any, compared to the scale effect. The results also support our choice of size as a main package characteristics to identify the sources of bidders' strategic behavior.

Chapter 4

Vickrey-Clarke-Groves Mechanism

4.1 Introduction

The Vickrey-Clark-Groves (VCG) mechanism generalizes the second-price auction in multi-unit settings. The units in a VCG auction can be homogeneous as in the Treasury Bill auctions or heterogeneous like in the FCC's spectrum auctions. It is well known in the literature that the VCG mechanism is strategy-proof: in a procurement setting, reporting one's *true supplying cost* is (weakly) better than any other bidding strategy, no matter how the competitors determine their bids (see Mas-Colell et al. (1995) for a formal proof). Given the reported bids, the VCG mechanism assigns the units to bidders so that the total reported cost – sum of the winning bids – is minimized.¹ Therefore, the strategy-proofness along with the allocation rule implies that the VCG mechanism always achieves the efficient allocation ex post; the VCG allocation is the one that minimizes the total supplying cost. As discussed in earlier chapters, achieving allocative efficiency is particularly attractive in public procurement applications such as the Chilean school meals program studied in Chapter 3.

Despite the aforementioned theoretical virtues, however, the VCG mechanism has been criticized for numerous drawbacks in settings with synergies, leading to a very rare use in practice.

¹Similarly, in a forward auction, the allocation is determined by maximizing the total reported valuation.

Ausubel and Milgrom (2006) explains this:

Why is the Vickrey auction design, which is so lovely in theory, so lonely in practice? The answer, we believe, is a cautionary tale that emphasizes the importance of analyzing practical designs from many perspectives. Vickrey's design has some impressive theoretical virtues, but it also suffers from weaknesses that are frequently decisive.

In particular, they show that in the face of complementarities, the total procurement cost under the VCG mechanism can be prohibitively high. Although the priority is generally given to allocative efficiency in a public procurement project, the procurement cost is also an important performance measure and can be a decisive factor to deny the use of the VCG mechanism. On top of this, the weaknesses of VCG also include vulnerability to collusion by a coalition of losing bidders and vulnerability to the use of multiple bidding identities by a single bidder (shell bidding), making the mechanism even less attractive in practice.

When we say that a mechanism may result in a prohibitively high procurement cost, an important question would be how high it should be to be *prohibitive*. For this, Ausubel and Milgrom (2006) and Milgrom (2004) adopt *the core*, a solution concept in cooperative game theory, as a competitive benchmark, and show that the VCG outcome may not be in the core of the transferable utility cooperative game played among the bidders and the buyer (auctioneer), resulting in high procurement cost. More specifically, it could be understood that if the outcome is not in the core, the payments are so high that there is a group of bidders who can offer a more favorable deal to the auctioneer. In addition, they also show that the other drawbacks mentioned above are also closely related to the core – the deficiencies disappear if the VCG outcome is in the core. These findings have motivated an active research agenda in recent years that studies alternative payment rules, giving rise to the so-called “core-selecting auctions” that determine the auction allocation given the reported bids so that the outcome is *always* in the core, alleviating the aforementioned drawbacks (e.g. Day and Milgrom (2008)). However, Goeree and Lien (2009) point out that those variants may fail to maintain the aforementioned virtues of VCG.

Therefore, an important research question would be whether we can characterize the cases where the VCG outcome is in the core (or close to it) and achieves reasonable procurement costs. In fact, it can be shown that when the units are substitutes, the VCG outcome is always in the core. However, this may not be the case when the units exhibit complementarities, where allowing package bidding has potential benefit. Hence, it is on itself interesting to see how the VCG mechanism would perform in a real-world application, especially in which the units exhibit significant cost synergies. Thus motivated, we conducted a counterfactual analysis to study the performance of VCG for the Chilean school meals program. Recall that in Chapter 3, we have estimated the supplying costs of the firms participating in the Chilean school meals program, using the structural estimation method developed in Chapter 2. Having estimated the bidders' supplying costs which would be their reported bids if VCG had been used, we are able to compute the procurement cost under this counterfactual scenario. Contrary to the theoretical predictions in the literature, the results show that VCG performs well in that particular application despite the significant cost synergies depicted in Figure 3.2 – the VCG procurement costs were very close to those of the first-price CAs in both 2003 and 2005. Consistent to this observation, we also found that the VCG outcome is *essentially* in the core. We believe this result is driven by the significant amount of competition introduced by the large number of package bids submitted by firms. In the Chilean school meals auction, most of the firms place bids on every unit, and therefore unit-wise the market is quite competitive. On top of that, firms also place many package bids.

Motivated by this important observation using real world data, in this chapter we address such apparent paradox between the theory and our empirical application. Focusing on the high procurement cost issue, we study the impact of competition on the revenue performance of the VCG mechanisms using an asymptotic analysis. Specifically, we find that the first order impact on the performance of VCG is measured by the amounts of bids rather than just the number of bidders; VCG mechanisms are expected to work better when the bidders' interests are not limited to a small subset of units and when they place ample combination bids that contain those units that they are interested in. In many practical applications, such as spectrum rights or transportation procurements,

it is expensive for bidders to correctly estimate their own valuations (or costs) on the combinations of units. In such an environment, bidders' interests could be restricted to a small subset of units, and more importantly, they are discouraged from placing many combination bids. This is the type of settings studied in the theoretical literature mentioned above. In the school meals procurement setting, on the contrary, it is relatively straightforward for bidders to estimate their supplying costs and they were able to place a large number of bids over most of the units in the auction. This scenario is precisely one in which our analytical results suggest that VCG should perform well. We note that the analysis we provide in this chapter is yet preliminary. However, the findings in this chapter adds useful insights for the practical usage of the VCG mechanism. Our analysis opens interesting future directions, which we discuss in Section 4.5.

The rest of the chapter is structured as follows. Section 4.2 describes the rules of the VCG mechanism and relates its outcome to the core. Section 4.3 provides the counterfactual analysis results using the estimates from the Chilean school meals auction data. We provide our main analysis on the asymptotic revenue properties of the VCG mechanism in Section 4.4, and Section 4.5 provides the main conclusions to the chapter.

4.2 VCG Mechanism and the Core

In this section, we provide a detailed description of the rules of the VCG mechanism in a procurement combinatorial auction setting. Then the connection between VCG outcomes and coalitional games is discussed in Section 4.2.2. Finally, Section 4.2.3 provide an example where a VCG outcome is not in the core, leading to a high procurement cost.

4.2.1 Rules of the VCG Mechanism

We first provide a formal description for the rules of the VCG mechanism. We begin with providing useful notation.

Notation. The auctioneer, indexed by 0, invites n risk-neutral bidders, indexed by $i = 1, 2, \dots, n$, to procure K units of (possibly distinct) items. We let N be the set of those n bidders and U be the set of K units. As in previous chapters, we use b_a^i to denote the bid price asked by bidder i for package a . We use $x := (x_1, x_2, \dots, x_n)$ to denote an allocation outcome which maps bidders to packages. Specifically, we let x_i to denote the (possibly empty) package that is allocated to bidder i . We will also assume that every bidder places a bid on the empty package, or the null package, with zero price. An allocation is feasible if every unit is assigned to exactly one bidder, and we let $X(S)$ denote the set of all feasible allocations for a given set of bidders S . Finally, we let \mathcal{A} denote the set of all possible packages out of the units in U and $\mathcal{A}_0 := \mathcal{A} \cup \{\emptyset\}$ be the set of all packages including the null package. For notational simplicity, we will use $N_{-i} := N \setminus \{i\}$ to denote the set of bidders excluding a particular bidder i .

Outcome. Given the set of bidders N and their reported bids, an outcome of the VCG mechanism is described by (i) an allocation; and (ii) payments to bidders. The allocation from the VCG mechanism is determined by solving for the following optimization problem:

$$\begin{aligned} & \text{minimize} && \sum_{j \in N} b_{x_j}^j && (4.1) \\ & \text{subject to} && x \in X(N), \quad x_j \in \mathcal{A}_0, \forall j \in N. \end{aligned}$$

Accordingly, we let x^* to denote the VCG allocation, that is:

$$x^* := \operatorname{argmin}_{x \in X(N)} \sum_{j \in N} b_{x_j}^j.$$

Note that this optimization problem can also be formulated as an integer program that is described in Section 2.2.1, where the definition of the allocation vector is slightly different. In fact, the allocation vector defined in this chapter leads to simpler notation for the theoretical analysis that we provide in Section 4.4.

The VCG payment made to bidder i , denoted by p_i , is defined as follows:

$$p_i := \min_{x \in X(N_{-i})} \sum_{j \in N_{-i}} b_{x_j}^j - \sum_{j \in N_{-i}} b_{x_j^*}^j.$$

Given the strategy-proofness of the VCG mechanism, the first term is same as the total cost of the minimum allocation *without* bidder i , whereas the second term is the total cost of the minimum allocation *minus* the winning bid by bidder i . After subtracting the winning bid of bidder i from the second term and adding it back, we get:

$$p_i := \left(\min_{x \in X(N-i)} \sum_{j \in N-i} b_{x_j}^j - \sum_{j \in N} b_{x_j^*}^j \right) + b_{x_j^*}^j,$$

which shows that the payment to this bidder is the summation of her *contribution* to the minimum possible total cost and the amount of her winning bid. Observe that this payment scheme shows how the VCG mechanism generalizes the second price auction. In the second price auction, the winner, who places the minimum bid, gets paid the amount of the second minimum bid, which can be decomposed into two parts: i) the second minimum bid subtracted by the minimum bid (which is the winner's contribution to the minimum possible total cost); and ii) the minimum bid (which is the amount of her winning bid). Again under strategy-proofness of the VCG mechanism, the profit that bidder i makes from the auction, denoted by π_i , is:

$$\pi_i := \min_{x \in X(N-i)} \sum_{j \in N-i} b_{x_j}^j - \sum_{j \in N} b_{x_j^*}^j, \quad (4.2)$$

which is the amount of her contribution to the minimum possible supplying cost. In fact, an important feature of this payment scheme is that the amount a winner receives depends on the bids of other bidders; this is essential for the strategy-proofness. We close this subsection by describing the auctioneer's payoff. First, denote by T the value of the procurement project. It could also be understood that T is the cost of the cheapest outside option for the auctioneer to procure the entire project. Since her total payment (i.e. the procurement cost) is $\sum_{i \in N} p_i$, the profit that the auctioneer (indexed by 0) makes from this auction is:

$$\pi_0 := T - \sum_{i \in N} p_i. \quad (4.3)$$

4.2.2 Connection to Coalitional Games

In the introduction of this chapter, we have discussed that the VCG mechanism may result in a prohibitively high procurement cost, which is one of the reasons why its practical use is so rare. An important question one may ask regarding this issue would be how high it should be to be declared as *prohibitively high*. Notably, Ausubel and Milgrom (2006) adopt *the core* of the transferable utility cooperative game played among the bidders and the buyer (auctioneer), as a competitive benchmark, and show that the VCG outcome may not be in the core, resulting in a high procurement cost. More specifically, it could be understood that if the outcome is not in the core, the payments are so high that there is a group of bidders who can offer a more favorable deal to the auctioneer. In this section, we discuss in more details the connection between the high procurement cost issue and the core of coalitional games. Closely following Day and Raghavan (2007), we begin by describing a formal definition of *the core*; Day and Milgrom (2008) also provide a useful description of this material.

Consider a coalitional game by the set of players $N_0(= N \cup \{0\})$. In a procurement setting, we define the coalitional value function, denoted by w , over a set of players S as follows:

$$w(S) := \begin{cases} \max_{x \in X(S)} \left[T - \sum_{i \in S} b_{x_i}^i \right], & \text{if } 0 \in S, \\ 0, & \text{if } 0 \notin S. \end{cases}$$

This means that the coalitional value function $w(S)$ takes the value of the maximum possible payoff of the auctioneer from the VCG mechanism if the player set S includes the auctioneer (and therefore the trade occurs), and zero otherwise. Given the value function w and the set of players N_0 , the core of a game is then defined as follows:

$$\text{Core}(N_0, w) := \left\{ \pi \geq 0 : (a) \sum_{i \in N_0} \pi_i = w(N_0), (b) \sum_{i \in S} \pi_i \geq w(S) \quad \forall S \subseteq N_0 \right\},$$

where $\pi := (\pi_0, \pi_1, \dots, \pi_n)$ denotes the payoff vector of the players. By definition, the core is the set of feasible and non-negative payoff profiles that are efficient with respect to the reported bids (condition (a)) and unblocked (condition (b)).

In the previous section, we saw that the VCG allocation is the one that minimizes the total reported values, and thus a VCG outcome is always efficient *with respect to the reported bids*.² Therefore, if a VCG outcome is unblocked, then it is in the core. Next, we examine this concept further.

In the context of auctions, if we assume that T is sufficiently large, and bidders' supply profiles are rich enough so that the trade always happens, then to examine the possibility of blocking coalitions, we may only consider the coalitions that contain the auctioneer. Thus, with a slight abuse of notation, we will use $S \subseteq N$ be a set of bidders, and to differentiate, we let $S_0 \subseteq N_0$ to be the coalition that contains the set of bidders S and the auctioneer. Accordingly, to simplify the analysis we define another value function over S , denoted by $v(S)$, as follows:

$$v(S) := \min_{x \in X(S)} \sum_{i \in S} b_{x_i}^i,$$

which corresponds to the minimum total cost that the set of bidders S can offer to the auctioneer. Notice that using the new value function, we can rewrite the payoff to bidder i (equation (4.2)) as $\pi_i = v(N_{-i}) - v(N)$, and also the total procurement cost as:

$$\sum_{i \in N} p_i = \sum_{i \in N} \pi_i + v(N). \quad (4.4)$$

In addition, we have the following relationship between the two value functions: $w(S_0) = T - v(S)$. Using these relations and the new value function, from condition (b) we know that a coalition S_0 *blocks* the current VCG outcome if the following inequality holds:

$$\sum_{i \in S_0} \pi_i < w(S_0).$$

After separating out the auctioneer's profit and applying equation (4.3) to the left-hand side, and also applying the above relationship between the two value functions to the right-hand side, the

²It is important to note that the fact that the outcome is *efficient with respect to the reported bids* does not directly mean allocative efficiency, which measures whether the allocation achieves the minimum possible total supplying cost. In the VCG mechanism, however, the latter is implied by the strategy-proofness.

previous inequality becomes:

$$T - \sum_{i \in N} p_i + \sum_{i \in S} \pi_i < T - v(S).$$

Finally, by canceling out T and rearranging terms after applying equation (4.4) to the left-hand side, we get the following condition:

$$v(S) + \sum_{i \in S} \pi_i < v(N) + \sum_{i \in N} \pi_i. \quad (4.5)$$

Now let us examine the physical meaning of the above inequality. First note that the right-hand side of inequality (4.5) is the total procurement cost in the current VCG outcome. The left-hand side is slightly more complicated. The first term, $v(S)$ is the cheapest possible total reported cost that the set of bidders S can achieve, and the second term is the total payoff that the bidders in set S is making in the *current VCG outcome*. Therefore, the inequality (4.5) means that even after being compensated their current VCG payoffs, the new set of bidders S is able to offer a cheaper procurement cost to the auctioneer than the set of current winners does, hence blocking the current VCG outcome. Note that this is equivalent to:

$$v(S) - v(N) < \sum_{i \in N \setminus S} \pi_i.$$

That is, when the VCG outcome is blocked, there is a group of bidders (S in this case) that have incentives to deviate from it and offer a better deal to the auctioneer. From the auctioneer's perspective, this is the case where the reduction in the procurement cost by giving less total profit to the new set of winners (the right-hand side) is larger than the increase in the total supplying cost, or equivalently the decrease in the allocative efficiency (the left-hand side). Overall, if the VCG outcome is not in the core, and therefore there exists at least one blocking coalition, it can be viewed that the current VCG outcome is *overpaying* to the VCG winners.

4.2.3 Example of a VCG outcome that is not in the Core

In this section, we provide an example of a VCG outcome in a very simple CA setting that is not in the core and illustrate how such an outcome can result in a high procurement cost. We believe

that this example provides insights on how the VCG performance is related to the core. A similar example is provided in Ausubel and Milgrom (2006) for a forward auction setting, and the following example is motivated by that.

Packages	Bidder 1	Bidder 2	Bidder 3
Unit A	\$ 15 M	-	-
Unit B	-	\$ 15 M	-
Bundle A & B	-	-	\$ 40 M
Profit	\$ 10 M	\$ 10 M	-
Payment	\$ 25 M	\$ 25 M	-

Table 4.1 – Summary of bids and the VCG outcome.

Consider the case where there are two units to procure, labeled as units *A* and *B*. Suppose that the auctioneer runs the VCG mechanism to allocate the procurement contracts, and three bidders, labeled as bidders 1, 2, and 3, have participated. The supplying cost for unit *A* alone by bidder 1 is \$ 15 million, and similarly bidder 2 can serve unit *B* at the cost of \$ 15 million. Bidder 3 can only serve both units together at \$ 40 million. They have infinite cost to serve any other units or the bundle, and therefore each of the bidders places only one bid corresponding to the true supplying cost. Table 4.1 summarizes the bidding result and the VCG auction outcome. Obviously, the cheapest possible allocation is assigning unit *A* to bidder 1 and unit *B* to bidder 2, at the total value of \$ 30 million. The cheapest allocation *without* bidder 1 is assigning bundle *AB* to bidder 3 with the total value of \$ 40 million, and therefore the payment to bidder 1 is (i) the value of its winning package (\$ 15 million) plus (ii) the difference in the minimum total costs with and without bidder 1 ($40 - 30 = \$10$ million). Similarly, the payment to bidder 2 is also \$ 25 million. Therefore the total procurement cost is \$ 50 million.

This outcome is not in the core. Note that bidder 3, a non-winner, will be better-off to serve the

two units at any payment between \$ 40 million – \$ 50 million. That is, there is an opportunity to jointly deviate for bidder 3 and the auctioneer at the price of, say, \$ 45 million for the procurement of the two units by bidder 3. This example illustrates that a competitive procurement cost would be \$ 40 million. In fact, this could have been achieved if some other mechanism had been used. For example, if the auctioneer combined the two units together, treating them as one unit, and had run a second-price single-unit auction, and if bidders 1 and 2 formed an alliance to place a bid on the bundle AB at the price of \$ 30 million, the auctioneer could have procured units A and B at the price of \$ 40 million. Hence, the VCG procurement cost can be viewed as an “overpayment” by \$ 10 million. More importantly, this overpayment can be arbitrarily large depending on the cost structure of bidder 3.

4.2.4 Core-selecting Auctions.

There has been an active body of work in recent years to propose auction mechanisms which ensure that the outcome is always in the core so that the mechanism does not suffer from the aforementioned problems of the VCG mechanism. Such auctions are called “Core-selecting Auctions.” It is worth noting that, given a set of placed bids, there can exist multiple outcomes that are in the core. For example, consider an auction with two units (units A, B) and three bidders (bidders 1, 2, 3). Suppose that bidder 1 placed a bid of \$ 10 on the bundle of the two, and bidders 2 and 3 place bids on units A and B, respectively, both at the price of \$ 7. If VCG was used, bidder 1 would win the two units and get paid \$ 14. Since \$ 14 is the best that bidders 2 and 3 can offer, they cannot block the current outcome. Therefore, the current VCG outcome is in the core. At the same time, consider an alternative outcome that also assigns the two units to bidder 1 but pays only \$ 7 instead of \$ 14. For the same reason, it is in the core as well. The latter payment scheme is in fact the first-price rule, which will not lead to truthful bidding. This example illustrates the fact that different payment rules can imply different bidding strategies, even if both produced core outcomes.

Therefore, the main issue in core-selecting auctions is which core outcome to select amongst multiple core outcomes. A core outcome is called *bidder-Pareto optimal* if there is no other core

outcome weakly preferred by every winning bidder in a given allocation, and Day and Raghavan (2007) and Day and Milgrom (2008) propose auctions that find efficient and bidder-Pareto optimal outcomes that are in the core. An attractive property of efficient core-selecting auctions that are also bidder-Pareto optimal is that, they minimize the bidders' incentives to unilaterally misreport true costs among all core-selecting auctions. In fact, they show that when the VCG outcome is in the core, the proposed core-selecting auction will give the same outcome, leading to the efficient allocation. Still, they may not be truthful, and therefore can lead to an inefficient allocation. For example, Goeree and Lien (2009) report instances where such bidder-Pareto optimal core-selecting auctions result in an inefficient allocation. In this chapter, we follow a different direction – we seek to find characteristics of applications which ensure that the VCG outcome is in (or close to) the core, so that it produces competitive procurement cost as well as the allocative efficiency.

4.3 VCG in the Chilean School Meals Program

Recall that the main objective of Chapter 2 was to develop a method to pin down bidders' cost information from observed bidding data in large-scale CAs. This was essential to directly measure the performance of a CA. In Chapter 3, we effectively applied the method to the data from a real world large-scale CA – the Chilean school meals auction. The uncovered cost information was then used to provide the markup information of the winning firms as well as the efficiency of the actual allocation. The results revealed that the first-price CA achieves high allocative efficiency with reasonable margins to the winning firms.

In this section, we conduct a counterfactual analysis to see how the VCG mechanism would perform in a real world application in terms of the procurement cost, and compare it with that of the first-price CA which has been already computed in the preceding chapter. Having estimated the costs of the bidders, we can directly use the estimated costs as the bids that bidders would report in this counterfactual scenario due to the strategy-proofness of the VCG mechanism. We use the same set of cost estimates used in Section 3.5.1 which include the costs estimated through an out-

of-sample extrapolation over the extended package set as well as the costs of relevant bids directly obtained by the structural estimation. We know VCG achieves the efficient allocation, which was previously computed in Section 3.5.1. From the bids, we can compute the individual VCG payments to the winning bidders, and by summing them, we obtain the VCG procurement cost.

We computed the procurement costs of both 2003 and 2005 auctions. The results show that the total procurement costs from the VCG mechanism are very close to those of the first-price CA. As seen in the previous chapter, the total annual procurement cost in the 2003 first-price CA is US\$ 70.5 million. The total annual procurement cost under the VCG mechanism is US\$ 70.3 million, which is about 0.32% cheaper than the first-price CA. Similarly, in 2005 the total annual procurement cost under VCG is computed to be US\$ 53.5 million, which is only 0.23% more expensive than the total procurement cost of US\$ 53.4 million under the first-price CA.

As discussed earlier, the performance of the VCG mechanism is closely related to the core, and we also examined further to see if the VCG outcomes in this counterfactual analysis are in the core or not. Recall that we found significant cost synergies among the units in our application, and as noted in Ausubel and Milgrom (2006) this is a type of environment where the VCG outcome may not be in the core. In our application, we find that indeed the VCG payoffs lie *essentially* in the core, which is consistent with the reasonable total procurement costs achieved by the VCG mechanism. In particular, following Day and Raghavan (2007) we computed the closest point in the core (with respect to the truthful bids) to the VCG payments under a suitable norm (more details of this algorithm are provided in Section 4.6.1). Specifically, we find that the difference in the total procurement costs between these two points is only 0.1% in 2003. Also, individual payments are very similar as well; half of the winners receive exactly the same payments in the core point as in VCG, and the rest receive payments that are no more than 0.7% apart. In 2005, the VCG payments are even closer to the core payments with respect to the truthful bids. The difference of the total procurement costs between these two points is less than 0.03% in 2005. Moreover, the individual payments are also closer; two-thirds of the nine winners receive exactly the same payments in the core point as in VCG and the rest three receive payments that are no more than 0.7% apart.

Our results are at odds with the theoretical literature mentioned in Section 4.1 describing the pitfalls of VCG; in our application, VCG achieves payments comparable to the first-price CA and induces a reasonable procurement cost. Moreover, despite the significant cost synergies among the units, the VCG outcomes are found to be essentially in the core. We believe this is a valuable observation. It may provide insights that can help to characterize the environments where the VCG mechanism is expected to perform well. In the remainder of the chapter, we develop our analysis in this direction.

4.4 Analysis

Although the units are significantly complementary in our application, the previous counterfactual analysis shows that the VCG mechanism achieves reasonable procurement costs, which are very close to those of the first-price CA. Being consistent with these results, a further analysis shows that the VCG outcomes in both years are *essentially* in the core. We believe these results are driven by the significant amount of competition introduced by the large number of package bids submitted by firms. In this case, a winning bidder is not that relevant; if her bids are eliminated, there is another allocation whose total supplying cost is close to that of the minimum-cost allocation, leading to reasonably low VCG payments. In contrast, in the example provided by Ausubel and Milgrom (2006) as well as the one in Section 4.2.3, competition is limited, resulting in high VCG payments. We believe that the VCG mechanism should achieve reasonable procurement costs in settings with a reasonable amount of bidders that are able to submit many package bids. The latter should be expected when it is relatively effortless for a bidder to evaluate the costs of many different packages.

These insights motivate the direction of our analysis in this section. In particular, we seek to understand how increased competition affects the VCG outcome, making it close to the core and leading to a competitive procurement cost. However, the main difficulty in such an analysis is that given a set of reported bids, the core cannot be expressed analytically. To bypass this challenge, we examine the asymptotic behavior of the procurement cost of a given VCG auction as the number of

bidders (or equivalently the number of bids) grows. This still provides a meaningful direction. If one can show that the procurement cost decreases sufficiently as the competition increases, and the payoffs to the winning bidders eventually vanish, then it is equivalent to showing that the converging point is in the core. Accordingly, we will focus on providing conditions under which the total payoff to the bidders in a VCG outcome converges to *zero* as the competition increases.

We begin by providing the model assumptions that will be kept throughout the analysis. Basically, we adopt independent private cost paradigm as is common in the auction literature. In particular, we assume that the costs of any given package placed by different bidders are independently and identically distributed. We have argued in Section 3.2.5 how the independent private cost assumption is justified in our application. However, the assumption of identical distribution might be strong – in general there could be significant heterogeneity among bidders which can potentially affect the nature of individual bidders' cost structures. The assumption, however, makes the analysis much simpler and still provide important insights on the relationship between increased competition and the VCG performance. Therefore we will take the assumption as our starting point and discuss about some of its relaxations later in this section.

Second, we assume that the units exhibit complementarity. More specifically, we assume that the costs are sub-additive – there are non-negative cost synergies between any two disjoint packages (including single unit packages). In fact, in our application, we do observe sub-additive costs roughly on most packages as reported in Section 3.4.2. We believe this is a natural assumption because in many multi-unit procurement settings, one can expect cost synergies amongst units and packages, and these are the environment where multi-unit package auctions have potential practical benefits.

To formalize the assumptions above, we first denote by c_a^i the cost of supplying package $a \in \mathcal{A}$ by bidder $i \in N$, and also let $c^i := \{c_a^i\}_{a \in \mathcal{A}}$ be the cost vector containing all costs from bidder i . Now, the following assumption describes the cost structure of the bidders.

Assumption 4.1. a) For all bidders $i \in N$, the random cost vectors, $\{c^i\}_{i \in N}$, are independent and

identically distributed from a joint distribution F that has a density.

b) For any package $a \in \mathcal{A}$, there exist non-negative numbers \underline{c}_a and \bar{c}_a such that the support of the marginal distribution of c_a^i , denoted by F_a , is $[\underline{c}_a, \bar{c}_a]$.

c) For each bidder and for all packages, the costs are sub-additive. That is, we have $c_{s \cup t}^i \leq c_s^i + c_t^i$ for any disjoint packages $s, t \in \mathcal{A}$, for all bidders $i \in N$.

In addition to the notation defined in Section 4.2.1, we will also use the following notation in our analysis. First, we define a partition of a set of units as follows:

Definition. Given a set of units U , a *partition* is a set of nonempty subsets of U (hence a set of packages in \mathcal{A}) such that every unit in U is in exactly one of these subsets (or packages).

Accordingly, we let \mathcal{S} be a set of all possible partitions of U . Notice that each partition $S \in \mathcal{S}$ can be viewed as a “type” of allocations. For example, in a three unit CA (units A, B , and C), a partition $\{A, B\}$ and $\{C\}$ characterizes a type of allocations that the first bundle is assigned to one bidder and the other unit to another. Then we let \underline{S} to denote the lower bound of the total costs arising from the type of allocations that correspond to this particular partition S . That is, $\underline{S} := \sum_{a \in S} \underline{c}_a$. Finally, we let P_N denote the total procurement cost that the set of bidders N generates in a VCG outcome. Note that using the value function defined in Section 4.2.2, the procurement cost can be written as:

$$P_N = \sum_{i \in N} [v(N_{-i}) - v(N)] + v(N).$$

In what follows, we analyze the relationship between the level of competition and the total procurement cost. We do this by looking at the asymptotic behavior of the total procurement cost of a given VCG auction as the number of bidders (or equivalently the number of bids) grows. We provide two different approaches: one that considers unit-wise costs and another that examines allocation-wise costs.

4.4.1 Unit-wise analysis

In this approach, we use the stand-alone bids (i.e. bids that are placed on the individual units) to analyze the total procurement cost in a VCG outcome. Our goal is to characterize the asymptotic properties of the VCG procurement cost by examining the behavior of the stand-alone bids as the number of bidders increases. We achieve this by constructing an upper bound for the VCG procurement cost only using the *second minimum* bids placed on each of the individual units, which requires that there are at least two stand-alone bids on each of the units. Moreover our analysis relies on the fact that the bidders' reported bids are sub-additive. Since sub-additivity of the supplying costs were already assumed in Assumption 4.1, this would be achieved if we assume that bidders place bids on all packages. To ensure this, we first make the following assumption. Later in this section, we will weaken this assumption.

Assumption 4.2. *Every bidder places bids on all units and packages.*

We now begin our *unit-wise analysis* by showing that the total procurement cost of a given VCG auction can be bounded by the “second order statistics” of the stand-alone bids. The following lemma is helpful in constructing such a bound.

Lemma 4.1. *Suppose Assumptions 4.1 and 4.2 hold. In a given VCG auction, any winning bid is the minimum among the bids placed on that particular package. That is, for any winner i and her winning package $x_i^* = a \in \mathcal{A}$, we have:*

$$b_a^i = b_a^{(1)} \equiv \min_{j \in N} b_a^j.$$

The proof of this lemma as well as the proofs of all other lemmas and propositions discussed in this chapter are provided in Section 4.6.2. Basically, the above lemma states that whatever package a bidder wins, the corresponding winning bid should be the minimum bid placed on that particular package. Note that this may not be true in general; the sub-additivity of the costs are essential for this result. For example, consider a two-unit two-bidder CA. Suppose bidder A places stand-alone bids both of which are cheaper than the bids by bidder B. However, if bidder A has significant

negative synergy in the cost of the bundle so that the costs are strictly super-additive, it could be optimal to assign only one unit to bidder A. In this case, bidder B wins the other unit even if her stand-alone bid on that particular unit is not the minimum.

Using this result, we now characterize an upper bound for the VCG procurement cost only using the *second minimum* bids placed on each of the individual units. The following proposition establishes the bound.

Proposition 4.1. *Suppose Assumptions 4.1 and 4.2 hold. Given a set of realized costs of the bidders, the total procurement cost of the VCG mechanism cannot be larger than the sum of the second minimum stand-alone bids. That is,*

$$P_N \leq \sum_{u \in U} b_u^{(2)}. \quad (4.6)$$

Remark. Proposition 4.1 establishes a useful bound that makes the convergence analysis in a random environment possible – it does not require the knowledge on the final VCG allocation, which is crucial for the analysis. In fact, there is another bound that could be tighter than bound (4.6), which is described as follows:

$$P_N = \sum_{i \in N} p_i \leq \sum_{i \in N} b_{x_i^*}^{(2)}.$$

Similar to bound (4.6), the above bound uses the fact that the payment to a winning bidder cannot exceed the second minimum bid placed on the package that the particular winning bidder wins. However, to use this bound, one needs to know the VCG allocation x^* given the placed bids. The difficulty arises by the fact that the winner determination problem of the VCG mechanism (4.1) has no analytical solution, and therefore x^* cannot be characterized *a priori* using the random bids.

Note that as the number of bidders increase, the second order statistic of each stand-alone bid gets closer and closer to the lower bound of its support, forcing the total procurement cost to the sum of the lower bounds of the stand-alone unit costs. Recall that our objective is to find conditions where the total payoff of bidders vanishes as competition increases. Therefore it would be sufficient if the procurement cost eventually approaches to the *lower bound* of the total supplying cost. One

possible problem in this unit-wise approach is that the lower bound for the total supplying cost may vary across allocations, but we are bounding the procurement cost only using the allocation that purely consists of the stand-alone costs. That is, if there exists another allocation that achieves a cheaper lower bound for its total cost than the sum of the lower bounds of the stand-alone costs, then inequality (4.6) itself does not imply that the total payoff of bidders will eventually vanish. For this, we further assume that all the feasible allocations have the same lower bounds – this could arise in an environment where the cost synergies tend to be less significant when a firm’s unit costs get more competitive. In Section 4.4.2, where we conduct the allocation-wise analysis, we relax this assumption. The following assumption formalizes this.

Assumption 4.3. *There exists a non-negative constant \underline{c} such that $\underline{S} = \underline{c}$ for any partition $S \in \mathcal{S}$.*

In real-world applications, a bidder may not attempt to win a certain set of packages, by placing no bids on such packages. It may be because she already knows that she is not particularly competitive on those packages and hence very unlikely to win one of those. Another possible reason why she may not place such bids would be that it is too expensive for her to estimate her own supplying costs on such packages. We note that this phenomenon may cause a problem in our analysis. Recall that we rely on the order statistics of the stand-alone bids to bound the total procurement cost, assuming that all the bidders place stand-alone bids on every unit. That is, our analysis may not work if there are bidders that do not bid on individual units but do only on multi-unit packages. As discussed in Chapter 3, however, we observed some notable bidding patterns in the Chilean school meals auction data. First, a firm’s bidding may be concentrated in some subset of units. Especially, bidders who have local cost advantage tend to focus on the units in which they are competitive, placing stand-alone bids on those local units as well as multi-unit bids on packages that contain those local units. Also, firms may not bid on large packages. In our application, for each bidder a maximum package size was imposed by the auctioneer so that a bidder cannot win packages that exceed her size limit. In other applications, it could also be determined by the bidder depending on her service or production capacity.

Our model assumptions on bidders' bidding behavior is based on these observations. Specifically, we assume that each bidder is interested in winning at most d units, and they randomly select d units equally likely, then submit bids on all possible packages that consist of those units only. Therefore, each bidder will place bids on $(2^d - 1)$ packages, that include the stand-alone bids on the d units. Under this selection scheme, a particular unit will be selected by a bidder with probability $p := d/K$ (again K is the total number of units in the auction). Before we formalize the assumptions on the bidders' bidding behavior and provide the main result, let us first examine an illustrating example which establishes the convergence results in a simple setting.

Example 1. In this example, we show that the expected total procurement cost of the VCG mechanism converges to zero if and only if the average number of interested bidders on each unit np grows to infinity, under the assumptions described so far as well as some additional assumptions – we assume that each of the unit costs follows a uniform distribution, and the lower bounds of the feasible allocations are all $c = 0$. Now for the analysis, we let $B^{(2:m)}$ be the second order statistic out of m i.i.d. observations from $Uni[0, 1]$. If each of the n bidders is randomly selecting d units equally likely, then the (random) number of bids placed on a given unit i , denoted by N_i , follows Binomial distribution with parameters (n, p) , where $p := d/K$. Therefore, we have:

$$\begin{aligned}
\mathbf{E}[P_N] &\leq \sum_{i \in U} \mathbf{E} \left[B^{(2:N_i^{(n,p)})} \right] = K \cdot \mathbf{E} \left[B^{(2:N_1^{(n,p)})} \right], && \text{by bound (4.6)} \\
&= K \cdot \sum_{k=2}^n \mathbf{E} \left[B^{(2:k)} | N_1^{(n,p)} = k \right] \cdot \mathbf{P} \left(N_1^{(n,p)} = k \right), \\
&= K \cdot \sum_{k=2}^n \frac{2}{k+1} \binom{n}{k} p^k (1-p)^{n-k}, \\
&= \frac{2K}{(n+1)p} \sum_{k=2}^n \binom{n+1}{k+1} p^{k+1} (1-p)^{n-k}, \\
&\leq \frac{2K}{(n+1)p} \sum_{k=-1}^n \binom{n+1}{k+1} p^{k+1} (1-p)^{n-k},
\end{aligned}$$

which can be simplified as:

$$\mathbf{E}[P_N] \leq \frac{2K}{p(n+1)} \leq \frac{2K}{np}.$$

This provides sufficiency of the convergence, that is: $np \rightarrow \infty$ implies $\mathbf{E}[P_N] \rightarrow 0$.

Now we turn our attention to the necessity of the convergence. We note the following observation that we had while analysing the estimated costs in our application in Chapter 3. First, the level of synergies tend to depend on the level of the unit costs – when the unit costs are less aggressive, the synergies were more variable and relatively larger on average. Second, when the package costs are very competitive, it was usually the case that the unit costs were also very competitive – and it was less likely that the unit costs are not that competitive but the synergies are significantly large. To be consistent with these observations, we make further assumption on the relationship between the synergy and the unit costs. Specifically, we assume that the maximum possible synergy level depends on the level of unit costs. Formally, we assume that there exists a constant $\alpha \in (0, 1]$ such that for any package $a \in \mathcal{A}$ and for any bidder i , we have $\sum_{u \in a} \alpha(c_u^i - \underline{c}_u) \leq c_a^i - \underline{c}_a$. We now explain the meaning of this assumption. Notice that by definition of the lower bound, we have $\underline{c}_a \leq c_a$ and by sub-additivity of the costs, we have $c_a \leq \sum_{u \in a} c_u$. That is, these two inequalities define the possible synergy level in this package. However, to capture the pattern of the second observation above, we assume that the smallest possible package cost c_a , and hence the maximum possible synergy, also depend on the unit costs $\sum_{u \in a} c_u$. More specifically, we assume that c_a is lower bounded by $\alpha \sum_{u \in a} c_u + (1 - \alpha)\underline{c}_a$. Then, since Assumption 4.3 implies $\underline{c}_a = \sum_{u \in a} \underline{c}_u$, we have $\sum_{u \in a} \alpha(c_u^i - \underline{c}_u) \leq c_a^i - \underline{c}_a$. In the special case of this example where we have set $\underline{c} = 0$, this condition can be simplified to $\sum_{u \in a} \alpha c_u^i \leq c_a^i$.

Now to show the necessity of the above convergence result in this environment, we will use the fact that for any VCG outcome, the total procurement cost is always larger than the total supplying cost. By Lemma 4.1 and the definition of α , the total supplying cost should be at least $\alpha \sum_{i \in U} B_i^{(1: N_i^{(n,p)})}$. Hence, taking expectation and using the fact that B_i 's are i.i.d samples from

$Uni[0, 1]$, the expected total procurement cost is lower-bounded by:

$$\mathbf{E}[P_N] \geq \alpha K \cdot \mathbf{E} \left[B_1^{(1:N_i^{(n,p)})} \right] \geq \frac{\alpha K}{np + 1},$$

where the second inequality is by Jensen using the fact that the function $g(k) = \frac{1}{k+1}$ is convex.

Therefore, if np does not grow to infinity, $\mathbf{E}[P_N]$ cannot vanish.

We are now ready to establish the desired asymptotic results, generalizing the above example. For analytical simplicity, however, we assume further that given n bidders, the number of interested bidders for each unit is deterministic – that is exactly $[np]$. We call p *intensity of the bidders' interests*. Under this assumption together with the assumptions made so far, we show in Theorem 4.1 that the total procurement cost of the VCG mechanism converges to the lower bound \underline{c} in expectation if and only if the number of bidders (who are interested in winning each unit) grows to infinity. Then the corollary that follows establishes the desired result; the convergence of the VCG payoff profile when competition increases. We first formalize the assumptions we discussed so far.

Assumption 4.4. *Given a VCG auction with the set of n bidders N as well as the set of K units U , we assume that:*

- a)** *Each bidder places a bid on each of the units with interest intensity $p \in (0, 1]$. That is, each bidder selects a unit with probability p and place a stand-alone bid on it.*
- b)** *Every package bid placed by a bidder only contains her interested units, which are the units on which she places stand-alone bids.*
- c)** *The selection is balanced. That is, it is done in a way that the number of interested bidders for each of the units is exactly $[np]$.*
- d)** *There exists a constant $\alpha \in (0, 1]$ such that for any package $a \in \mathcal{A}$ and for any bidder i , we have $\sum_{u \in a} \alpha (c_u^i - \underline{c}_u) \leq c_a^i - \underline{c}_a$.*

We now present the main results from our unit-wise analysis.

Theorem 4.1. *Consider a fixed set of units U , and suppose Assumptions 4.1, 4.3, and 4.4 hold. Then the the expected total procurement cost of the VCG mechanism converges to the lower bound*

\underline{c} if and only if the number of bidders n grows to infinity.

The above theorem states that the necessary and sufficient condition for the expected total procurement cost of the VCG mechanism to converge to the lower bound \underline{c} is that for each of the units there should be infinitely many “interested” bidders. The intuition behind this result is as follows. Recall that the profit that a winning bidder makes in the VCG mechanism is same as her contribution to the minimum possible total cost. As the competition increases and therefore as the number of placed bids increases, a particular winning bidder’s contribution decreases – when the bidding is ample on each of the units, even without the specific winning bidder’s bids, we can find another allocation that achieves the total cost also quite close to that of the optimal (efficient) allocation. Our analysis upper-bounds the total cost of such an alternative allocation using the second order statistics, which approaches to the lower bound \underline{c} as the number of bids increases. Note that once the procurement cost approaches to the lower bound of the total cost \underline{c} , the payoff that the winning bidders make will vanish. That is, even in the complementary environment, increased competition can lead to a competitive VCG outcome. The following corollary provides this result.

Corollary 1. *Suppose assumptions 4.1, 4.2, and 4.4 hold. Then the total VCG payoff that is given to the winning bidders converges to zero in expectation if and only if the number of bidders increases to infinity.*

Our analysis highlights the impact of competition on the revenue performance of the VCG mechanism. Although the result is intuitive, we believe our analysis provides helpful insights to enhance the practicality of the VCG mechanism. High competition merely measured by the number of bidders may not be enough – in our analysis, ensuring enough competition in each of the unit-wise markets is the key to a competitive revenue performance of the VCG mechanism.

4.4.2 Allocation-wise analysis

In the previous analysis, we used the order statistics of the stand-alone bids on the individual units to bound the total procurement cost. The main challenge in the analysis stems from the fact that

in a VCG outcome it is not feasible to analytically characterize the final allocation using the costs. The unit-wise analysis, which did not require the knowledge of the final allocation, was therefore useful in obtaining the desired asymptotic results. However, we also had to impose some strong assumptions. For example, we had to assume that all the bidders place bids on individual units. Additionally, it was assumed that the lower bounds of all the feasible allocations are identical. These assumptions might not be applicable especially when there are significant cost synergies among the units, and individual unit costs are relatively very large. Moreover, we also assumed that bidders' interest on the auctioned units are spread all over the units equally likely, ensuring that the competition grows unit-wise as the number of bidders increases. This assumption may not capture the cases where the bidders' interest could be concentrated on a subset of units. In this section, we develop another approach in which we do not need to make such assumptions.

Whereas the previous analysis relies on the costs of each unit, in this approach we examine the total cost of each allocation, hence called the *allocation-wise analysis*. Specifically, we consider the set of feasible allocations and examine the properties of the allocations to characterize the asymptotic behavior of the VCG procurement cost. As alluded to above, we will relax the assumption that all the lower bounds of the feasible allocations are identical. For example, consider a simple two-unit CA. When the cost synergy between the two is very strong it is possible that the lower bound of the cost for allocating the bundle to a single bidder (i.e. the lower bound of the cost for supplying the bundle) could be smaller than the lower bound of the cost for the allocation where the two units are given to two distinct bidders (i.e. the sum of the lower bounds of the individual costs for each of the units). To extend our analysis to these types of cost structures, we first consider the allocations that achieve the minimum possible lower bound among all the feasible allocations. The following definition will be useful for our analysis in this direction.

Definition. Given the set of units U and its associated set of feasible allocations \mathcal{S} , the *minimum lower bound* of feasible allocations, denoted by \underline{S}^* , is the minimum possible value of the lower

bounds for the total supplying cost among all feasible allocations. That is,

$$\underline{S}^* := \min_{S \in \mathcal{S}} \underline{S}.$$

We call the allocations that achieve the minimum lower bound *optimal allocations*. The set of all the optimal allocations, denoted by \mathcal{S}^* , can be formally defined as follows:

$$\mathcal{S}^* := \arg \min_{S \in \mathcal{S}} \underline{S}.$$

Using these definitions, we begin our analysis by providing the following lemma which establishes a bound that is applicable on *any* feasible allocation.

Lemma 4.2. *Take any allocation $S \in \mathcal{S}$. Then we have*

$$P_N - \underline{S}^* \leq K \left[\sum_{a \in S} b_a^{(2)} - \underline{S}^* \right].$$

Note that the left-hand side, the total VCG procurement cost subtracted by the minimum lower bound, corresponds to the maximum possible total payoff to the winning bidders in the particular VCG outcome. The above lemma states that this maximum total payoff can be bounded by only using the values of second order statistics of the reported bids placed on a feasible allocation S . This is an important bound since it will hold for any feasible allocation we choose. The main reason why it works for any allocation is that, no matter what the final VCG allocation is, the term $\sum_{a \in S} b_a^{(2)}$ will bound the payment to *one* VCG winner. The right-hand side is multiplied by the number of the units K because it is the maximum number of winners in a VCG outcome. Using this result, we now provide another asymptotic result, stated in the following theorem.

Theorem 4.2. *The total procurement cost converges in probability to the minimum lower bound \underline{S}^* as the number of bidders grows to infinity if and only if there exists a partition $S \in \mathcal{S}^*$ such that the number of bids placed on each of the packages in S also grows to infinity.*

Formally, after letting $n_a(N)$ be the number of bids placed on package a given the set of bidders N , and similarly defining $n_S(N) := \min_{a \in S} n_a(N)$, the above theorem is equivalent to

the following mathematical expression:

$$P_N \xrightarrow{\mathbf{P}} \underline{S}^* \iff \exists S \in \mathcal{S}^* \text{ s.t. } n_S(N) \rightarrow \infty.$$

This asymptotic result also captures the importance of increased competition. Whereas the results from the unit-wise analysis emphasizes the importance of competition at the individual unit level, the results from this allocation-wise analysis highlights the importance of the amount of package bidding. In particular, in an environment where cost synergies are strong, this analysis suggests that having enough package bidding is essential for a reasonable performance of the VCG mechanism.

4.5 Conclusions

In this chapter, we studied the revenue properties of the VCG mechanism. First, we performed a counterfactual analysis on the Chilean school meals auction – based on the cost estimates obtained in Chapter 3, we computed the total procurement cost if the VCG mechanism had been used in stead of the first-price sealed-bid CA. Contrary to the recent theoretical work, the results showed that VCG performs well in that particular application with procurement costs very close to those of the first-price CA. Given its rare practical use as well as the criticism it has received in the literature in terms of the potentially poor revenue performance, we believe it is an important observation to report the nice revenue performance of the VCG auction in a real world application setting.

Motivated by this observation, this chapter focuses on the revenue properties of the VCG mechanism, addressing such apparent paradox between the theory and practice. More specifically, we examine the impact of competition on the revenue performance of the VCG mechanism using an asymptotic analysis. The main insight that our analysis provides is that the first order impact would be the competition measured by the amount of bids rather than measured simply by the number of bidders; our results emphasize that the VCG mechanism is expected to perform better when the bidders' interests are not limited to a small subset of units, leading to high unit-wise competition, and when they place ample combination bids that contain such units they are interested in.

We believe the findings in this chapter add useful insights that could enhance the practical usage of the VCG mechanism. In many practical applications, such as spectrum rights or transportation procurements, it could be expensive for bidders to correctly estimate their own valuations on the combinations of units. In such an environment, it might be the case that bidders' interests are restricted only to a small subset of units, and more importantly, they are discouraged from placing many combination bids. In the school meals procurement setting, on the contrary, it is relatively straightforward for bidders to estimate their supplying costs, and they were able to place a large number of bids over most of the units in the auction. This scenario is precisely one in which the results in this chapter suggest that VCG should perform well.

Our analysis opens interesting future directions. Throughout the analysis, we kept the number of units fixed, purely measuring the impact of the increased bidders, and our results highlight the importance of the amount of bids on the performance. However, in a multi-unit setting, the number of possible bids per bidder also increases as the number of units increases, and therefore is expected to have a potential impact on the performance as well. For this reason, the simultaneous consideration of increased number of bidders and scale of the auction seems to provide useful characterization of the bidding environment where the VCG mechanism is expected to perform well. Another interesting future direction points to the characterization of the environments where the VCG outcome always lies in the core. Our results are limited to the asymptotic case; we show that the VCG outcome converges to a point that lies in the core. We believe finding the properties of environments where the VCG outcome is always in the core, especially when the costs of the units exhibit complementarity, will be a very important contribution to the literature as well as practice.

4.6 Appendix for Chapter 4

4.6.1 Finding Core Payments

In Section 4.3, we reported that the VCG outcomes in our counterfactual analysis lie essentially in the core, meaning that the payment vectors are very close to core payments – the individual payments are no more than 0.7 % apart in both 2003 and 2005 auctions. In this section, we summarize the algorithm by Day and Raghavan (2007) that we used to compute core payments.

Given the efficient allocation (i.e. the VCG winners and their winning packages), the objective is to find a payment vector that is i) in the core; ii) and close to the VCG payment vector. Starting from the VCG payment vector, we will adjust the payment vector iteratively until no blocking coalition is found. Specifically, in each iteration we first check if there is a blocking coalition, then find a new payment vector to the *efficient winners* adjusted to satisfy the core constraint that was violated in the current payment scheme. We now formally describe this procedure.

We first provide some useful notation. Consider the coalition of winning bidders in the current VCG allocation. We let W to denote the set of the efficient winners (i.e. the winning bidders in the current VCG allocation) and π^{VCG} to denote the vector of VCG payments to these winners. At the start of iteration t , we let π^t be the current payment vector for the efficient winners W . If this outcome is not in the core, there exists a coalition of bidders W' and a payment vector π' that can block the current outcome. For each of the efficient winners $f \in W$, let γ_f be the indicator whether bidder f joins this (blocking) coalition, taking the value of one if it does and zero otherwise. Finally, we denote by a_f the winning package of bidder $f \in W$ in the VCG outcome and also denote by a'_f the winning package of bidder $f \in W'$ in the blocking outcome. Later we will let W^t to denote the coalition found at iteration t that blocks the current payment outcome π^t .

Conditions on the Payments to the Blocking Coalition. In each iteration, given the current payment vector π^t to the efficient winners, we seek to find a core constraint that is *most-violated*. A core constraint is an inequality comparing the total payment to the bidders so that the new coalition

will not block the current outcome. Hence, given a new coalition, we first compute the minimum level of payment to each of the bidders so that they indeed join the coalition. This *minimum* payment vector will be used to see if the new coalition can violate the associated core constraint. Hence, let us first explain how to find this minimum payment vector.

Consider a coalition of bidders W' and its associated allocation. Now we first look at the bidders that belong to the set $W' \setminus W$. Because such a bidder currently makes zero profit in the efficient outcome, she is willing to join the coalition W' as long as her profit is non-negative. Therefore, the minimum level of payment to bidder $f \in W' \setminus W$ so that she joins the coalition W' will be $\pi'_f = b_{a'_f, f}$, which is the true cost of her winning bid a'_f in the new allocation. We now turn our attention to the bidders in $W \cap W'$. Note that such a bidder is only willing to join the coalition W' if the new profit is at least as much as her current profit. Therefore, the minimum level of payment to bidder $f \in W \cap W'$ so that she joins the coalition W' will be $\pi'_f = b_{a'_f, f} + (\pi_f^t - b_{a_f, f})$, where the second term accounts for her current profit. In Day and Raghavan (2007), these minimum requirements for the payments to a blocking coalition are referred to as “coalitional contribution.” Similarly, from the auctioneer’s perspective, he will only be interested in the new coalition if he pays at most the current total payment. Therefore, combining all these observations, the core constraint associated with this new coalition W' can be expressed as follows:

$$\pi_0^t = \sum_{f \in W} \pi_f^t \leq \pi'_0 = \sum_{f \in W'} \pi'_f = \sum_{f \in W'} b_{a'_f, f} + \sum_{f \in W \cap W'} (\pi_f^t - b_{a_f, f}).$$

Each iteration consists of two steps. In the first step, we solve an integer program to find a core constraint that is most-violated at the current payment vector. If there is no violating core constraints found, then we are at a core outcome. Otherwise, we proceed to the second step, where we solve a linear program to adjust the payment vector so that the violated core constraints found so far are all satisfied and the total payment is as close to the VCG procurement cost as possible. We provide the details of the two steps next.

Step 1: Core-constraint separation problem (SEP). The main idea of SEP is to see whether we can find a coalition W' and its associated allocation that results in π'_0 that is *strictly* less than the current procurement cost π_0^t through an integer program. If so, then the new coalition with hypothetical payments π' blocks the current outcome. The objective of SEP is to find the most-violated core constraint, hence it minimizes the total payment in a new coalition. Using the notation described in Sections 2.2.1, SEP is formulated as follows:

$$\begin{aligned}
z(\pi^t) = \text{minimize} \quad & \sum_{f \in F} \sum_{a \in \mathcal{A}_f} b_{af} x_{af} + \sum_{f \in W} (\pi_f^t - b_{af,f}) \gamma_f \\
\text{subject to} \quad & \sum_{f \in F} \sum_{a \in \mathcal{A}_f: i \in a} x_{af} \geq 1, \quad \forall i \in U \\
& \sum_{a \in \mathcal{A}_f} x_{af} \leq 1, \quad \forall f \in F \setminus W \\
& \sum_{a \in \mathcal{A}_f} x_{af} \leq \gamma_f, \quad \forall f \in W \\
& x \in X, x_{af}, \gamma_f \in \{0, 1\}, \quad \forall a \in \mathcal{A}, f \in F
\end{aligned}$$

where X is the set of all feasible allocation satisfying all the allocative constraints such as market share and competition constraints used in our application.

The optimal objective value, $z(\pi^t)$, is the minimal procurement cost π'_0 we can find using the reported bids and the current payment vector π^t . If $z(\pi^t) = \pi_0^t$, then there is no possible blocking coalition with the current payment vector π^t , indicating that the current outcome is in the core. Therefore, the algorithm terminates with an outcome that is in the core. If $z(\pi^t) < \pi_0^t$, however, we just found a new payment vector π' where all the bidders in the new coalition $W' := W^t$ and the auctioneer become better-off compared to the current outcome. This means that the associated core constraint $\pi_0^t \leq z(\pi^t)$ is violated, blocking the current outcome. In this case, we proceed further to find a new set of payments to the efficient winners so that the violated core constraint that we found in this iteration is satisfied. Next we describe how to adjust the payment vector for this purpose.

Step 2: Finding New Payments. Finding bidder-Pareto-Optimal core payments can be done by solving a linear program, which is referred to as BPO in Day and Raghavan (2007). The main

idea is as follows. If there is a blocking coalition, an efficient winner who is *not* in the blocking coalition W^t can be viewed as an overpaid winner since they can be replaced by some other bidders to achieve a cheaper procurement cost. So we adjust their payments. Let us call them “expensive” winners. Also note that the objective value of SEP, $z(\pi^t)$, is linear in the payments to the winners who remain in the new coalition W^t , referred to as non-expensive winners. Therefore, we only need to consider the total value $z(\pi^t)$ subtracted by the total money given to these non-expensive winners to bound the payments to the expensive winners. This gives one inequality that the new payment should satisfy regarding the payments to the expensive winners. Since this only considers the blocking coalition W^t found at the current iteration t , to find a new payment we need to consider all such constraints found so far. The following linear program formulates BPO.

$$\begin{aligned}
& \text{minimize} && \sum_{f \in W} \pi_f \\
& \text{subject to} && \sum_{f \in W \setminus W^s} \pi_f \leq z(\pi^s) - \sum_{f \in W \cap W^s} \pi_f^s, \quad \forall s \leq t. \\
& && b_{a_f, f} \leq \pi_f \leq \pi_f^{VCG}.
\end{aligned}$$

Here we try to find a payment vector that maximizes the total payments (but bounded from above by the VCG payments) because it will minimize the incentive of bidders not to report truthfully. The solution to this problem then becomes our new candidate payment vector π^{t+1} and we return to SEP to check if the new one is in the core.

4.6.2 Proofs

Proof of Lemma 4.1. Fix a winning bidder i and let a be the package that bidder i wins. Suppose, for the sake of contradiction, that b_a^i is not the smallest among the bids on a . Then the minimum bid $b_a^{(1)}$ must be placed either by a non-winner or another winner. If it is placed by a non-winner, it contradicts the optimality of current allocation – replacing bidder i with this particular non-winner to assign package a in current allocation forms a feasible allocation that results in a cheaper total cost. If the minimum bid is from a winner other than bidder i , say bidder j with her current winning

package a' , then by the sub-additivity of the supplying costs described in Assumption 4.1, we have $b_{a \cup a'}^j \leq b_a^j + b_{a'}^j < b_a^i + b_{a'}^j$ since the supplying costs will be the bidders' reported bids in VCG. Therefore, assigning the package $a \cup a'$ to bidder j while the rest assignments remain unchanged will result in a cheaper total cost than the current allocation, which also contradicts the optimality of the current allocation. Hence, b_a^i should be the minimum bid on a , which completes the proof. ■

Proof of Proposition 4.1. Choose any winner i . Given the realized costs, let a be her winning package in the VCG auction, that is x_i^* . Since the choice of winner i was arbitrary, it suffices to show $p_i \leq \sum_{u \in a} b_u^{(2)}$. Equivalently, if we show $v(N_{-i}) \leq v(N) + \sum_{u \in a} b_u^{(2)} - b_a^i$ instead, the proof will also be complete by the relationship: $p_i - b_a^i = \pi_i = v(N_{-i}) - v(N)$. Therefore we will show the latter inequality. We show it by considering an alternative allocation without bidder i , constructed as follows: for each unit $u \in a$, assign it to a bidder whose stand-alone bid for that particular unit is the minimal when all the bids by bidder i are excluded. For example, if bidder i 's bid is the minimum bid placed on unit u , then assign it to the bidder who placed the second minimum stand-alone bid on unit u . If not, assign it to the bidder who placed the minimum bid. In this way, the units in package a may be assigned to multiple bidders possibly including some of the currently winning bidders. It is clear that this alternative allocation is feasible and does not involve bidder i . Notice further that the total (reported) cost of this *new* allocation, denoted by TC , will be:

$$TC \leq \sum_{j \in N_{-i}} b_{x_j^*}^j + \sum_{u \in a} b_u^{(2)}. \quad (4.7)$$

This inequality is by the sub-additivity of the supplying costs (again truthful bidding is assumed). To verify inequality (4.7), first note that a winner in this alternative allocation can be categorized into three types: 1) a bidder who wins a package in the VCG allocation only; 2) a bidder who wins some of the units in package a only in the alternative allocation; and 3) a bidder who wins a package in the VCG allocation and who also wins additional units in package a in the alternative allocation. The cost of the winning package in the alternative allocation for a type 1) winner is just her winning bid in the VCG allocation – this is a part of the first term on the right-hand side. The cost of the winning package in the alternative allocation for a type 2) winner is at most the sum of

the individual costs on the units she wins, by sub-additivity of the supplying costs. Since the units that a type 2) winner wins are all in package a , and initially we chose this particular winner such that her stand-alone bids on these units are the minimum bids when bidder i 's bids were excluded, her supplying cost on the package she wins is upper bounded by the sum of the second minimum bids on the units in that particular package. This accounts for a part of the second term on the right-hand side. Finally, for a type 3) winner, her winning package in the new allocation contains a package she previously won in the VCG allocation and some of the units in package a . Like the type 1) case, this winner's reported cost on the package she won in the VCG allocation is just her winning bid in that allocation. Similar to type 2) case, this bidder's reported cost on the units in a that she wins in the alternative allocation is at most the sum of the second minimum stand-alone bids on these units. Therefore, again by the sub-additivity of the supplying costs, the cost of the package that she wins in this alternative allocation is upper bounded by the sum of those two quantities.

At the same time, we can also bound TC from below. Note that TC represents the total reported cost of a feasible allocation where bidder i does not win any package. Therefore, by the optimality in the value function $v(\cdot)$ we have, $v(N_{-i}) \leq TC$. Combining this with inequality (4.7), we get:

$$v(N_{-i}) \leq \sum_{j \in N_{-i}} b_{x_j^*}^j + \sum_{u \in a} b_u^{(2)}.$$

Using the definition of $v(N)$ and the fact that $a = x_i^*$, we then have:

$$\begin{aligned} v(N_{-i}) &\leq \sum_{j \in N_{-i}} b_{x_j^*}^j + b_a^i + \sum_{u \in a} b_u^{(2)} - b_a^i \\ &= v(N) + \sum_{u \in a} b_u^{(2)} - b_a^i, \end{aligned}$$

which establishes the desired inequality and therefore completes the proof. ■

For the proof of Theorem 4.1, the following lemma is useful.

Lemma 4.6.1 (Corollary 3 in Albright and Derman (1972)). *Let $\{X_i\}_{i=1}^n$ be n i.i.d samples from a distribution F which has a finite mean and is absolutely continuous with density f . Then the numbers for the associated stochastic sequential assignment problem, $\{a_{i,n+1}\}_{i=1}^n$, which are defined in*

Albright and Derman (1972), has the following asymptotic property: For any fixed i , $1 \leq i \leq n$,

$$\lim_{n \rightarrow \infty} a_{i,n+1} = \sup\{x : F(x) = 0\}.$$

Proof of Theorem 4.1. (Sufficiency): First, we let $b_u^{(j:m)}$ denote the j^{th} smallest bid price out of m bids placed on unit u . By Proposition 4.1 and part c) of Assumption 4.4, we have $\mathbf{E}[P_N] \leq \sum_{u \in U} \mathbf{E} \left[b_u^{(2:[np])} \right]$. Also, by the payment rule of the VCG mechanism, it is clear that $\mathbf{E}[P_N] \geq \underline{c}$. Since Assumption 4.3 implies that $\underline{c} = \sum_{u \in U} \underline{c}_u$, we have:

$$\sum_{u \in U} \underline{c}_u \leq \mathbf{E}[P_N] \leq \sum_{u \in U} \mathbf{E} \left[b_u^{(2:[np])} \right].$$

Therefore, it suffices to show that $\lim_{n \rightarrow \infty} \mathbf{E} \left[b_u^{(2:[np])} \right] = \underline{c}_u$ for all $u \in U$. For that, choose an arbitrary unit $u \in U$, and consider a stochastic sequential assignment problem with the stand-alone bids placed on this unit, $\{b_u^j\}_{j=1}^{[np]}$, as the arriving values for the job. By the meaning of the quantities $a_{1,n+1}$ and $a_{2,n+1}$ as described in Albright and Derman (1972), we have that:

$$\mathbf{E} \left[b_u^{(1:[np])} + b_u^{(2:[np])} \right] \leq a_{1,[np]+1} + a_{2,[np]+1}.$$

(This is implied by the fact that the expected performance of the assignment under incomplete information is inferior to the performance from the full information assignment.) From this inequality together with the fact that $\mathbf{E} \left[b_u^{(1:[np])} \right] \geq \underline{c}_u$ and $\mathbf{E} \left[b_u^{(2:[np])} \right] \geq \underline{c}_u$, Lemma 4.6.1 then implies that:

$$\lim_{n \rightarrow \infty} \mathbf{E} \left[b_u^{(2:[np])} \right] = \underline{c}_u.$$

Since the choice of unit u was arbitrary, the proof is now complete.

(Necessity): First, note that the payoff that each bidder makes in the VCG mechanism is always non-negative, and therefore $\mathbf{E}[P_N] \geq \mathbf{E}[v(N)]$. Given a set of realized costs, the minimum supplying cost will be: $v(N) = \sum_{j \in N} b_{x_j^*}^j \geq \sum_{j \in N} \alpha \sum_{u \in x_j^*} (b_u^j - \underline{c}_u) + \underline{c} \geq \alpha \sum_{u \in U} \left(b_u^{(1:[np])} - \underline{c}_u \right) + \underline{c}$, where the first inequality is by part d) of Assumption 4.4, and the second inequality comes from the fact that $b_u^{(1:[np])}$ is the minimum bid among all the bids placed on unit u . Since this is true for any realization, we get:

$$\mathbf{E}[P_N] \geq \mathbf{E}[v(N)] \geq \alpha \sum_{u \in U} \left(\mathbf{E} \left[b_u^{(1:[np])} \right] - \underline{c}_u \right) + \underline{c},$$

and therefore, we have:

$$\mathbf{E}[P_N] - \underline{c} \geq \alpha \sum_{u \in U} \left(\mathbf{E} \left[b_u^{(1:[np])} \right] - \underline{c}_u \right)$$

Hence, it suffices to show that the right-hand side cannot vanish if n is finite. For that, first suppose that there exist a number $M \in \mathbb{N}$ such that $[np] \leq M$. Then by definition of order statistics, it is clear that:

$$\mathbf{E}[P_N] - \underline{c} \geq \alpha \sum_{u \in U} \left(\mathbf{E} \left[b_u^{(1:[np])} \right] - \underline{c}_u \right) \geq \alpha \sum_{u \in U} \left(\mathbf{E} \left[b_u^{(1:M)} \right] - \underline{c}_u \right).$$

So we will show that the right-hand side of above inequality is strictly positive for any n such that $[np] \leq M$. First note that, by part a) of Assumption 4.1, the distribution of the first order statistic $b_u^{(1:M)}$ should have a density. In addition, by part b) of Assumption 4.1, the support of the distribution for $b_u^{(1:M)}$ is also $[\underline{c}_u, \bar{c}_u]$. Therefore, for any $\epsilon \in (0, 1)$ we can find $c(\epsilon) \in (\underline{c}_u, \bar{c}_u)$ such that $1 - F_u^{(1:M)}(c(\epsilon)) = \epsilon$, where $F_u^{(1:M)}$ is the distribution of $b_u^{(1:M)}$. But then $\mathbf{E} \left[b_u^{(1:M)} \right] = \int_{\underline{c}_u}^{\bar{c}_u} (1 - F_u^{(1:M)}(t)) dt \geq \int_{\underline{c}_u}^{c(\epsilon)} \epsilon dt = \epsilon(c(\epsilon) - \underline{c}_u) > 0$. Since this is true for any $u \in U$, the right-hand side of the above inequality is strictly positive if n is finite, and the proof is now complete. ■

Proof of Lemma 4.2. Let $v' := \sum_{a \in S} b_a^{(2)}$ and fix a winner i . We first show $v(N - i) \leq v'$. Note that for any package $a \in S$, one of the bids $b_a^{(1)}$ and $b_a^{(2)}$ is placed by a bidder other than i . Now we choose a feasible allocation as follows: for each package $a \in S$, if $b_a^{(1)}$ is placed by a bidder other than i , assign package a to this bidder. If $b_a^{(1)}$ is placed by bidder i , then assign package a to the bidder who placed $b_a^{(2)}$. Then the total supplying cost of this new allocation is less than v' by the sub-additivity of the costs. At the same time, since this is a feasible allocation with bidder i winning nothing, the total supplying cost from this new allocation is at least $v(N - i)$ by the optimality of the value function $v(\cdot)$. Hence $v(N - i) \leq v'$, which then implies $\pi_i = v(N - i) - v(N) \leq v' - v(N)$. Since this is true for any winner i and there can be at most K winners, we get $P_N = \sum_{i \in N} \pi_i + v(N) \leq K(v' - v(N)) + v(N)$. Finally, using $v(N) \geq \underline{S}^*$, we get the desired result. ■

For the proof of Theorem 4.2, the following lemma is useful.

Lemma 4.6.2. *For any package $a \in \mathcal{A}$, we have $b_a^{(2)} \xrightarrow{\mathbf{P}} \underline{c}_a$. Moreover for any partition $S \in \mathcal{S}$ we have $\sum_{a \in S} b_a^{(2)} \xrightarrow{\mathbf{P}} \underline{S}$.*

In the proof of Theorem 4.1 (Sufficiency part), we have shown that the second order static converges in \mathbf{L}_1 to a finite lower bound. Lemma 4.6.2 described above is a direct consequence of that result as convergence in the mean implies convergence in probability. The second statement comes from the fact that convergence in probability is preserved through a finite addition of random variables.

Proof of Theorem 4.2. (Sufficiency): Pick any $S \in \mathcal{S}^*$. By Lemma 4.2, we have $P_N - \underline{S}^* \leq K[\sum_{a \in S} b_a^{(2)} - \underline{S}^*]$. Moreover, it is clear by definition that $P_N - \underline{S}^* \geq 0$. Therefore, for any $\epsilon > 0$, $\{|P_N - \underline{S}^*| > \epsilon\}$ implies $\{K|\sum_{a \in S} b_a^{(2:N)} - \underline{S}^*| > \epsilon\}$. Hence, $\mathbf{P}(|P_N - \underline{S}^*| > \epsilon) \leq \mathbf{P}(K|\sum_{a \in S} b_a^{(2:N)} - \underline{S}^*| > \epsilon) = \mathbf{P}(|\sum_{a \in S} [b_a^{(2:N)} - \underline{c}_a]| > \epsilon/K)$. Then by Lemma 4.6.2, we have the right-hand side of the inequality above converges to zero, which completes the proof for the sufficiency.

(Necessity): Suppose that there exists a constant $M \in \mathbb{N}$ such that $n_S(N) < M$ for all $S \in \mathcal{S}^*$ for any set of bidders N . Then for each partition $S \in \mathcal{S}^*$ we can find a package $a(S) \in S$ such that $n_{a(S)}(N) < M$ for any N . Now fix $\epsilon > 0$, and consider $\mathbf{P}(P_N - \underline{S}^* > \epsilon)$. Since $\mathbf{P}(P_N - \underline{S}^* > \epsilon) \geq \mathbf{P}(v(N) - \underline{S}^* > \epsilon)$, we will show that $\mathbf{P}(v(N) - \underline{S}^* > \epsilon)$ does not converge to zero for some $\epsilon > 0$ if there exists such a constant M .

First, let $\underline{S}' := \min_{S \in \mathcal{S} \setminus \mathcal{S}^*} \underline{S}$ if $\mathcal{S} \neq \mathcal{S}^*$, otherwise set $\underline{S}' := +\infty$. It is clear that the only possible allocations that can achieve a total supplying cost in $[\underline{S}^*, \underline{S}')$ are those in \mathcal{S}^* . Thus, we pick $\epsilon \in (0, \underline{S}' - \underline{S}^*)$, and restrict our attention to the partitions in \mathcal{S}^* only. Now, we define $v_S(N)$ be the value function of the minimum supplying cost by only considering allocations corresponding to partition S . Then, $\mathbf{P}(\min_{S \in \mathcal{S}^*} v_S(N) > \epsilon + \underline{S}^*) = \mathbf{P}(v(N) > \epsilon + \underline{S}^*)$. Moreover, for each S , $\{b_{a(S)}^{(1:N)} > \epsilon + \underline{c}_{a(S)}\}$ implies $\{v_S(N) > \epsilon + \underline{S}^*\}$. Therefore, we get:

$$\mathbf{P}(v(N) > \epsilon + \underline{S}^*) = \mathbf{P}(v_S(N) > \epsilon + \underline{S}^*, \forall S \in \mathcal{S}^*) \geq \mathbf{P}(b_{a(S)}^{(1:N)} > \epsilon + \underline{c}_{a(S)}, \forall S \in \mathcal{S}^*).$$

We now fix a bidder i , and consider $\mathbf{P}(b_{a(S)}^i > \epsilon + \underline{c}_{a(S)}, \forall S \in \mathcal{S}^*)$. This bidder may place bid on all of such packages $a(S)$'s, only some of them, or none of them. And the probability is minimum if she places bids on all of them. Let $p(\epsilon)$ be the value of this probability when bidder i places bids on all of such $a(S)$'s for all $S \in \mathcal{S}^*$. By Assumption 4.1, this probability should be strictly positive as long as $\epsilon > 0$. That is, $\mathbf{P}(b_{a(S)}^i > \epsilon + \underline{c}_{a(S)}, \forall S \in \mathcal{S}^*) \geq \underline{p}(\epsilon) > 0$, for any $\epsilon \in (0, \underline{S}' - \underline{S}^*)$. Finally, notice that the number of bidders who place bids at least one of those packages $a(S)$ is bounded by $M \times |\mathcal{S}^*|$ which is also finite given that U is fixed. Therefore when the number of potential bidders in N is larger than $M \times |\mathcal{S}^*|$, the probability $\mathbf{P}(b_{a(S)}^{(1:N)} > \epsilon + \underline{c}_{a(S)}, \forall S \in \mathcal{S}^*)$ should be no smaller than the probability that there are exactly $M \times |\mathcal{S}^*|$ number of bidders who place bids on all of the packages $a(S)$'s which are all larger than the lower bound of the packages $a(S)$'s by ϵ . Hence by the independence of costs across the bidders we get: $\mathbf{P}(v(N) > \epsilon + \underline{S}^*) \geq \mathbf{P}(b_{a(S)}^{(1:N)} > \epsilon + \underline{c}_{a(S)}, \forall S \in \mathcal{S}^*) \geq \underline{p}(\epsilon)^{M \times |\mathcal{S}^*|} > 0$. Since this is true for any $N \geq M \times |\mathcal{S}^*|$, we get $\liminf_{N \rightarrow \infty} \mathbf{P}(v(N) > \epsilon + \underline{S}^*) \geq \underline{p}(\epsilon)^{M \times |\mathcal{S}^*|} > 0$. Hence P_N cannot converge to \underline{S}^* if there exists such a constant M , and the proof is complete for the necessity. ■

Chapter 5

General Conclusions

This dissertation develops empirical and theoretical methods to analyze the performance of CAs in procurement settings. In Chapter 2, we propose an empirical approach to evaluate the revenue and efficiency performance of first-price CAs. One important feature of our method is that it can be applicable to large-scale CAs which are frequently found in practice. In Chapter 3, we apply our method to the bidding data from a real-world application. In particular, our results show that the first-price CA in that application performs well, achieving high allocative efficiency with a reasonable procurement cost. In addition, via a counterfactual analysis, we found that the VCG mechanism, if it had been used in that application, would also perform well in terms of the procurement cost. This observation, which is contrary to theoretical predictions in the literature, motivated an analytical investigation of the VCG mechanism, and in Chapter 4, we examine the impact of competition on the revenue performance of the VCG mechanism.

The main objective of Chapter 2 is to develop a structural estimation approach for large-scale first-price CAs. An important methodological contribution of our work is that we propose methods to overcome the curse of dimensionality that may arise when dealing with bids from large-scale CAs. First, the standard structural approach followed by CP and GPV involves estimating the distribution of competitors' bids, which are high dimensional random vectors in large-scale CAs. To address the complexity arising in this step, we propose a parametric approach to model the

competitors' bid distribution. We believe that our model provides a parsimonious, yet flexible parametric description of the distributions for the competitors' bids which can be especially useful in CAs that involve geographically dispersed and heterogeneous units that are subject to discounts due to scale and density. Second, high dimensionality of the winning probability vectors, which have to be estimated, may also bring a severe computational challenge for large-scale CAs. We bypass this issue by introducing a restricted markup model in which bidders are assumed to determine their markups based on a reduced set of package characteristics. With this simplification, the first-order conditions of the bidder's problem become computationally and econometrically tractable. We impose reasonable restrictions to the structure of the markups that reduce the complexity of the bidders problem but still provide sufficient flexibility to capture strategic behavior that can undermine the performance of a CA.

We hope that the methods we propose in this work will provide researchers and practitioners a useful tool to analyze auction data from large-scale CAs in many different contexts. Many of the CAs used in practice are indeed in large-scale, and therefore their econometric analysis has been limited due to the complexity. Even though our method may need modification to accommodate different pricing and auction rules in those different settings, we believe that our idea of using characteristic-based pricing as well as our parametric model of bid prices can be a useful starting point to reduce the complexity of econometric analysis in such large-scale settings.

We apply our structural estimation approach to the large-scale Chilean school meals CA in Chapter 3. We find that cost synergies in this auction are significant, and the current CA mechanism, which allows firms to express these synergies through package bidding, seems appropriate. In particular, the current CA achieves high allocative efficiency with a reasonable procurement cost. We believe that this is the first empirical analysis documenting that a CA performs well in a real-world application. In addition, based on the cost estimates we also performed an interesting counterfactual analysis – we computed the total procurement cost if the VCG mechanism had been used in this particular application in stead of the first-price CA. Contrary to the recent theoretical work, the results show that the VCG mechanism also performs well in this particular application –

the procurement cost of the VCG mechanism were very close to that of the first-price CA. Given the criticism that the VCG mechanism has received in the literature in terms of the potentially poor revenue performance, we believe it is an interesting and important observation to report.

Motivated from this observation, in Chapter 4 we address such apparent paradox between the theory and practice. Specifically, we study the impact of competition on the revenue performance of the VCG mechanism using an asymptotic analysis. The main insight that our analysis provides is that the first order impact would be the competition measured by the amount of placed bids rather than the one measured by the number of bidders; our results emphasizes that the VCG mechanism is expected to perform better when the bidders' interests are not limited to a small subset of units (so that it has high competition on every unit) and when bidders place ample combination bids that contain such units in which they are interested. The findings in this chapter adds useful insights for the practical usage of the VCG mechanism. In many practical applications, such as spectrum rights or transportation procurements, it could be expensive for bidders to correctly estimate their own valuations on the combinations of units. In such an environment, bidders' interests might be restricted to a small subset of units, and more importantly, they could be discouraged from placing many combination bids. In the school meals procurement setting, on the contrary, it is relatively straightforward for bidders to estimate their supplying costs, and they were able to place a large number of bids over most of the units in the auction. This scenario is precisely one in which the results in this chapter suggest that the VCG mechanism should perform well.

The chapters of this dissertation illustrate the constructive interplay between theoretical analysis and empirical investigation. The structural estimation approach developed in Chapter 2 assumes that the data we observe are generated by a well specified economic model, which is the bidders' profit maximization problem. It shows how the economic theory can be useful in extracting useful information from data that is essential for the performance evaluation of a given CA. In turn, the empirical results using the method then led to an important observation which are at odds in the literature that (theoretically) study the properties of the VCG mechanism, motivating a further analytical investigation on the revenue performance of the VCG mechanism. More broadly, our

results highlight the importance of the simultaneous consideration of the suppliers' operational cost structure and their strategic behavior for the successful design of a CA. In this way, we hope that this research agenda enhances the understanding of the performance of CAs and thereby provide insights to improve their design.

Bibliography

- Aksin-Karaesmen, Z., B. Ata, S. Emadi, C. Su. 2013. Structural estimation of callers' delay sensitivity in call centers. *Management Science* **59**(12) 2727–2746.
- Albright, C., C. Derman. 1972. Asymptotic optimal policies for the stochastic sequential assignment problem. *Management Science* **19**(1) 46–51.
- Allon, G., A. Federgruen, M. Pierson. 2011. How much is a reduction of your customers' wait worth? an empirical study of the fast-food drive-thru industry based on structural estimation methods. *Manufacturing & Service Operations Management* **13**(4) 489–507.
- An, N., W. Elmaghraby, P. Keskinocak. 2005. Bidding strategies and their impact on revenues in combinatorial auctions. *Journal of Revenue and Pricing Management* **3**(4) 337–357.
- Athey, S., P. A. Haile. 2006. Empirical models of auctions. R. Blundell, W. K. Newey, T. Persson, eds., *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, Volume II*. Cambridge University Press, 1–45.
- Ausubel, L., O. Baranov. 2010. Core-selecting auctions with incomplete information. Working Paper, University of Maryland.
- Ausubel, L., P. Milgrom. 2002. Ascending auctions with package bidding. *Frontiers of Theoretical Economics* **1**(1) Article 1.
- Ausubel, L. M., P. Cramton, R. P. McAfee, J. Mcmillan. 1997. Synergies in wireless telephony: Evidence from the Broadband PCS auctions. *Journal of Economics and Management Strategy* **6**(3) 497–527.
- Ausubel, L. M., P. Milgrom. 2006. The lovely but lonely vickrey auction. P. Cramton, Y. Shoham, R. Steinberg, eds., *Combinatorial Auctions*. MIT Press, 17–40.
- Baranov, O. 2010. Exposure vs. free-riding in auctions with incomplete information. Working paper, University of Colorado at Boulder.

- Bichler, M., A. Davenport, G. Hohner, J. Kalagnanam. 2006. Industrial procurement auctions. P. Cramton, Y. Shoham, R. Steinberg, eds., *Combinatorial Auctions*. MIT Press, 593–612.
- Cantillon, E., M. Pesendorfer. 2006a. Auctioning bus routes: The London experience. P. Cramton, Y. Shoham, R. Steinberg, eds., *Combinatorial Auctions*. MIT Press, 573–592.
- Cantillon, E., M. Pesendorfer. 2006b. Combination bidding in multi-unit auctions. Working Paper, Université Libre de Bruxelles.
- Caplice, C., Y. Sheffi. 2006. Combinatorial auctions for truckload transportation. P. Cramton, Y. Shoham, R. Steinberg, eds., *Combinatorial Auctions*. MIT Press, 539–572.
- Catalán, J., R. Epstein, M. Guajardo, D. Yung, C. Martínez, F. Espejo. 2009. Solving multiple scenarios in a combinatorial auction. *Computers & Operations Research* **36** 2752–2758.
- Chapman, J.T.E., D. McAdams, H.J. Paarsch. 2005. Multi-unit, sealed-bid, discriminatory-price auctions. Working Paper, University of Iowa.
- Chu, C. S., P. Leslie, A. Sorensen. 2011. Bundle-size pricing as an approximation to mixed bundling. *American Economic Review* **101**(1) 263–303.
- Cramton, P., Y. Shoham, R. Steinberg. 2006. *Combinatorial Auctions*. MIT Press.
- Day, R., P. Cramton. 2012. Quadratic core-selecting payment rules for combinatorial auctions. *Operations Research* **60**(3) 588–603.
- Day, R., P. Milgrom. 2008. Core-selecting package auctions. *International Journal of Game Theory* **36**(3–4) 393–407.
- Day, R., S. Raghavan. 2007. Fair payments for efficient allocations in public sector combinatorial auctions. *Management Science* **53**(9) 1389–1406.
- Day, R., S. Raghavan. 2009. Matrix bidding in combinatorial auctions. *Operations Research* **57**(4) 916–933.
- de Vries, S., R. V. Vohra. 2003. Combinatorial auctions: A survey. *INFORMS Journal on Computing* **15**(3) 284–309.
- Elmaghraby, W., P. Keskinocak. 2004. Combinatorial auctions in procurement. C. Billington, T. Harrison, H. Lee, J. Neale, eds., *The Practice of Supply Chain Management: Where Theory and Application Converge*. Springer US, 245–258.

- Epstein, R., L. Henríquez, J. Catalán, G. Y. Weintraub, C. Martínez. 2002. A combinatorial auction improves school meals in Chile. *Interfaces* **32**(6) 1–14.
- Epstein, R., L. Henríquez, J. Catalán, G. Y. Weintraub, C. Martínez, F. Espejo. 2004. A combinatorial auction improves school meals in Chile: A case of OR in developing countries. *International Transactions in Operational Research* **11** 593–612.
- Fox, J. T., P. Bajari. 2013. Measuring the efficiency of an fcc spectrum auction. *American Economic Journal: Microeconomics* **5**(1) 100–146.
- Gandal, N. 1997. Sequential auctions of interdependent objects: Israeli cable television licenses. *The Journal of Industrial Economics* **XLV**(3) 227–244.
- Glynn, P. 1989. Optimization of stochastic systems via simulation. *Proceedings of the 1989 Winter Simulation Conference* 90–105.
- Goeree, J., Y. Lien. 2009. On the impossibility of core-selecting auctions. Working Paper, California Institute of Technology.
- Guerre, I., I. Perrigne, Q. Vuong. 2000. Optimal nonparametric estimation of first-price auctions. *Econometrica* **68**(3) 525–574.
- Hendricks, K., R. H. Porter. 2007. An empirical perspective on auctions. M. Armstrong, R. Porter, eds., *Handbook of Industrial Organization, Volume 3*. North-Holland, 2073–2143.
- Hohner, G., J. Rich, E. Ng, G. Reid, A. Davenport, J. Kalagnanam, H. Lee, C. An. 2003. Combinatorial and quantity-discount procurement auctions benefit Mars, Incorporated and its suppliers. *Interfaces* **33**(1) 23–35.
- Horn, R. A., C. R. Johnson. 1985. *Matrix Analysis*. Cambridge University Press.
- Hortaçsu, A., D. McAdams. 2010. Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market. *Journal of Political Economy* **118**(5) 833–865.
- Kastl, J. 2011. Discrete bids and empirical inference in divisible good auctions. *Review of Economic Studies* **78**(3) 974–1014.
- Krasnokutskaya, E. 2011. Identification and estimation of auction models with unobserved heterogeneity. *Review of Economic Studies* **78**(1) 293–327.

- Krishna, V., R. W. Rosenthal. 1996. Simultaneous auctions with synergies. *Games and Economic Behavior* **17** 1–31.
- Levin, J. 1997. An optimal auction for complements. *Games and Economic Behavior* **18** 176–192.
- Li, J., N. Granados, S. Netessine. 2011. Are consumers strategic? structural estimation from the air-travel industry. Forthcoming in *Management Science*.
- Lu, Y., M. Olivares, A. Musalem, A. Schilkrut. 2012. Measuring the effect of queues on customer purchases. Forthcoming in *Management Science*.
- Maréchal, F., P. Morand. 2009. Free riding in combinatorial first-price sealed-bid auctions. *The B.E. Journal of Theoretical Economics* **9**(1) Article 41.
- Mas-Colell, A., M. D. Whinston, J. R. Green. 1995. *Microeconomic Theory*. Oxford University Press.
- McDuff, R.D. 2003. Analyzing package bidding in the FCC auction no. 31: Upper 700 mhz band. URL <http://cramton.umd.edu/econ415/project-fcc-package-auctions.pdf>.
- Metty, T., O. Raskina, R. Schneur, J. Kanner, K. Potts, J. Robbins. 2005. Reinventing the supplier negotiation process at motorola. *Interfaces* **35**(1) 7–23.
- Milgrom, P. 2000. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy* **108**(2) 245–272.
- Milgrom, P. 2004. *Putting Auction Theory to Work*. Cambridge.
- Moreton, P. S., P. T. Spiller. 1998. What's in the air: Interlicense synergies in the federal communications commission's broadband personal communication service spectrum auctions. *The Journal of Law and Economics* **XLI** 677–716.
- Musalem, A., M. Olivares, E. Bradlow, C. Terwiesch, D. Corsten. 2010. Structural estimation of the effect of out-of-stocks. *Management Science* **56**(7) 1180–1197.
- Olivares, M., C. Terwiesch, L. Cassorla. 2008. Structural estimation of the newsvendor model: An application to reserving operating room time. *Management Science* **54**(1) 41–55.
- Olivares, M., G. Y. Weintraub, R. Epstein, D. Yung. 2012. Combinatorial auctions for procurement: An empirical study of the Chilean school meals auction. *Management Science* **58**(8) 1458–1481.

- Paarsch, H. J., H. Hong. 2006. *An Introduction to the Structural Econometrics of Auction Data*. MIT Press.
- Pekec, A., M. H. Rothkopf. 2003. Combinatorial auction design. *Management Science* **49**(11) 1485–1503.
- Reguant, M. 2011. The welfare effects of complementary bidding mechanisms: An application to electricity markets. Working Paper, Stanford University.
- Rothkopf, M. H., A. Pekec, Ronald M. Harstad. 1998. Computationally manageable combinatorial auctions. *Management Science* **44**(8) 1131–1147.
- Sandholm, T. 2006. Optimal winner determination algorithms. P. Cramton, Y. Shoham, R. Steinberg, eds., *Combinatorial Auctions*. MIT Press, 337–368.
- Sheffi, Y. 2004. Combinatorial auctions in the procurement of transportation services. *Interfaces* **34**(4) 245–252.
- Wooldridge, J.M. 2002. *Econometric analysis of cross section and panel data*. The MIT press.