Risk Management for Combinatorial Auctions

ALAN HOLLAND

A Thesis Submitted to the National University of Ireland

in Fulfillment of the Requirements for the Degree of

Doctor of Philosophy in the Faculty of Science.

July, 2005

Research Supervisor: Dr. Barry O'Sullivan
Head of Department: Professor Gregory M. Provan

Department of Computer Science,
National University of Ireland, Cork.

Contents

List of Figures

List of Tables

Abstract

Auction theory has traditionally regarded bids in auctions as enforceable commitments. We relax this important, yet often incorrect, assumption that is common to almost all prior literature on the subject. This work addresses the possibility of winning bids being withdrawn, or reneged upon, before a transaction is completed successfully.

In particular, we examine the significance of winning-bid withdrawal in a combinatorial auction setting. We find that it may be difficult or even impossible for the bid-taker to find a repair solution of adequate revenue without causing undue disturbance to the remaining winning bids in the allocation. We have called this the *bid-taker's exposure problem* and we also show that it is exacerbated for a risk averse bid-taker.

It is preferable for the bid-taker to preempt uncertainty by choosing a solution that is *robust* to bid-withdrawal and provides a guarantee that possible withdrawals may be repaired easily with a bounded loss in revenue. We discuss the computational difficulties posed by risk management and investigate a constraint programming approach to tackling the problem. We also analyze the drawbacks of this approach and motivate useful extensions to the framework.

We then propose a new framework that facilitates solution robustness for constraint programs in a wide range of settings. We briefly demonstrate its versatility with an application to job-shop scheduling. We then apply this new framework to combinatorial auctions in order to investigate the trade-off between robustness and revenue. We also introduce a new auction model that improves solution reparability by facilitating backtracking on winning bids by the bid-taker. We demonstrate experimentally that fewer winning bids partake in robust solutions, thereby reducing any associated overhead in dealing with extra bidders.

Finally, we consider the case in which the bid-taker wishes to optimize some social objective, thereby necessitating truthful bidding. We have provided some impossibility results pertaining to truthful mechanism design that incorporate robust solutions. However, we also propose a means of circumventing this problem for a restricted class of combinatorial auctions. We develop an approximate allocation algorithm that incentivizes truthful bidding whilst attaining an allocation that minimizes the risk of revenue loss in the event of a winning bid being withdrawn.

Declaration

This dissertation is submitted to University College Cork, in accordance with the requirements for the degree of Doctor of Philosophy in the Faculty of Science. The research and thesis presented in this dissertation are entirely my own work and have not been submitted to any other university or higher education institution, or for any other academic award in this university. Where use has been made of other people's work, it has been fully acknowledged and referenced. Parts of this work have appeared in the following publications which have been subject to peer review.

1. Alan Holland and Barry O'Sullivan, Weighted Super Solutions for Constraint Programs , *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI-05)*, July 2005, Pittsburgh, Pennsylvania.
(Citeseer Conference Impact Rating: top 9.17%, Acceptance Rate: 18.4%.)

2. Alan Holland and Barry O'Sullivan, Robust Solutions for Combinatorial Auctions , *Proceedings of the Sixth ACM Conference on Electronic Commerce*, June 2005, Vancouver, Canada.
(Citeseer Conference Impact Rating: top 5.15%, Acceptance Rate: 29%.)

3. Alan Holland and Barry O'Sullivan, Super solutions for combinatorial auctions, *Proceedings of the Workshop on Modelling and Solving Problems with Constraints*, ECAI 2004, Valencia, Spain.

4. Alan Holland and Barry O'Sullivan, Super solutions for combinatorial auctions, *ERCIM/Colognet Constraints Workshop*, 2004, Springer LNAI, Lausanne, Switzerland.

5. Alan Holland and Barry O'Sullivan, Towards Fast Vickrey-Pricing using Constraint Programming , *Artificial Intelligence Review*, 21, 335-352, 2004.

6. Alan Holland and Barry O'Sullivan, Fast Vickrey-pricing for the Assignment Problem, *Ercim/Colognet Constraints Workshop*, 2003, Budapest, Hungary.

The contents of this dissertation extensively elaborate upon previously published work and mistakes (if any) are corrected.

Alan Holland
July 2005.

Dedication

This dissertation is dedicated to Annie.

Acknowledgements

There are many people for whom I am eternally grateful for making this work possible. Firstly, I would like to thank my PhD supervisor, Dr. Barry O'Sullivan, for all his help and advice. He has always been extremely generous with his time throughout this research. In particular, his effort and hard work during the writing of this dissertation is very much appreciated.

I would like to thank Emmanuel Hebrard, Brahim Hnich and Toby Walsh whose work provided inspiration for some of the contents of this dissertation. They also provided useful advice and assistance during the course of this work. Other members of the Cork Constraint Computation Centre who provided welcome advice and criticism at various stages include Eugene Freuder, Rick Wallace and Nic Wilson. I am also very grateful to Professor Tuomas Sandholm of Carnegie Mellon University for his help during the early stages of my research when he acted as my Doctoral student adviser during CP 2003 in Kinsale, and later during the the course of discussions at ACM E-commerce 2004 and 2005.

I would like to acknowledge the financial support of Enterprise Ireland and Science Foundation Ireland, whose funding made this research possible.

I would also like to thank my fellow students Barry O'Callaghan, Colin Quirke, Jerome Kelleher, Peter Blackburn and Sarah O'Connell who all helped to make student-life far more entertaining and enjoyable.

I am also very grateful to my family, and in particular my parents, for their help and support throughout this work.

And finally, my deepest thanks must go to my partner in love and life, Annie, for always being there during the highs and lows of my PhD. I couldn't have done it without you.

Chapter 1

Introduction

1.1 Auctions

An auction is an economic mechanism by which a single seller explicitly permits bidders, through their bids, to determine both payments to the seller and an allocation of the item(s) being sold. It is an ancient means of trading goods where the precise value of items for sale is unknown and the auction mechanism is used to determine the transaction price(s).

There are four basic forms of auctions for the sale of a single item: first-price ascending (English), first-price descending (Dutch); first-price sealed-bid and second-price sealed-bid (Vickrey) auction. The English auction is the quintessential auction format in which the price of the item for sale increases until there is only one bidder willing to pay that amount. Variations on these auction rules are important in some instances. Consider the Amsterdam flower auctions, for example, where the goods are perishable. It is important that the auction is executed expeditiously, therefore, a descending Dutch auction, where the price of flowers decreases from an initial high in timed decrements, places a bound on the length of time required to conduct the auction. Other considerations, such as the geographic dispersion of bidders, motivate the use of sealed-bids.

The reasons for the increased popularity of auctions are the perceived improvements in efficiency and revenue for the seller over less competitive procedures such as "beauty contests", in which a panel of judges determine the winners using a comparative selection procedure. As the nature of the items for sale have increased in both value and complexity, so too has the need to ensure efficiency and the ability to withstand strategic manipulation. Some disastrous auction results in recent history have highlighted the need to carefully manage auctions. A notable example was the sale of spectrum rights in New Zealand in 1990 [89]. A second-price sealed-bid auction was used to sell licenses which resulted in one particular auction where the highest bid was for NZ$7 million and the second highest bid was only NZ$5,000 so that the winning bidder only paid the latter amount. This problem highlights the importance of intelligent auction design. The political repercussions for a government following such a revenue failure could be severe, so an auction format whose expected income has less variance is desirable.

In recent years there has been a proliferation in the use of auctions across many sectors in modern commerce. Large-volume auctions are used to sell and procure item such as:

Spectrum Rights. National Governments sell the rights to broadcast on certain frequencies in the electromagnetic spectrum in given geographic locations to telecommunications companies.

Electricity Production. Electricity generators and consumers partake in auctions for the supply of power to a grid.

Transportation. Transport providers bid for contracts to transport items for consumers in reverse (procurement) auctions.

There has also been a significant increase in research focused on designing new auctions suited toward these emerging applications [83, 85]. However, more advanced auction mechanisms can also introduce new problems. Combinatorial auctions have emerged in recent years as a popular auction type for the sale of multiple distinguishable items among which bidders perceive complementarities or substitutabilities. Such auctions, otherwise known as package auctions, permit bidders to select subsets of items they wish to bid upon, thus facilitating full expression of super- or sub-additive valuations over sets of items. Such auctions incur difficulties in terms of preference elicitation from bidders and winner determination that is computationally challenging. However, there has been significant advances in tackling these problems in recent years [18, 65, 100, 101, 115].

1.2 The Topic of this Research

Combinatorial auctions are becoming more popular in modern commerce because of the improvements in efficiency that are achievable when bidders can express their preferences over sets of distinguishable items for sale. When selling a farm, for example, it may be sold as a single item or divided into separate packages such as the farmhouse, outhouses, sites overlooking the beach, arable and non-arable land. When deciding on how to package the sale, it is impossible for the seller to know in advance what packages would maximize revenue. Bidders may view the complementarities between items differently, and as the size of the auction grows it quickly becomes impossible for the auctioneer to know how the items should be packaged. It is more economically efficient for the bidders to bid on combinations of items. Unfortunately, the number of possible combinations of items of interest to each bidder grows exponentially as the number of items increases. The bids also need to be communicated to the auctioneer in a concise manner, which becomes increasingly difficult in large auctions.

This dissertation examines the hitherto unexplored problem of solution robustness in such auctions. We first of all determine the implications of bid-withdrawal upon revenue in a number of economically motivated scenarios. We examine possible preemptive actions that may prevent large losses in revenue and then explore different means of tackling this problem. In particular, we study constraint programming methodologies for establishing robust solutions for combinatorial auctions.

1.3 The Goals of this Research

The overarching goal of this work is to investigate a problem that presents computational difficulties to state-of-the-art solving techniques, demonstrate its significance, and find a novel solution to this problem. In doing so, we strive to maintain a cohesive argument throughout this dissertation that supports the central thesis presented in § 1.4.

In particular, we aim to investigate the effects of solution uncertainty in combinatorial auctions and to provide a means of establishing robust solutions that lead to controlled losses in revenue, given a bid-withdrawal. This may be seen as a form of risk management for the bid-taker in which there is a trade-off between risk and revenue. This trade-off must be achieved in a computationally feasible manner and ideally minimize the revenue sacrificed in order to reduce risk to an acceptable level.

Sometimes, bid-takers do not seek revenue maximizing solutions, but instead wish to allocate items according to some social objective, *e.g.* maximize social welfare. An important result from economics, called the *Revelation Principle* [50, 82], implies that only truthful auction mechanisms need to be considered when optimization of a social choice function is desired. However, truthful mechanisms pose computational challenges for mechanism designers that we aim to overcome so that we can facilitate risk management for a non-revenue-maximizing bid-taker.

1.4 A Statement of the Thesis

The thesis presented and defended in this dissertation can be stated as follows:

"*Winning-bid withdrawal in combinatorial auctions following winner determination can cause large losses in revenue. The bid-taker can become stranded in a local optimal solution of low revenue from which satisfactory repair solutions cannot be attained, because of the awarding of items to other non-reneging bidders. Preemptive steps are necessary to reduce this risk so that a solution can be repaired to form another solution of acceptable revenue given a possible withdrawal. Robust solutions that guarantee a minimum level of revenue, given possible failures, are particularly desirable when the bid-taker is risk averse. It is possible to find robust solutions for combinatorial auctions in a computationally feasible manner. Furthermore, it is possible to incentivize truthful bidding for risk managed allocations, given certain restrictions.*"

1.5 Contributions of this Research

This dissertation makes the following contributions. Firstly, we identify the "bid-taker's exposure problem" in combinatorial auctions and show that it can have a serious impact on revenue for the bid-taker. We also elaborate upon the difficulty of solving this problem and how finding robust solutions can be preferable to maximizing expected revenue for a risk averse bid-taker.

Secondly, we propose an initial solution to this problem using an existing constraint programming technique that bounds the number of changes required to form a repair solution following the withdrawal of a winning bid [63]. We then analyze the deficiencies of this approach when applied in a real-world context. We proceed to develop a novel framework that finds robust solutions with probabilistic failures, and varying costs of repair for changes to the solution [64]. This framework has been developed as a versatile means of establishing robust solutions for any constraint program and its flexibility has been demonstrated using job-shop scheduling problems.

We then apply this new framework to combinatorial auctions involving thousands of bids with very positive results [62]. We examine two separate auction mechanisms that permit withdrawal. The first mechanism has no penalty payments associated with bid-withdrawal, and items assigned to winning bidders cannot be withdrawn by the bid-taker to help find a repair solution. The second mechanism allows bid-withdrawal for a given penalty and uses this penalty payment to fund a reassignment of items to bidders with the bid-taker paying compensation to previously winning bidders for bid revocation. This improves solution reparability and increases the expected revenue of optimal robust solutions.

We also examine the strategic implications of such a non-incentive-compatible and non-optimal winner determination algorithm. We prove an impossibility result regarding truthfulness and robust solutions. We show that this result also applies to auctions that permit item revocation by the bid-taker. Finally, we propose a truthful approximation scheme for robust solutions that circumvents these difficulties and incentivizes truthful bidding.

1.6 Structure of this Dissertation

This dissertation follows a logical path that begins with the exposition of a significant problem in the field of auction theory and proceeds towards its resolution by development of a novel technique in constraint programming, a subfield of artificial intelligence. At each stage of this process, the preface of each chapter will summarize the progress of this resolution process in preceding chapters and motivate the work in the relevant following chapter. The dissertation is constructed as follows.

Chapter 2 outlines the background relating to auction theory and describes how bidders behave in a strategic manner to maximize their expected surplus from auctions. We also introduce the field of mechanism design and show how it can be applied to the design of auctions so that the bid-taker can design allocation and payment schemes to incentivize desired outcomes. We introduce combinatorial auctions and outline the motivation for such auctions. We then examine the possible implications for uncertainty in such auction settings and describe how the discipline of risk management quantifies such risks so that an objective comparison of the utility of different outcomes can be made possible.

Chapter 3 expounds the "bid-taker's exposure problem" and analyzes the implications for revenue given a bid-withdrawal in a combinatorial auction. Risk averse bid-takers will seek to hedge against the danger of a large loss in revenue and are, therefore, particularly vulnerable to this problem. We discuss possible measures to counteract this risk and discuss the computational implications of the different approaches.

Chapter 4 introduces constraint programming (CP) and discusses a framework that can be used to find robust solutions for constraint programs. Combinatorial auctions are then modeled as constraint optimization problems and this framework is then applied to these problems. The results of these experiments are discussed and shortcomings of the approach are highlighted and used to motivate a new extended framework.

We then proceed to introduce a novel technique for finding robust solutions in Chapter 5, called the weighted super solutions framework, and discuss its usefulness and versatility. We show how this framework can model any constraint program whose assignments may exhibit probabilistic failure rates over time. We also propose a metric for repair costs that can capture the true costs of changing a solution in a more realistic manner. We provide job-shop scheduling as an example of a possible alternative application domain.

Chapter 6 applies the new framework to combinatorial auctions and demonstrates its usefulness in a variety of economically motivated settings. A new auction mechanism for improving solution reparability, called “mutual bid bonds”, is also presented and it is shown that this mechanism increases the expected revenue of optimal robust solutions. A number of interesting and unexpected results arise from the empirical analysis conducted in this chapter.

Robust solutions for combinatorial auctions present a difficulty for bidders in that computing the equilibrium bidding strategy is a difficult problem. Chapter 7 examines mechanisms for robust solutions that incentivize truthful bidding, thereby easing determination of bidding strategy. Truthful bidding also allows the bid-taker to maximize a social objective. We present a negative result surrounding truthfulness and WSSs but we circumvent the problem by relaxing revenue constraints, thus enabling the creation of a truthful mechanism for approximately utility maximizing solutions in the presence of exogenous probabilities of bid withdrawal.

We conclude with a chapter that discusses the relevance of our contributions to the fields of electronic commerce and constraint programming. We also propose natural extensions to this work and suggest possible future work.

Chapter 2

Background: Auctions and Risk Management

Auction theory is one of the triumphs of game theory and has proved particularly useful since the introduction of auctions for spectrum rights sales [83]. This chapter provides an overview of auction theory and introduces the four most common auction forms: English, Dutch, sealed-bid first-price and sealed-bid second-price (Vickrey). We also discuss the different valuation models for bidders that influence the strategic decision-making process of choosing what bid amount maximizes the expected profit. This is particularly pertinent in situations when there is uncertainty surrounding the true underlying value of an item. For example, the intrinsic value of the oil being acquired in an auction for oil-drilling rights is the same for all bidders although the estimates of the oil available may differ. This problem for bidders can become quite complicated and is resolved using game-theoretic analysis. We also present an important result from the field of auction theory: the *Revenue Equivalence Theorem*.

Mechanism design can be used to define allocation and payment rules for auctions so that rational bidders follow strategies dictated by game-theoretic analysis. It is currently a very active research topic [8, 29, 45]. Mechanism design may be seen as form of "inverse game theory". More specifically, mechanisms that incentivize truthful revelation of valuations become particularly interesting because of a famous result known as the *Revelation Principle* [50, 82] that allows us to restrict our attention to truthful mechanisms when trying to maximize some social objective.

Of particular interest is an emerging auction form, the combinatorial or ‘package auction’. The popularity of the combinatorial auctions has been increasing in recent years due to its perceived economic efficiencies, as well as improvements to algorithms that have addressed computational impediments to winner determination [51, 115, 119]. We discuss the advantages and disadvantages of such auctions and present a combinatorial auction (CA) simulation tool that mimics bidder behavior in different economically motivated scenarios. We shall use this tool in subsequent chapters to analyze the effects of bid withdrawal [78].

Finally, we introduce the field of risk management in this chapter, because we aim to investigate the effects of solution uncertainty for CAs. Winning bid withdrawal constitutes a considerable risk for the bid-taker in such auctions and a quantitative approach to evaluating tolerable levels of risk is desirable. Risk is a subjective notion and we describe how a bid-taker’s risk attitude affects his preference over outcomes in different scenarios.

2.1 Auction Theory

An auction is a market mechanism in which an item, or items, are exchanged on the basis of bids submitted by participants. The auction mechanism consists of a set of rules that govern the sale, or purchase in a reverse auction, of an item according to the most favorable bid. Auction theory concerns the design of auctions and how their rules should be determined so that desired goals are achieved. The goals of auction design, *e.g.* revenue or social welfare maximization, may differ according to the requirements of the bid-taker and the bidders’ valuations for the item(s) on sale.

Auction theory has emerged as an important field of research in recent years [72, 73]. The proliferation of large-volume, state-run auctions for items such as spectrum-rights has focussed attention on how auctions should be designed [83, 85]. Subtle changes to auction rules can cause marked differences in outcomes.

2.1.1 Common Auction Types

In this section we present an overview of four of the most common types of auction used in commercial settings.

English auction. This is the classical auction type that is typically used in the sale of property or collectibles. Participants bid openly against one another, with each successive bid being higher than the previous one. The auction continues as long as there is someone willing to outbid the current asking price. The auction then ends when no participant is willing to bid further, at which point the highest bidder pays his declared bid amount. The seller may set a reserve price so that if the auctioneer fails to raise a bid higher than this reserve the sale may not go ahead, but the seller typically pays a fee to the auctioneer in any case.

Dutch auction. The Dutch auction is also known as an "open-outcry descending price" auction or a "clock auction" [85]. The auctioneer starts at a high price and subsequently lowers the price repeatedly as time passes. This is sometimes achieved by using a clock which gradually decrements the price. The first bidder who communicates that he will accept the current price wins the item at that price for a quantity they may specify publicly to the auctioneer. The Dutch auction has been made famous by the Amsterdam flower auctions, where speed is essential because of the perishable nature of the products for sale.

Sealed-Bid first-price auction. In this type of auction all bidders simultaneously submit bids in such a way that no bidder knows the bid of any other participant. The highest bidder pays the price he submitted. In this form of auction, bidders strategize about the valuations of other agents when evaluating what they should bid in order to maximize their expected surplus.

Bidders in the traditional Dutch auction and sealed first-price auction will bid below their true valuation for an item so that they can maximize their expected profit. This tactic is known as bid shading. These two auctions are also theoretically equivalent, but in practice Dutch auctions will produce less revenue than sealed first-price auctions. This is one of the most important results of experimental economics [121, 122].

Sealed-Bid Second-Price (Vickrey) auction. This auction is conducted in the same manner as a sealed-bid first-price auction with the difference that the winner only pays the amount equal to the second highest bidder. The auction format has appealing strategic properties in that it is a dominant strategy for each bidder to bid his true valuation. A recent auction of note was the initial public offering (IPO) of Google shares, that used a modified multi-unit, homogenous, Vickrey auction [32, 60, 136]. In a pure Vickrey auction the optimal strategy for bidders is to bid their true maximum value. So if a bidder is willing to pay $100 per share then that should consititute his bid. If the Google IPO auction was a pure Vickrey (assuming all potential bidders were rational and permitted entry to the auction) then the post-IPO market price would equal the IPO clearing price. This means that the post-auction share price should not fluctuate excessively in the immediate aftermath of the flotation. However, the auction was not a pure Vickrey auction and all the bidders were not completely rational. For example, some bidders may not have been able to calculate their optimal strategy correctly. Google also reserved the right to set the offering price lower than the auction-clearing price.

2.1.2 Private, Common and Affiliated Values in Auctions

The behavior of bidders in an auction depends upon whether they can draw conclusions from competing bids. In an *independent private-values* model, each bidder knows how much he values the object for sale but his value is not dependent upon the bids of others [48, 126]. Each bidder derives a value from only his own personal tastes and not external factors such as re-sale value. The revelation of other bids does not influence each bidder's valuation.

The *common-values* model was subsequently introduced, where the true actual value is the same for everyone, *e.g.* the oil in a drilling rights auction, but bidders have different private information ("signals") about the true actual value [111, 138, 140]. In this model a bidder would change his estimate of the value if he learned another bidder's signal. Common valuations often occur in auctions for rights to natural resources [21]. If a bidder's signal was significantly more than all other bids for example, he may re-estimate the value of the item, therefore, his *ex-post* valuation may be decreased. If it decreases to below his bid amount, he is then a victim of the "*winner's curse*", a term first coined by Capen *et al.* [21]. The winners in common-value auctions are necessarily the most optimistic bidders when payment is conducted using a first-price scheme. This can sometimes result in winning an item at a cost of more than the ex-post valuation [19, 74, 139], which may result in serious losses for the winner [56]. Section 3.1.3 deals with this problem in more detail.

In practice, many goods have neither purely private values nor purely common values. Collectible items are likely to have an important private value aspect, where each individual bidder alone knows how much he likes or dislikes the work, but also a common value aspect because the winner may decide to resell the work at a later date. The combination of private and common value models is referred to as *affiliated* values and was first introduced by Milgrom and Weber [86].

The presence of *asymmetric information* is another important feature in auction theory. Environments with asymmetric information describe situations in which some agents hold private information that is relevant to all parties [138]. This information can be *directly* relevant in that it directly affects the payoffs of the bidders. For example, when the bid-taker knows the quality of the items for sale but the bidders do not. Asymmetric information can also be *indirectly* relevant by helping each agent to gauge the expected rational behavior of others and thereby solve his strategic uncertainty.

2.1.3 Game-Theoretic Bidding Behavior

Game theory is a mathematical theory of strategic interaction where multiple players must make decisions that may affect the interests of other players. An auction is an example of a game in which bidders are competing agents, each of whom is seeking to maximize their utility. The bid-taker sets the rules for the game in such a way as to achieve his objective, which is often revenue maximization but may also be the fulfilment of some social objective. The bidders, on the other hand, will strategize so that their expected surplus is maximized [74, 84, 85, 96].

A *strategic equilibrium* is a profile, or combination, of strategies such that if other players conform to the equilibrium strategies (*i.e.* other bidders are rational), no player has an incentive to unilaterally deviate from his equilibrium strategy [123]. Game theory provides several solution concepts to compute the outcome of a game with self-interested agents. A *solution concept* is used to predict the strategies agents will choose in order to maximize their utility, thus determining an equilibrium position for the game. These concepts assume knowledge about agent preferences, rationality, and shared information about one other. The best known concept is that of a Nash equilibrium, which states that in equilibrium every agent will select a utility-maximizing strategy given the strategy of every other agent [91]. A Nash-equilibrium is self-referential and constitutes a profile of strategies that form "optimal reactions" to other agents' "optimal reactions". Nash equilibrium is the pure form of the basic concept of strategic equilibrium; as such, it is useful mainly in normal form games with complete information. When allowing for randomized strategies, at least one Nash equilibrium exists in any game with regular payoff functions [123]. A game may possess one or more Nash equilibria. We discuss two example auction forms below and outline their equilibrium bidding strategies.

Second-price (Vickrey) auctions with private values. In this type of auction, the winner is the bidder with the highest bid but he only pays the second highest bid amount for the item. The optimal bidding strategy is to submit a bid equal to one's actual true value for the item. This is a *weakly dominant* strategy, *i.e.* no matter what other bidders do, no other strategy can achieve a superior outcome. Note that a *dominant* strategy in a game for a player gives a better payoff than any other strategy, regardless of what the other players are doing. It weakly dominates another strategy if it is always at least as good [50]. In an English auction the winner pays the amount at which the second highest bidder dropped out, suggesting that it lends itself towards analysis via the *second-price* auction format [125]. Vickrey's analysis can, therefore, be applied to such auctions when attempting to predict bidder behavior [126].

Sealed-bid first-price auctions with common values. Consider the sale of drilling rights in an oil-field. A strategic bidder recognizes that winning the auction implies that other bidders estimated a lower valuation for the oil in the field. These other estimates would cause the bidder to be nervous about their valuation and result in a lower estimate of the value of the oil. The equilibrium bidding strategy is, therefore, to reduce the bid amount slightly to take the "winner's curse"[†] into consideration. This is an example of a *Bayesian game* in which the players have incomplete information about the game. This can be described in terms of a game in which a player may have one of many (or even infinite) "types", and that the type of any player is known to that player but unknown to others [50]. The type of a player determines the payoffs that player receives from any outcome of the game. The common equilibrium notion for such games is *Bayes-Nash Equilibrium* (BNE). In a BNE, each player picks a strategy function, rather than a simple strategy [50, 125]. The strategy function then selects a particular strategy for the player's type. A BNE is a profile of strategy functions such that no single player can improve his expected utility by changing his function. BNE is a solution concept that is often applied to auctions [28, 105].

2.1.4 Revenue Equivalence Theorem

An important result in auction theory is the *Revenue Equivalence Theorem* (RET) which states that, if all bidders are risk neutral, and have independent private values for the auctioned items, then all four of the standard single unit auctions (mentioned in § 2.1.1) have the same expected revenue.

Theorem 2.1.1 (Revenue Equivalence Theorem (RET) [90, 108])**.** *Assume each of a given number of risk neutral potential buyers has a privately-known valuation independently drawn from a strictly-increasing atomless distribution, and that no buyer wants more than one of the k identical indivisible prizes. Then any mechanism in which (1) the prizes always go to the k buyers with the highest valuations and (2) any bidder with the lowest feasible valuation expects zero surplus, yields the same expected revenue (and results in each bidder making the same expected payment as a function of her valuation).*

Proof. See [90] for a proof. □

The RET applies broadly to many auction types, beyond just the English, Dutch, first- and second-price auctions, to include many other auction mecha-

[†]This effect will be described in more detail in § 3.1.3.

nisms. It is a remarkable result, as various auctions may have completely different strategies and rules. It is important to note, however, that although the expected revenue is the same for auctions that satisfy the conditions in Theorem 2.1.1, it may vary in different specific instances. Some auctions return revenues with greater variance even though the mean revenue is the same. A risk averse bid-taker may not welcome variance and would, therefore, prefer a more predictable outcome [132]. This offers a partial explanation for the prevalence of first price auctions over second price auctions.

2.2 Mechanism Design for Auctions

Mechanism Design can be considered as "inverse game theory" whereby the rules of the game are decided by an authority so as to fulfil some objective. Two typical goals of auction design are either revenue maximization or maximization of social welfare. The goal of maximizing revenue is an obvious one and features in auctions where the identity or private valuations of the winning bidder(s) matter little when compared to the revenue received by the bid-taker. In some circumstances, however, the bid-taker may wish to achieve certain social objectives, but because these individuals' actual preferences are not publicly observable, the analysis of such auctions can become more complicated. The *mechanism design problem* is to elicit these preferences so that they may be aggregated into social preferences to form a collective decision [50, 82].

2.2.1 Mechanism Design Goals

The traditional goal of mechanism design is to determine the rules of a game in which an overall equilibrium (or equilibria) is reached according to some desirable system-wide properties, given that all participating agents are self-interested [82]. A *social choice function* (SCF) describes the properties that the outcome should possess. Some typical desirable properties include the following:

- *Individual Rationality*: No agent attempts to take part in a trade that fails to increase, or at least leaves constant, his own utility [79]. This is an important property if agent participation is voluntary.

- *Efficiency*: The outcome must maximize overall agent utility, thereby maximizing social welfare.
- *Revenue Maximizing*: A single agent, an auctioneer for example, maximizes her revenue (utility).
- *Budget Balance*: The sum of all agent payments equals zero, therefore, no money is extracted or injected into the system. This is particularly important for any self-sustaining mechanism where no external benefactor exists to subsidize the system.

However, these desirable properties may directly conflict with one another. For example, budget balance and efficiency conflict in Vickrey auctions, which achieve only the latter property.

The *Revelation Principle* [50, 82] is a fundamental tenet of mechanism design and implies that in a wide variety of settings, only "truthful revelation mechanisms" in which agents truthfully announce their types need to be considered. This result is not immediately obvious but means that there are no manipulable mechanisms that, when agents strategically report their types, attain superior outcomes according to the social objective than any non-manipulable mechanism. Conitzer and Sandholm [30] have demonstrated, however, how this principle may fail in certain extreme instances when computational or communication complexity hinders strategic manipulation.

2.2.2 Generalized Vickrey Auction

The seminal contributions of Vickrey [126], Clarke [25] and Groves [55] (VCG) to the field of mechanism design provide an important standard by which other mechanisms are judged. VCG is a general method of designing truthful mechanisms and a proof of its truthfulness may be found in [82].

For a given mechanism, a *solution concept* is used to predict the strategies that agents will choose in order to maximize their utility, thus determining an equilibrium position for the game. The Generalized Vickrey Auction (GVA) uses the VCG mechanism to determine payments and has a *dominant strategy* equilibrium. This means that the best response strategy for each agent remains the same, irrespective of its knowledge about other agents or their actions. Agents bid their truthful valuation for an item and do not profit from under-estimating or exaggerating their valuation even if they know all other bids. This is a powerful solution concept and makes it unprofitable for agents to concern themselves with other agents' bids.

The GVA also ensures optimal efficiency, thereby maximizing social welfare according to some objective such as fairness. However, it is not budget balanced, so a benevolent external party, such as a government, may be required to supplement the budget. To determine payments to each agent participating in the overall optimal solution, the overall revenue is determined without each agent present in turn. This involves $m + 1$ optimization problems if m agents participate in the optimal solution to the allocation problem.

Although optimization of a social objective becomes possible using such a truthful mechanism, there are several notable disadvantages of the GVA that are worth highlighting. If non-optimal solutions are found to the optimization problems that determine the prices paid (based on the Vickrey-Clarke-Groves mechanism [25, 55, 126]) the mechanism is no longer guaranteed to be truthful. This is a major disadvantage of the GVA. Various polynomial-time approximation algorithms can provide good or near optimal solutions very quickly. However, Nisan and Ronen [94] showed, constructively, that a non-optimal solution can in fact result in payments arbitrarily far from optimum. If an auctioneer seeks to approximate optimal solutions in a GVA using polynomial-time algorithms the results may not be reliable and agents may have an incentive to lie.

Other limitations include reduced revenue compared with other auctions and susceptibility to a fraudulent auctioneer. It is possible for an auctioneer to introduce fake bids just below the value of the winning bids to increase revenue. For this reason a trustworthy auctioneer is imperative in a GVA.

A GVA is also sensitive to bidder collusion. Bidders may coordinate their bid prices so that the bids remain artificially low. In this manner, the bidders get the item at a lower price than they normally would. The GVA self-enforces some typical collusion agreements and makes it easier for agents to conspire by allowing the coordination of bids. Therefore, from the perspective of deterring collusion, Dutch or first-price sealed-bid auctions are preferable because colluding agents require more trust in one another for the collusion to succeed.

2.3 Combinatorial Auctions

The popularity of online auctions has increased in recent years because the internet promises to promote competition, thereby increasing revenue for the auctioneer. In separate single-item auction, where the items exhibit complementarities or substitutabilities, there exists a phenomenon known as the *exposure problem* when items are sold separately [85]. This occurs when bidders seek a certain set of items but do not want to end up with a subset of the items that they may find valueless. This encourages cautionary bidding tactics that result in depressed bidding. Combinatorial auctions may alleviate the exposure problem by permitting bids on an arbitrary combination of items that suits the bidders needs. In this manner such auctions improve efficiency where items exhibit complimentarities/substitutabilties for the bidders.

Such auctions have been used in many real-world scenarios such as procurement for the Mars corporation [61] and the sale of spectrum licences in America's Federal Communications Commission (FCC) auctions [83]. The London Transport Authority also operated a CA in their procurement of bus services from private operators [85, 102]. The Chilean government have also adopted CAs for the supply of school meals to children [41]. In the latter case, the quality of suppliers was considered as well as the bid amount in deciding the winner, and the system also ensured there was no monopoly in any individual region. The reported supply costs have fallen by 22% since the adoption of the program.

In auctions where complementarities or substitutabilities are exhibited between items, there is a compelling argument for the introduction of combinatorial bidding to improve overall efficiency. Their application is spreading to other areas such as Supply Chain Management [131].

Combinatorial auctions fall into two categories, forward and reverse. In a reverse auction the auctioneer is seeking to procure goods in order to minimize cost. If it is possible to purchase more items than are strictly necessary, the problem is a Set Covering Problem. The buyer may choose to stipulate in the auction rules that no surplus items are to be purchased. With the introduction of this constraint this problem becomes a Set Partition Problem.

In the remainder of this dissertation we focus our attention on forward auc-

Table 2.1: Combinatorial auction example.

	Items		
Bidders	A	B	AB
x_1	**0.60**	0.00	0.00
x_2	0.00	**0.60**	0.00
x_3	0.00	0.00	1.15
x_4	0.00	0.00	1.10

tions, where items are being sold and the objective is to maximize revenue.

2.3.1 The Winner Determination Problem

It is important that the bid-taker can determine the winners from all bids received in a timely manner. This is known as the *winner determination problem* (WDP) which can be represented as a Set Packing Problem (SPP) for forward auctions. In this case it is not necessary to sell all items in order to maximize revenue.

Example 2.3.1. *Consider a simple example where an auctioneer is selling two items and there are four interested parties, bidders x_1, x_2, x_3 and x_4 (Table 2.1). Bidders x_1 and x_2 are interested in the first and second items, respectively, bidding $0.60million for the those items. Bidders x_3 and x_4 seek both items only and bid $1.15million and $1.10million, respectively, for the pair but $0 for each item individually[†]. The revenue maximizing solution for the auctioneer is to sell the items separately to x_1 and x_2 thus securing $1.2million. It is impossible for the auctioneer to know in advance whether combining the items in a single sale would be profitable or not. Combining both items in this case would have resulted in $0.05million of lost revenue, since the winning bid would have been for $1.15million rather than $1.2million. Instead, CAs allow the bidders decide on parcels of items that suit their needs, thereby improving overall efficiency.* △

Combinatorial auctions involve a single bid-taker allocating multiple distinguishable items amongst a group of bidders. The bid-taker has a set of m items,

[†]Such arbitrary complementarities amongst different bidders are often seen in property sales.

$M = \{1, 2, \ldots, m\}$, for sale and bidders submit a set of bids $\mathcal{B} = \{B_1, B_2, \ldots, B_n\}$. A bid is a tuple $B_j = \langle S_j, p_j \rangle$ where $S_j \subseteq M$ is a subset of the items for sale and $p_j \geq 0$ is a price. The WDP for a CA is to label all bids as either winning or losing so as to maximize the revenue from winning bids without allocating any item to more than one bid. The following is the Integer Programming formulation for the WDP:

$$max \sum_{j=1}^{n} p_j x_j$$

$$s.t \sum_{j|i \in S_j} x_j \leq 1,\ \forall i \in \{1 \ldots m\}, \quad x_j \in \{0, 1\}.$$

This problem is $\mathcal{NP}$-complete [110] and inapproximable [115], and is otherwise known as the Set Packing Problem. The above problem formulation assumes the notion of *free disposal*. This means that the optimal solution need not necessarily sell all of the items. If the auction rules stipulate that all items *must* be sold, the problem becomes a Set Partition Problem [33]. The WDP has been extensively studied in recent years. The fastest search algorithms that find optimal solutions (*e.g.*, CABOB [115] or CPLEX [66]) can, in practice, solve very large problems involving thousands of bids very quickly.

2.3.2 Combinatorial Auctions in Practice

Combinatorial auctions permit bidders to express complementarities amongst items. Auctions for different types of items tend to exhibit different patterns in complementarity. For example, in spectrum rights auctions there is a clear super-additive valuation in securing rights in adjacent geographical areas because there is the potential for fewer transmitters than would be necessary to cover those regions on their own.

Leyton-Brown *et al.* introduced a bid simulation tool for CAs called the 'Combinatorial Auction Test Suite' (CATS) [78]. This tool can simulate bid distributions for different economically motivated auction scenarios listed below:

`arbitrary`. Bidders view complementarity between items in an arbitrary manner, as in electronic component or pollution emitting rights auctions. The sets of desired items for different bidders are random and possess no structural similarities.

`matching`. Temporal relationships between pairs of items are exhibited, *e.g.* take-off and landing slots in airports, so complementarities are more highly correlated for different bidders than for `arbitrary` auctions.

`paths`. This bid distribution simulates the sale of routes between points in 2-D space. This is applicable in logistics, network routing and gas pipeline auctions. Complementarities are very highly correlated.

`regions`. The sale of items whose complementarity relates to adjacency in 2-D space, e.g. spectrum rights or property. Complementarities exhibit moderate levels of correlation.

`scheduling`. Simulation of auctions for time slices on machines in distributed job-shop scheduling problems [135]. Complementarities exhibit moderate levels of correlation.

The above bidding distributions reflect realistic bidding patterns that we shall use in the course of this work to examine the effects of solution robustness in different types of auctions. Prior to the introduction of CATS, artificial bid distributions were used to examine the efficiency of algorithms for the WDP [2, 33, 119]. However, we are more interested in exploring realistic bid distributions, because algorithmic efficiency is not the central focus of this work.

2.4 Risk Management

Risk management is a discipline that enables individuals and organizations to cope with uncertainty by taking steps to protect its vital assets and resources. It is an integral part of good management and encourages stability of wealth [17]. When potential winning bids may be withdrawn from a solution to a CA, this constitutes a considerable risk for the bid-taker because item retraction from other (non-reneging) winning bidders may not be possible. We are interested in managing this risk in a controlled fashion so that we can balance revenue against the risk of losses caused by unreliable bidders. An uncertain solution to an auction can be seen as a lottery over outcomes, whereby potential losses vary according to the probability of bids being withdrawn and the extent to which a solution can be repaired following such a withdrawal [63].

In this section we present expected utility theory, that incorporates both Bernoulli utility functions over wealth and von Neumann-Morgenstern preferences over lotteries [68, 127]. We also describe risk aversion and the concave nature of its associated Bernoulli utility functions. This is an important concept because the risk attitude of the bid-taker has a fundamental role in deciding upon a utility maximizing allocation, as we shall see in future chapters.

2.4.1 Expected Utility Theory

In 1944, von Neumann and Morgenstern published their seminal book, *Theory of Games and Economic Behavior* [127], in which they defined an expected utility function over lotteries, or gambles. Theirs is an axiomatic derivation, meaning a set of assumptions over individuals' preferences is required before one can construct a utility function.

According to von Neumann and Morgenstern a lottery, or gamble, is simply a probability distribution over a known, finite set of outcomes. In this framework, we know with certainty the probability of occurrence of each outcome. Based on the probability associated with each outcome, we can define a simple lottery as a set of outcomes, $\mathcal{O} = \{o_1, \ldots, o_n\}$, each of which occurs with some known probability p_i. We can also construct compound lotteries, which are probability distributions over lotteries, so that an outcome of a lottery may itself be an entry to another lottery.

The Preference Axioms

In order to construct a utility function over lotteries, or gambles, a number of assumptions about individuals' preferences are required. Jensen [68] has provided a proof for von Neumann and Morgenstern's expected utility hypothesis [127] using an axiomatic derivation whose axioms include the following:

Completeness. All lotteries can be ranked so that for any 2 gambles g_1 and g_2 in G, $g_1 \preceq g_2$ or $g_2 \preceq g_1$ (where $\succeq$ denotes the binary preference relation "is weakly preferred to" and $\succ$ denotes "is strongly preferred to").

Transitivity. For any three gambles g_1, g_2 and g_3 in G, $g_1 \preceq g_2$ and $g_2 \preceq g_3 \Rightarrow g_1 \preceq g_3$.

Archimedean Axiom. For any three gambles g_1, g_2 and g_3 in G, where $g_1 \preceq g_2 \preceq g_3$, there exists $\alpha, \beta \in (0, 1)$ whereby $\alpha g_1 + (1 - \alpha)g_3 \prec \alpha g_2$ and $g_2 \prec \beta g_1 + (1 - \beta)g_3$.

Independence Axiom. For any two gambles g_1, g_2 and g_3 in G, $g_1 \preceq g_2$ if and only if $\alpha g_1 + (1 - \alpha)g_3 \preceq \alpha g_2 + (1 - \alpha)g_3$, $\forall \alpha \in [0, 1]$.

The Archimedean Axiom implies *continuity* so that the upper and lower contour sets of a preference relation over lotteries are closed. The *most* and *least* preferred lottery can be combined so that the compound is preferable to any intermediate gamble.

The Independence Axiom describes the *independence of irrelevant alternatives*. It implies that the preference between any two lotteries is unaffected by combining them in the same way with a third lottery. This enables the reduction of compound lotteries to simple lotteries.

The Expected Utility Property

A utility function u is said to have the *expected utility property* if, for a gamble g with outcomes $o_1, o_2, \ldots, o_n$ and effective probabilities $p_1, p_2, \ldots, p_n$ respectively, we have:

$$u(g) = p_1 u(o_1) + p_2 u(o_2) + \ldots + p_n u(o_n), \tag{2.1}$$

where $u(o_i)$ is the decision-maker's utility from outcome o_i. An individual who chooses one gamble over another if an only if the expected utility is higher is an expected utility maximizer. Von Neumann and Morgenstern proved that, as long as all the preference axioms described previously hold, then a utility function exists, and it satisfies the expected utility property.

Expected utility theory is widely used in theoretical and practical analysis. The insurance industry is an obvious example of an application domain. It is an imperfect theory and it has been demonstrated experimentally that people can sometimes violate the apparently rational underlying axioms. This paradoxical behavior is described by the Allais [1], Ellsberg [39] and Machina [80] paradoxes.

The St. Petersburg Paradox and Decreasing Marginal Utility

The St. Petersburg paradox relates to the valuation of an offer to gamble on the tossing of a fair coin that would be tossed continuously until it turned up tails [81]. If the coin produced a tails on the n^{th} toss, you would receive \$$2^n$, *i.e.* if tails first came up on the fourth toss, you would receive \$$2^4$ = \$16.

The probability of tails first appearing on toss n is $\frac{1}{2^n}$. The expected value of the gamble is thus:

$$\sum_{i-1\ldots\infty} \frac{2^i}{2^i} = \infty.$$

However, in reality nobody would be willing to pay a very large amount for the offer of this bet, even though the expected value of this gamble is infinity and herein lies the paradox. This is commonly referred to as the St. Petersburg Paradox, and was studied by Swiss mathematician Nicholas Bernoulli. In 1738 his cousin, Daniel Bernoulli, provided a solution to the problem [12]. Daniel Bernoulli's solution lay in the realization that people's "utility" for units of money was a subjective internal valuation that was determined by their current wealth. Bernoulli hypothesized that a person's utility from money was a logarithmic function of the amount of money, of the form:

$$u(x) = k\log(x) + c,$$

where x represents the amount of money, k is a parameter and c is a constant. Bernoulli used a logarithmic function to represent utility of wealth because its value increases at a decreasing rate as the value of its argument increases, *i.e.* a concave function whereby the first derivative is positive and the second derivative is negative ($u'(x) > 0, u''(x) < 0$). This is the famous idea of *diminishing marginal utility* [12].

There is an important distinction to be made between Bernoulli and von Neumann-Morgenstern utility functions. By convention, we use the term *Bernoulli Utility Function* to refer to a decision-maker's utility over wealth - since of course it was Bernoulli who originally proposed the idea that people's internal, subjective value for an amount of money was not necessarily equal to the physical value of that money [12]. The term *von Neumann-Morgenstern Utility Function* is used to refer

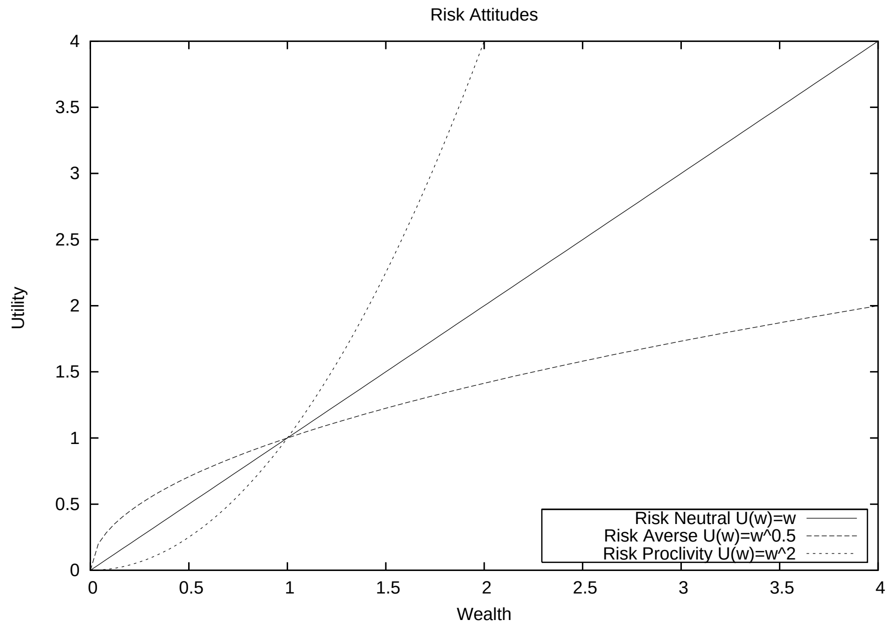

Figure 2.1: Risk attitudes.

to a decision-maker's utility over lotteries, or gambles. Preferences over lotteries may also be thought of as preferences over probability distributions [127].

2.4.2 Risk Aversion, Neutrality and Proclivity

A decision makers risk attitude characterizes his willingness to engage in lotteries or gambles that exhibit probabilistic returns. Risk attitudes can be categorized in three distinct ways: risk aversion, neutrality and proclivity. Risk aversion is the most common risk attitude displayed in commercial activity with diminishing marginal returns expected for increased wealth. Risk neutral agents display an objective predilection for risk whilst agents with a proclivity for risk are willing to engage in gambles where the utility of the expected return is less than the expected utility.

Consider a gamble g_c where a fair coin is tossed and, there is a 0.5 probability of receiving w_1 or w_2, where $w_1 < w_2$, is dependent upon the outcome of the toss. The question of valuing the right to this gamble depends upon a person's risk attitude. Figure 2.1 illustrates how different risk attitudes can be modeled as Bernoulli utility functions. Our main interest in this work is risk aversion, but we also outline the concepts of risk neutrality and proclivity below.

Risk aversion. If an individual's utility of the expected value of a gamble is greater than their expected utility from the gamble itself, they are said to be risk averse. The intuition behind risk aversion being that when faced by comparable returns, agents tend to choose the option that involves less risk. This constitutes a more precise definition of Bernoulli's principle. Risk aversion is described by a concave Bernoulli utility function, such as a square root or logarithmic function. Consider the gamble g_c being offered to a risk averse person, whose Bernoulli utility function takes the form $u(w) = w^{\frac{1}{2}}$, where w was the outcome. The expected utility (from Equation 2.1 [127]) for g_c is:

$$E(u) = (0.5 \times \sqrt{w_1}) + (0.5 \times \sqrt{w_2}) = \frac{\sqrt{w_1} + \sqrt{w_2}}{2},$$

whilst the utility of the expected value of the gamble is:

$$u[E(w)] = \sqrt{\frac{w_1 + w_2}{2}}.$$

Figure 2.2 illustrates a concave utility function. The gamble involves a fair coin toss, therefore, $E(w)$ lies halfway between w_1 and w_2. The expected utility, $E(u)$, is shown by point E on the chord connecting A=$w_1, u(w_1)$ and B=$w_2, u(w_2)$ is less than the utility of the expected outcome, $u[E(w)]$. The certainty equivalent may be considered as another lottery whose outcome involves no risk. It results in an income of $CE(w)$ with absolute certainty where $u(CE(w)) = E(u)$. However, when the Bernoulli utility function is concave (i.e risk averse), $CE(w) < E(w)$. The agent is indifferent between the two lotteries and the difference between the two values:

$$\pi(w) = E(w) - CE(w) \tag{2.2}$$

is known as the *risk premium*. This is the maximum amount the agent is willing to pay for a risk free lottery. It follows that a risk averse person's certainty equivalent will always be less than the expected value of the gamble, and the risk-premium

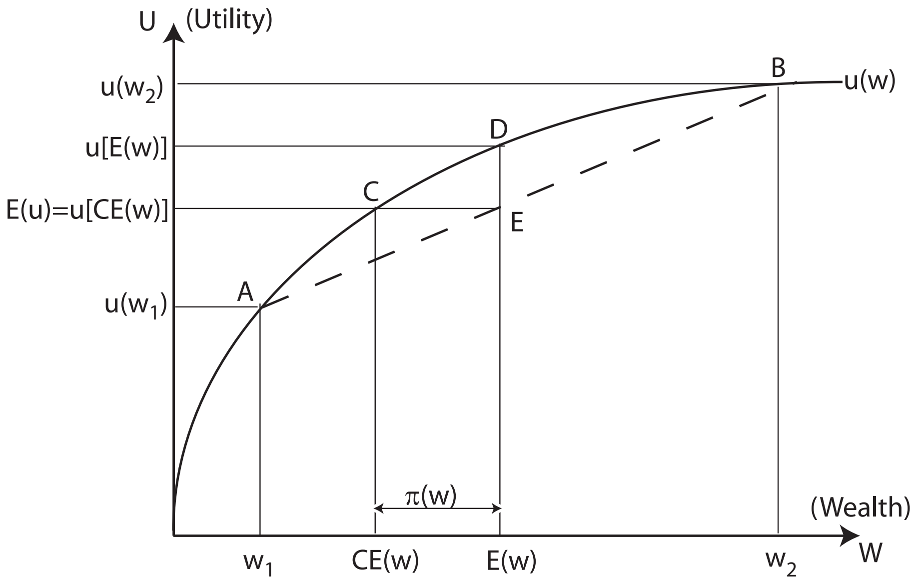

Figure 2.2: Risk aversion and certainty equivalence

will have a positive value.

Risk neutrality. If an individual's utility of the expected value of a gamble is exactly equal to their expected utility from the gamble itself, they are said to be risk neutral. Risk neutral behavior is captured by a linear Bernoulli function. For the above gamble, a risk neutral individual whose Bernoulli utility function takes the form $u(w) = 2w$ will have an expected utility over the gamble of:

$$E(u) = (0.5 \times 2 \times w_1) + (0.5 \times 2 \times w_2) = w_1 + w_2,$$

while their utility of the expected value of the gamble is

$$u[E(w)] = 2 \times \frac{w_1 + w_2}{2} = w_1 + w_2.$$

A risk neutral agent will have a zero risk premium, and a certainty equivalent equal to the expected value of the gamble.

Risk proclivity (risk loving). If an individual's utility of the expected value of a gamble is less than their expected utility from the gamble itself, they are said to be risk-loving. Note, however, that this does not capture normal gambling behavior of the kind observed in casinos the world over. By this definition, a truly risk-loving individual ought to be willing to stake all of their assets, everything they own, on a single roll of dice, since any uncertain outcome is preferred to any certain one. A convex Bernoulli utility function captures risk-loving behavior, for example, an exponential function. For the above gamble, a risk-loving person whose Bernoulli utility function took the form $u(w) = w^2$ would have an expected utility over the gamble of:

$$E(u) = (0.5 \times w_1^2) + (0.5 \times w_2^2) = \frac{w_1^2 + w_2^2}{2},$$

while their utility of the expected value of the gamble is:

$$u[E(w)] = (\frac{w_1 + w_2}{2})^2.$$

A risk-loving agent will have a negative risk-premium, since he would need an extra incentive to accept the certain outcome, not the risky gamble, and the certainty equivalent would be greater than the expected value of the gamble.

Jensen's inequalities [68] summarize risk attitudes as follows:

- $E(u(w)) < u(E(w)) \Rightarrow$ risk aversion,
- $E(u(w)) = u(E(w)) \Rightarrow$ risk neutrality,
- $E(u(w)) > u(E(w)) \Rightarrow$ risk proclivity.

Friedman and Savage [49] noted that it is not necessarily true that an individual's utility function has the same kind of curvature everywhere; there may be levels of wealth, for example, where inflection points may occur that infer risk loving takes over from risk aversion. Risk attitudes typically depend upon the wealth of an agent. An insurance company is risk averse but tends towards risk neutrality as its relative wealth grows. Purchasers of insurance are naturally more risk averse because of their inability to withstand large losses. The insurance company earns profits since the value of the premiums it receives is at best equal to, most probably higher than, the expected value of the loss.

2.4.3 Measuring Risk Aversion

The most famous measures for risk aversion were introduced by Arrow [7] and Pratt [106]. The Arrow-Pratt measure of risk aversion is:

$$r(w) = -\frac{-u''(w)}{u'(w)} > 0 \tag{2.3}$$

where $r(w)$ is referred to as the coefficient of *absolute risk aversion* (ARA). Highly risk averse agents will have a greater curvature in $u()$. A person with a constant ARA ($r(w) = c$) only cares about absolute losses, regardless of their level of wealth. As the size of a gamble grows, people generally become more risk averse. Consider a coin toss with winnings of $1000 \times w_1$ and $1000 \times w_2$. It is natural to assume a greater level of risk aversion, which implies that $\frac{\partial r(w)}{\partial w} < 0$. A natural extension of the ARA thus becomes *Arrow-Pratt measure of relative risk aversion* RRA:

$$rr(w) = -w \times \frac{-u''(w)}{u'(w)}. \tag{2.4}$$

If the RRA of a person is a constant, $rr(w) = c$, a person will pay a constant *share* of wealth to attain the certainty equivalent over a given *fraction* of their income.

The relative risk aversion (RRA) coefficient of a person whose consumption is close to the subsistence level may be very high. For example, if consumption is barely above the subsistence level, no risk taking may be tolerable. If this is the case, then the RRA coefficient must be a decreasing function of wealth for poorer individuals. Therefore, decreasing RRA (DRRA) is a realistic measure of risk aversion:

$$\frac{\partial \left(\frac{-w \times u''(w)}{u(w)} \right)}{\partial w} \leq 0. \tag{2.5}$$

The measurement of risk aversion is most commonly associated with the insurance industry and the calculation of premia. Customers of insurers must also calculate how much insurance they should purchase given the fact that insurance companies are profit seeking and premia are not generally *actuarially fair*, *i.e.* the premium is not equal to the expected claims.

2.5 Summary

We provided background information concerning the field of auction theory and gave examples of basic auction forms and how bidders behave in a game-theoretic manner in such auctions depending upon whether their valuation model is private, common or affiliated. We also described how mechanism design can be utilized by the bid-taker to design auctions so that social objectives can be met when bidders strategize according to game-theoretic principles.

We then introduced CAs, that provide a means for bidders to describe perceived complementarities between multiple distinguishable items for sale. This helps avoid the problem of inefficient bundling of items into single item auctions because the bid-taker cannot know in advance the optimal bundling strategy to gain maximum revenue. Unfortunately, winner determination poses considerable computational problems.

We are interested in addressing solution uncertainty in CAs. Winning-bid withdrawal can pose a considerable risk to the bid-taker. We need to be able to make decisions about how the bid-taker should make choices in the presence of such uncertainty and in particular when he:

1. values outcomes using a concave (risk averse) Bernoulli utility function [12],
2. has preferences about the risk profiles of those outcomes.

Von Neumann and Morgenstern proposed a model for understanding and analyzing risk preferences, expected utility theory [127]. In the presence of uncertainty surrounding solutions to CAs, risk management is necessary to hedge against the danger of encountering objectionable losses in revenue for bid-takers with different levels of risk aversion. We aim to quantify risk in auction solutions using expected utility theory and, therefore, make an accurate and informed decision regarding what outcome is best for the bid-taker.

Chapter 3 presents potential sources of risk in solutions for such auctions and describes the computational difficulties involved in reducing these risks.

Chapter 3

Bid-taker's Exposure Problem

In this chapter we present the hitherto unrecognized "bid-taker's exposure problem" for combinatorial auctions in which winning-bid withdrawal is a possibility. We argue that non-withdrawal rules in auctions constitute a fallacy because enforcement of such rules may be impossible. The root cause of the failure may be outside the control of the bidder, *e.g.* natural disaster or competition regulator intervention. Bid withdrawal is, therefore, an unavoidable threat to revenue and the bid-taker should consider this risk when choosing an allocation of items to bidders.

We demonstrate the significance of this problem using auction simulation tools that replicate bidding in different types of economically motivated scenarios. We also show that this problem is exacerbated for a risk averse bid-taker.

Risk management for this problem is computationally challenging and we discuss the merits of proactive and reactive approaches to combating such uncertainty. The search for a solution that maximizes expected utility for a risk averse bid-taker is not straightforward. We examine the feasibility of conventional integer linear programming approaches to tackling this problem and outline a goal of establishing robust solutions that can better withstand the risk posed by winning-bid withdrawal in combinatorial auctions.

3.1 Problem Description

Combinatorial auctions are gaining popularity in sectors such as supply chain management where trades are more frequent and of lower value [61, 102, 135]. In such auctions, bidder reliability may be questionable and some bidders may be more trusted than others.

Revenue is the most obvious optimization criterion for such auctions, but another desirable attribute is *solution robustness*. This quality, in its most general sense, can be considered as a solution's ability to withstand future uncertainty. In terms of CAs, the uncertainty may surround the bid amounts, sets of desired items, sets of items for sale or the ability of the bidders to execute the transaction. The work in this dissertation concerns the latter form of uncertainty, which may lead to the withdrawal of a winning bid. We argue that potential winning-bid withdrawal is the most serious uncertainty facing the bid-taker because other uncertainties such as the alteration of a bid amount may be deemed illegal by auction rules, resulting in disqualification. It may not be practicable for the bid-taker to enforce non-withdrawal rules. The cause of failure in auction solutions may be outside the control of the bidder who instigates a bid withdrawal. The myriad possible reasons for bid withdrawals include natural disasters, competition-regulator intervention or the failure of a supplier's sub-contractor to fulfil a dependent obligation. Auction rules regarding non-withdrawal of bids may constitute a fallacy because of the inability to impose absolute enforcement. We, therefore, concentrate upon winning-bid withdrawal as the most serious threat to solution stability and revenue.

We focus upon the withdrawal of winning, rather than losing, bids only for the following reasons. When a winning bid is withdrawn, there is a loss in revenue for the bid-taker, who is left with unallocated items that may be sought-after by other bidders. There is an opportunity for the bid-taker to reduce the resulting loss in revenue by allocating these items to another bidder in a "repair" solution. However, if a losing bid is withdrawn, there is no need for the bid-taker to repair the solution because all items are already allocated.

We do *not* ensure that repair solutions are robust, but instead find an initial solution that is robust for the following reason. Given that the probability of bid

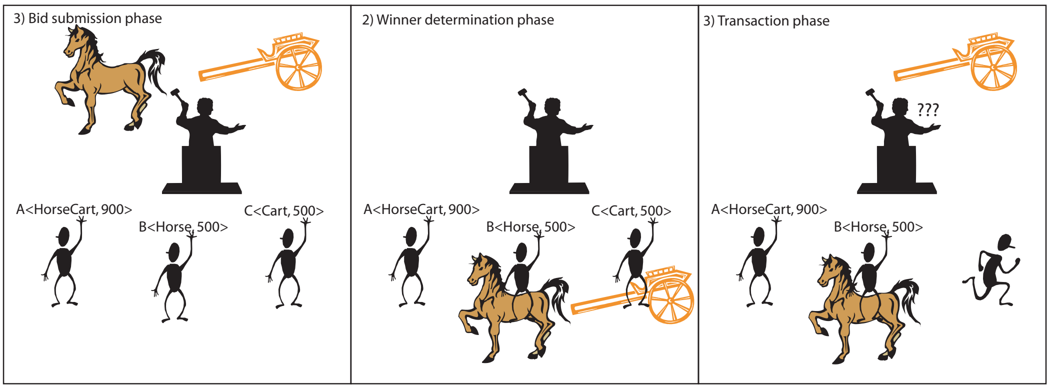

Figure 3.1: Bid-taker's exposure problem

withdrawal is low, (*i.e.* low enough for almost all prior work in the domain to assume it is zero), we contend that it is sufficient to ensure the reparability of the initial allocation without ensuring that repair solutions are also robust. The probability of a withdrawal by a winning-repair bid, conditional on a prior withdrawal by a winning-bid, tends towards zero.

Example 3.1.1. *Consider a simple CA where a horse and cart are for sale. Bidders can bid on either the combination of horse and cart or each item in isolation. Figure 3.1 shows how three bids,* $A\langle HorseCart, 900\rangle$, $B\langle Horse, 500\rangle$ *and* $C\langle Cart, 500\rangle$, *were entered in the bid submission phase.*

The second phase in an auction involves winner determination. Traditionally, this has focused upon maximizing revenue using the set packing or partition formulation, depending upon the rules governing free disposal [33, 119]. The optimal solution to the WDP in this case is to award the horse and cart to bids B *and* C*, respectively. However, a difficulty may arise in the transaction phase if either of the winners renege upon their bid and do not complete the transaction. In this example, bid* C *is withdrawn and the bid-taker is left with a cart that is valueless to all remaining bidders. The bid-taker is unable to revoke the award of the horse to* B*. The bid-taker's exposure problem is realized because instantiation of a repair solution of satisfactory revenue is impossible (assuming that 500 constitutes an unsatisfactory auction revenue).* $\triangle$

We seek a *robust solution* that extenuates the effects of winning-bid withdrawal (a *break* in the solution) by facilitating the formation of a repair solution of adequate revenue, without causing undue disturbance to other bidders. A *brittle solution* to a CA is one in which an unacceptable loss in revenue is unavoidable if a winning bid is withdrawn, as we have seen in Example 3.1.1. In such situations the bid-taker may be left with a set of items deemed to be of low value by all other bidders. These bidders may associate a higher value for these goods if they were combined with items already awarded to others, hence the bid-taker is left in an undesirable local optimum in which a form of backtracking is required to reallocate the items in a manner that results in sufficient revenue.

We have called this problem the "bid-taker's exposure problem" because it bears similarities to the "exposure problem" faced by bidders seeking multiple items in separate single-unit auctions but holding little or no value for a subset of those items. Such a bidder may end up winning a subset of desired items, thereby incurring a loss. This "exposure problem" causes depressed bidding, as described in § 2.3. Combinatorial auctions resolve this problem by permitting bids on sets of items [85], but when bid-withdrawal is considered possible, the bid-taker becomes exposed to the problem of holding items that do not complement one another and would have returned greater revenue if combined with other items already sold in the auction.

Definition 3.1.1 (Bid-taker's exposure problem for combinatorial auctions)**.** *The problem faced by a bid-taker after one or more winning bids are withdrawn, whereby a solution of satisfactory revenue cannot be attained because of the irrevocable awarding of items to other winning bidders.*

Reallocating items may be regarded as disruptive to a solution in many real-life scenarios. Consider an example where industrial procurement is conducted using a CA. It would be highly undesirable to revoke contracts from a group of suppliers because of the failure of a third party. A robust solution that is tolerant of such breaks is preferable. Robustness may be regarded as a preventative measure protecting against future uncertainty by sacrificing revenue in return for solution stability and reparability. We assume a probabilistic approach whereby the bid-taker has knowledge of the reliability of bidders from which the likelihood of an incomplete transaction may be inferred. The assumption of knowledge over probability distributions for bidders valuations, or *priors*, has traditionally been used in economics to choose reservation prices for optimal auction design [20, 90]. On these grounds, we contend that prior knowledge of bidder reliability is also a reasonable assumption.

3.1.1 Related Work

Although the bid-taker's exposure problem has not received prior attention, there is some related work concerning auctions and risk management that we mention here.

Waehrer *et al.* [128] analyzed the preferences of a risk averse bid-taker over the class of standard auctions, presented in § 2.1.1, as opposed to preferences over solutions for a given auction type in our case. Risk averse sellers dislike variance in expected revenue, therefore the Revenue Equivalence Theorem no longer applies. They found that a sealed-bid first-price auction with an appropriately set reserve price was preferable to all other auction formats. They also found that in both first- and second-price auctions, a risk averse seller optimal reserve price is lower.

Harstad and Rothkopf [57] presented withdrawable bids as a form of insurance against the winner's curse. They showed that in a sealed-bid auction model with common or strongly affiliated values for a single item, the possibility of withdrawing a winning bid for a given penalty can provide sufficient conditions for increased bidder aggressiveness. They show that the bid-taker is better off on average providing such insurance.

Kastner *et al.* [70] examined how to handle perturbations in CAs given a solution whilst minimizing necessary changes to that solution. These perturbations may include winning-bid withdrawals, change of valuation/items of a bid or the submission of a new bid. They looked at the problem of finding incremental solutions to restructure a supply chain whose formation is determined using combinatorial auctions [131]. Following a perturbation in the optimal solution they proceed to impose involuntary item withdrawals from winning bidders. They formulated an incremental integer linear program (ILP) that sought to maximize the valuation of the repair solution whilst preserving the previous solution as much as possible. This contrasts to our proposed proactive approach where we seek to find an initial solution that is robust against such failures.

Porter [103] examined the FCC's ascending multi-round auction format that permitted bid withdrawals. The FCC permitted withdrawal for a given penalty in multiple single item auctions. This design was intended to reduce the effects of the exposure problem experienced by bidders who try to accumulate a package of items in a piecemeal fashion. Porter examined the trade-off between revenue and efficiency. He found that the increased efficiency does not outweigh the higher prices paid so that bidder surplus falls in the presence of the withdrawal rule [103].

3.1.2 Winning-bid Withdrawal

We assume an auction protocol with a three-stage process involving a) the submission of bids, b) winner determination, and finally c) a transaction phase, as witnessed Example 3.1.1. We are interested in winning-bid withdrawals that occur between the announcement of winning bids and the end of the transaction phase. All bids are valid until the transaction is complete, so we anticipate an expedient transaction process.

Note that in some instances the transaction period may be so lengthy that consideration of non-winning bids as still valid may not be fair. Breaks that occur during a lengthy transaction phase are more difficult to remedy and may require a subsequent auction. For example, if the item is a service contract for a given period of time and the break occurs after partial fulfilment of this contract, the other bidders' valuations for the item may have decreased in a non-linear fashion.

An example of a winning-bid withdrawal occurred in an FCC spectrum auction [137]. Of course, the impact of such withdrawals depends upon other bidders valuations. In this case, it did not materially affect revenue because another bid eventually superseded the withdrawn bid. However, large revenue losses are possible for bid-takers, as occurred when book publisher Scholastic withdrew a winning bid of $8m for inventory in a bankruptcy sale that resulted in subsequent revenue of only $5.4m [40].

Withdrawals, or *breaks*, may occur for various reasons. Winning-bid withdrawal may be instigated by the bid-taker when Quality of Service agreements are broken or payment deadlines are not met. We refer to winning bid revocation by the bid-taker as item withdrawal in this dissertation so as to distinguish between the actions of a bidder and the bid-taker. Harstad and Rothkopf [57] outlined several reasons for breaks in single item auctions that include:

1. an erroneous initial valuation/bid;
2. unexpected events outside the winning bidder's control;
3. information obtained or events that occurred after the auction but before the transaction that reduce the value of an item;
4. a desire to have the second-best bid honored;
5. the revelation of competing bidders' valuations infers reduced profitability, a problem known as the "*winner's curse*".

Reasons 1, 2 and 3 are non-strategic causes for bid-withdrawal that are unaffected by the revelation of winning bids. The bid-taker cannot factor the likelihood of such failures upon the outcome, so the probabilities of failure are exogenous. However, reasons 4 and 5 describe events that are outcome-dependent. Therefore, it is possible to apply game-theoretic analysis to analyze the probabilities of failure due to a desire to have the second-best bid honored, or the winner's curse, given prior valuation distributions. It is possible to circumvent the possibility of a bidder benefiting from withdrawal by disallowing the award of items in a repair solution to the guilty party. However, strategic bid withdrawal can still occur because of the winner's curse.

3.1.3 Strategic Bid Withdrawal and the Winner's Curse

The winner of an item whose value is unknown is ultimately the bidder with the highest bid. If the winner has bid truthfully, he therefore has the most optimistic valuation. All other bids in the auction are all less than those of the winners, so if he can infer that the true valuation for the item was probably less than his payment for the item, he may have overestimated the value of the item. This is known as the winner's curse. Knowledgeable bidders compensate for this effect by bidding below their true valuation.

Definition 3.1.2 (Winner's Curse)**.** *The winner's curse states that the winning bidder in an auction with common-values (i.e. the value of the item to all bidders is the same) may bid an amount that exceeds the item's intrinsic value, thereby suffering a loss.*

Wilson [139] first identified the winner's curse problem that is associated with auctions in which bidders have uncertainty surrounding the valuation of the item†. This problem was initially of most interest to the oil industry because auctions for drilling rights display a common values model in which the intrinsic value of the oil rights is not known accurately.

Upon purchase of an item, it is usually the case that other potential buyers have decided against purchasing and have lower valuations for the item than the winner. A winner may reassess the value of the item, upon reflection, and question their original bid. The possible answers depend on the reasons for buying the item.

If the only reason for purchasing the item was the satisfaction of owning it, and there was no possibility of resale, then the winner can be sure that he did not pay more for the property than it is worth so long as the purchase price did not exceed his own valuation. This is an example of independent private values, as discussed in § 2.1.2, and implies that the winner's curse is not applicable.

The winner's curse is most clearly evident in auctions where bidders have common or strongly affiliated values models [21, 69]. If the item was purchased with the intention of resale, other bidders' valuations are an important indicator of potential resale value. The more bidders that a winner has outbid, the more worrying may be the fact that he has won. The reason for this concern is the following. Suppose that a bidder was competing against a group of other potential buyers for an item and the precise intrinsic value is unknown, as in auctions for oil drilling rights [21], but the bidders have estimated its value. This scenario occurs in many economic situations, even markets as diverse as the transfer system for baseball players [16, 23]. Each bidder has a probability distribution over the uncertain value of the item. Typically, some of the estimates will be lower than the actual value of the item and some will be higher. In a competitive environment, such as an auction, the buyer will usually have the highest estimate of value. This

†Capen *et al.* [21] coined the term "winner's curse" in a subsequent paper on bidding strategies in auctions for drilling rights.

highest estimate of the value is likely to exceed the intrinsic value of the item. Bidders avoid the significant risk of paying too much for an item by discounting his bid. It may be impossible to know in advance whether an estimate is high or low, so all bidders must mitigate against the risk of the winner's curse. Determination of optimal bidding strategies whilst considering the winner's curse can be a difficult problem for bidders and is based in the properties of order statistics [21] using numerical integration in single item auctions.

Significantly, however, Harstad and Rothkopf [57] showed that for single item auctions the bid-taker can expect higher revenues on average when bid withdrawal is permitted because the effects of bid discounting due to the winner's curse are reduced. They used a game-theoretic analysis that assumed the only reason for withdrawal was the winner's curse. Kagel and Levin found that withdrawable bid auctions are attractive in markets where bidders are unable to compensate fully for the effects of bid withdrawal [69]. Harstad and Rothkopf conclude that exogenous probabilities of failure increase the need for possible bid withdrawal by risk averse bidders, thereby making it possible for the bid-taker to provide a form of insurance to the bidder for a given premium [57].

3.1.4 Solution Brittleness in Combinatorial Auctions

An analysis of the sensitivity of optimal allocations to winning-bid withdrawal provides a strong motivation for the research in this dissertation. A closed-form analytical solution to expected loss in revenue following withdrawal is not feasible for combinatorial auctions because of the highly non-trivial nature of finding solutions. The WDP is $\mathcal{NP}$-complete [110] and inapproximable [115]. A Monte-Carlo style simulation technique is required to gauge losses in revenue from winning-bid withdrawal.

We performed a sensitivity analysis by simulating auction problems using a bid simulation tool. We computed the revenue of optimal solutions and then examined the effects of a single winning bid being withdrawn in turn. After a winning bid is withdrawn we attempted to repair the solution by awarding all remaining unsold items to bids that are only interested in a subset of those items. We assumed that the bid-taker could not revoke items already awarded to other winning bidders. In later chapters, we shall attempt to relax this assumption in a reasonable manner so that we can improve solution reparability.

We used the Combinatorial Auction Test Suite (CATS) [78] to generate sample auction data. We created 100 instances of auctions in which there are 20 items for sale and 100-2000 bids that may be dominated in some instances†. A dominated bid is one whose set of items is the same as, or a superset of, another whose bid amount is less than this bid. It, therefore, cannot partake in an optimal solution to the WDP that does not consider solution robustness. CATS uses economically motivated bidding patterns to generate auction data in various scenarios, as described in § 2.3.2. We performed sensitivity analysis upon the following four bid distributions: airport take-off/landing slots (`matching`), electronic components (`arbitrary`), property/spectrum-rights (`regions`) and transportation (`paths`). We chose these distributions because they represent a diverse selection of bidding patterns in different application domains.

The method used is as follows. Firstly, we determined the optimal solution using `lp_solve`, a mixed integer linear program solver [11]. We then simulated a single winning-bid withdrawal and re-solved the problem with the other winning bids remaining fixed, *i.e.* there were no involuntary dropouts. The optimal repair solution was then determined. All bids are assumed to be brittle, so this process is repeated for all winning bids in the overall optimal solution.

Figure 3.2 shows the average revenue of such repair solutions as a percentage of the optimum. Also shown is the average worst-case revenue following a single break in each auction. We implemented an auction rule that disallows bids from the reneging bidder from participation in a repair. This removes the possibility of a bidder strategically withdrawing a bid so that a lower bid from the same bidder can win instead. This addresses one of the possible strategic reasons for bid withdrawal listed previously in §3.1.2. We assumed that all bids in a given XOR bid were from the same bidder. An XOR bid in the CATS simulation tool incorporates the same dummy item in the set of desired items for a set of bids to explicitly state that only one of these bids may win. CATS provides no other

†The CATS flags included int_prices with the bid_alpha parameter set to 1000. The normal distribution of values for items was used for the `arbitrary` (`arbitrary-npv`) and `regions` (`regions-npv`) distributions.

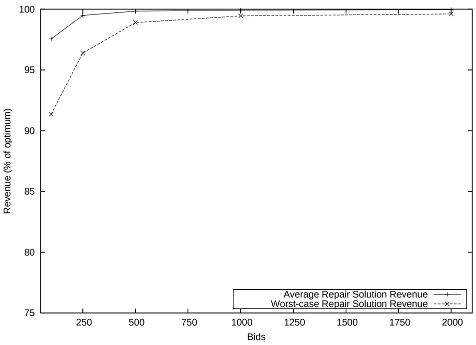

(a) `paths`

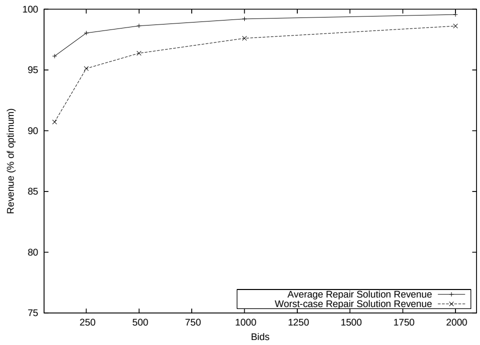

(b) `matching`

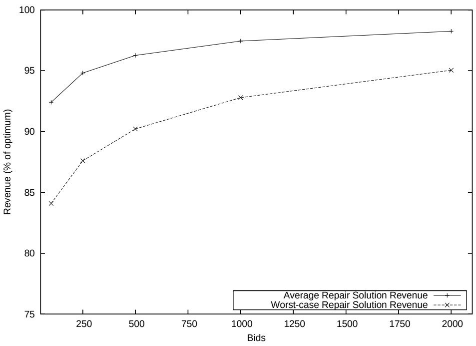

(c) `regions`

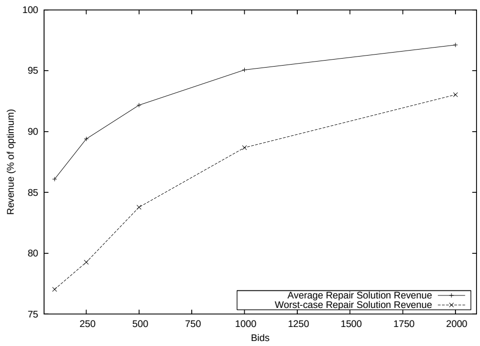

(d) `arbitrary`

Figure 3.2: Sensitivity of bid distributions to single winning-bid withdrawal.

means of differentiating between bids from separate bidders.

Figure 3.2(a) illustrates how the `paths` distribution is inherently the most robust distribution since when any winning bid is withdrawn the solution can be repaired to achieve an average of over 98.5% of the optimal revenue for auctions with more than 250 bids. However, there are some cases when such withdrawals result in solutions whose revenue is significantly lower than optimum. Even in auctions with as many as 2000 bids there are occasions when a single winning-bid withdrawal can result in a drop in revenue of over 5%, although the average worst-case drop in revenue is only 1%. Figure 3.2(b) shows how the `matching` distribution is more brittle on average than `paths` and also has an inferior worst-case revenue on average. This trend continues as the `regions` (Figure 3.2(c)) and `arbitrary` (Figure 3.2(d)) distributions are more brittle still. These distributions are clearly sensitive to bid withdrawal when no other winning bids in the solution may be involuntarily withdrawn by the bid-taker.

3.2 Risk Averse Bid-taker

In this Section we consider at the effects of the bid-taker's risk attitude on solution preference in CAs when the bid-taker's exposure problem is considered. The Revenue Equivalence Theorem, as described in § 2.1, stated that the expected revenue of all four basic auction types for a single item is identical. This principle assumes that the bid-taker is risk neutral. Given a risk averse bid-taker, however, minimal variance of expected revenues is desirable. Observe that the winner in a second-price or ascending auction pays the price dictated by the runner-up, and, by revenue equivalence, must bid the expectation of this price in a first-price auction. The revenue in a first-price auction is fixed according to the winning bid, but has the same mean as the second-price auction. An additional source of risk in the second-price auction is that it is unconditional on the first-price and, therefore, exhibits greater variance. The strategic actions of bidders serve to reduce the variance of bidding patterns in first-price auctions. This implies that there is a greater level of risk for the bid-taker in the second-price auction. For a complete analysis of preferred auction forms for a risk averse bid-taker, the reader is referred to [128].

The sensitivity analysis conducted in § 3.1.4 indicated that the revenue loss for repair solutions given a winning-bid withdrawal in an optimal solution can be substantial. Therefore, in sealed-bid combinatorial auctions, considerable risk is posed by such uncertainty [62, 70]. We seek to address this risk by finding robust solutions whose repair revenue variance is less than that of the optimal solution so that expected utility can be increased.

3.2.1 Risk Management for Combinatorial Auctions

When a risk averse bid-taker faces a choice of solutions incorporating various probabilities of bid withdrawal, this may be viewed as a choice over lotteries. We learned in § 2.4 that von Neumann and Morgenstern's expected utility theory can be used to maximize expected utility for a given risk attitude and, therefore, govern the choice of solution.

We consider a CA with an optimal solution to the WDP, S_o, and v winning bids whose probability of withdrawal is p_i, where $i \in \{1, \ldots, v\}$. Post-withdrawal solutions cannot confiscate items from non-reneging winners, therefore, the set of items previously awarded to the withdrawn bid, b_i, must be redistributed to form a repair solution, S_{r_i}. For simplicity, let us initially only consider the possibility of a single winning-bid withdrawal. The von Neumann-Morgenstern [127] expected utility model implies that we should seek to maximize the expected utility:

$$E(u) = p_{nw} \times u(S_o) + \sum_{i=1}^{v} p_i \times u((S_{r_i}), \tag{3.1}$$

where $p_{nw} = (1 - \prod_{i=1}^{v} p_i)$ is the probability of no withdrawal†. The utility function of the bid-taker, u, determines the risk attitude of the bid-taker. In this dissertation we examine the preferences of both risk neutral and risk averse bid-takers. The optimal solution to the WDP in terms of revenue, S_o, is no longer necessarily optimal in the presence of bid withdrawal, even for a risk neutral bid-taker. A solution whose repair solutions have higher revenue may return a higher expected revenue than S_o.

Example 3.2.1. *Consider an auction whose bids comprise those listed in Table 3.1. We seek the optimal expected revenue in the situation where a single bid may be withdrawn, assuming that the bid-taker is risk neutral. There are three submitted bids for items A and B, the third being a combination bid for the pair of items at a value of 190. The optimal solution has a value of 200, with the first and second bids as winners. When we consider the probabilities of failure, in the fourth column of Table 3.1, the problem of which solution to choose becomes more difficult.*

Computing the expected revenue for the solution with the first and second bids winning the items, denoted $\langle 1, 1, 0 \rangle$*, gives:*

$$(200 \times 0.9 \times 0.9) + (2 \times 100 \times 0.9 \times 0.1) + (190 \times 0.1 \times 0.1) = 181.90.$$

If only a single bid can be withdrawn there is probability of 0.18 of a revenue of 100, given the fact that we cannot withdraw an item from the other winning

†Note that the conditional probability of multiple withdrawals is 0, so that only a single bid is likely to be withdrawn.

Table 3.1: Example combinatorial auction.

	Items			
Bids	A	B	AB	Withdrawal prob
1	100	0	0	0.1
2	0	100	0	0.1
3	0	0	190	0.1

bidder. The expected revenue for $\langle 0, 0, 1\rangle$ *is:*

$$(190 \times 0.9) + (200 \times 0.1) = 191.00.$$

We can say that $\langle 0, 0, 1\rangle$ *is preferable to* $\langle 1, 1, 0\rangle$ *based on expected revenue, therefore, preferable for a risk neutral bidder.* △

Risk aversion accentuates the problem of suffering low utility from repair solutions. Variance in potential revenue constitutes a risk, and consequently risk averse bid-takers prefer an initial solution whose repairs display a smaller deviation in revenue from the initial solution.

Example 3.2.2. *Consider Example 3.2.1, where the bid amount for bid* 3 *is instead 175. The expected revenue of* $\langle 1, 1, 0\rangle$ *(181.75) becomes greater than that of* $\langle 0, 0, 1\rangle$ *(177.50). However, a risk averse bid-taker may prefer the latter solution because the revenue is never less than 175, but the former solution returns revenue of only 100 with probability 0.18. The expected utility of* $\langle 1, 1, 0\rangle$ *for a risk averse bid-taker whose Bernoulli utility function is* $u(w) = \sqrt{w}$*, is:*

$$(\sqrt{200} \times 0.9 \times 0.9) + (2 \times \sqrt{100} \times 0.9 \times 0.1) + (\sqrt{190} \times 0.1 \times 0.1) = 13.39,$$

whereas the utility of $\langle 0, 0, 1\rangle$ *is:*

$$(\sqrt{190} \times 0.9) + (\sqrt{200} \times 0.1) = 13.81.$$

A risk averse bid-taker with such a risk attitude would prefer the second solution, preferring to sacrifice revenue for reduced risk. △

Example 3.2.2, shows how the risk attitude of the bid-taker governs the selection of utility maximizing solutions. Unfortunately, finding robust solutions poses computational challenges that we shall examine in § 3.2.2.

3.2.2 Computational Feasibility

Maximizing expected utility via the von Neumann Morgenstern expected utility property is a combinatorial optimization problem. If we restrict the probabilities of failures to be conditional so that the probability of two or more failures is zero, we can approximately maximize expected utility. An integer programming formulation for this problem, assuming only a single, winning-bid withdrawal, is the following:

$$\max\left(\sum_{j=1}^{n}(1-p_j)u(v_j)x_j + \sum_{k=1}^{n}\sum_{j=1}^{n}p_k x_k u(v_j) r_{kj}\right) \tag{3.2}$$

subject to the following constraints:

$$\sum_{j|i\in S_j}^{m} x_j \leq 1, \forall i \in \{1,\ldots,m\},\ x_j \in \{0,1\}, \tag{3.3a}$$

$$\sum_{j|i\in S_j, k\neq j}^{m} x_j + r_{kj} \leq 1, \forall i \in \{1,\ldots,m\}, \tag{3.3b}$$

$$\forall k \in \{1,\ldots,n\}, x_j, r_{kj(k\neq j)} \in \{0,1\}, r_{kk} \in \{0\},$$

where p_j is the probability of bid j being withdrawn, x_j the decision variable for bid j, v_j the amount of bid j and r_{kj} is the repair value of bid j given a withdrawal of bid k. Therefore, r_{kj} indicates whether or not bid j is successful in the repair solution following the withdrawal of the winning-bid k. Equation 3.2 seeks to maximize expected utility according to von Neumann-Morgenstern's expected utility property with the conditionally dependent probability of multiple bid withdrawals set to zero.

A major disadvantage of this approach is that $\mathcal{O}(n^2)$ variables are necessary to represent the status of n bids in repair solutions for n potential bid withdrawals, thus making the formulation impractical. The constraints in Equation 3.3a are the same as those in the WDP and stipulate that bids sharing the same item have only a single winning bid in that set. Equation 3.3b describes constraints for repair solutions after bid j is withdrawn. No repair bid may be awarded an item already assigned to another non-reneging bidder in the original solution. Also, the repair solution for bid x_k failing has $r_{kk} = 0$ because this bid is withdrawn. In this formulation, $\mathcal{O}(mn)$ constraints are necessary. The space complexity may be further increased by considering multiple, y, breaks that require $\mathcal{O}(n^{y+1})$ variables and $\mathcal{O}(mn^y)$ constraints.

When a bid is withdrawn there are restrictions on how the solution can be repaired. If the bid-taker is able to revoke the awarding of items to other bidders, the solution can be repaired easily by reassigning all the items to the optimal solution of the auction without the withdrawn bid. Alternatively, the bidder who reneged upon a bid may have all his other bids disqualified and the items could be reassigned based on the optimum solution without that bidder present. However, the bid-taker is often unable to freely reassign the items already awarded to other bidders. The optimization problem of Equation 3.2, subject to the constraints of 3.3, applies the same auction rules as were used in the sensitivity analysis of § 3.1.4. When items cannot be withdrawn from winning bidders, following the failure of another bidder to honor his bid, this imposes a restriction on reparability. A repair solution includes all non-reneging winning bids and a WDP decides upon a set of winning bids amongst all bids whose items only include those in the withdrawn bid or those that were not allocated in the original solution. We are free to award items to any of the previously unsuccessful bids when finding a repair solution.

Determining the maximum expected utility in the presence of such uncertainty becomes computationally infeasible as the number of brittle bids grows. A WDP needs to be solved for all possible combinations of bids that may fail. If any k bids may fail, we need to compute $\binom{n}{k}$ possible repair solutions.

PROPOSITION 1. *Finding a repair solution that maximizes expected revenue following a single winning-bid withdrawal is $\mathcal{NP}$-complete.*

Proof. The following proof follows trivially from the complexity of the WDP. The set of items in the withdrawn bid, $\mathcal{M}$, is reassigned to other bids whose desired set of items is $\subseteq \mathcal{M}$, in a winner-determination problem to establish a repair solution. The WDP is $\mathcal{NP}$-complete [110]. $\Box$

Proposition 1 implies that optimization of expected utility with an exponential number of possible repair problems is computationally intractable.

3.3 Robust Solutions

A *robust solution* to a combinatorial auction should be able to withstand winning-bid withdrawal by not incurring an unacceptable loss in revenue. Such a solution should be able to tolerate such uncertainty by making small changes to the allocation of items to form a repair solution, without causing undue disturbance to the rest of the solution. The revenue of repair solutions for different potential winning-bid withdrawals should also exhibit a low variance because risk averse bid-takers prefer predictable incomes, as described in § 3.2.

In the following section we examine the benefits of robust solutions for risk averse bid-takers. We also explore the use of integer linear programming (ILP) techniques for finding robust solutions. Winner determination for CAs typically uses ILP solvers such as CPLEX [66] for solving such combinatorial optimization problems [33].

3.3.1 Added Value of Robust Solutions

When there is uncertainty about a solution, possible bid withdrawal in our case, then the solution may be regarded as a lottery. Before considering the value of a robust solution we shall determine the value of the lottery, l_o, that contains the optimal solution, S_o, and its optimal repairs, r, as possible outcomes. The certainty equivalent (CE) of l_o is the exact payment that when offered to the bid-taker with zero risk offers the same expected utility as l_o. Figure 3.3 shows an example optimal solution with the revenue of repairs on the abscissa.

The bid distributions presented in the CATS bid simulation tool [78], discussed previously in § 3.1.4, exhibit different bidding patterns. Consequently, the variance in solution revenues is dependent upon the bid distribution. This also affects

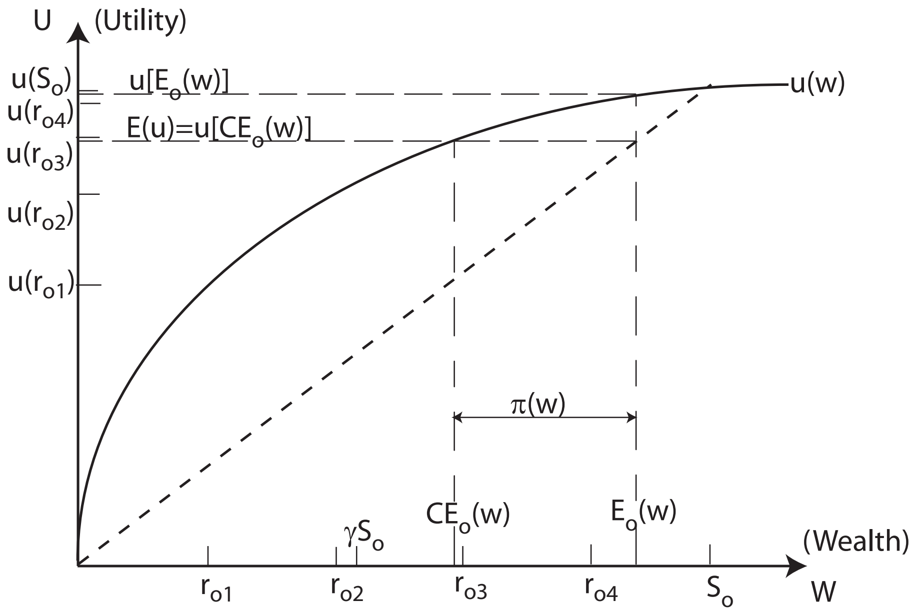

Figure 3.3: Certainty equivalence for the optimal solution.

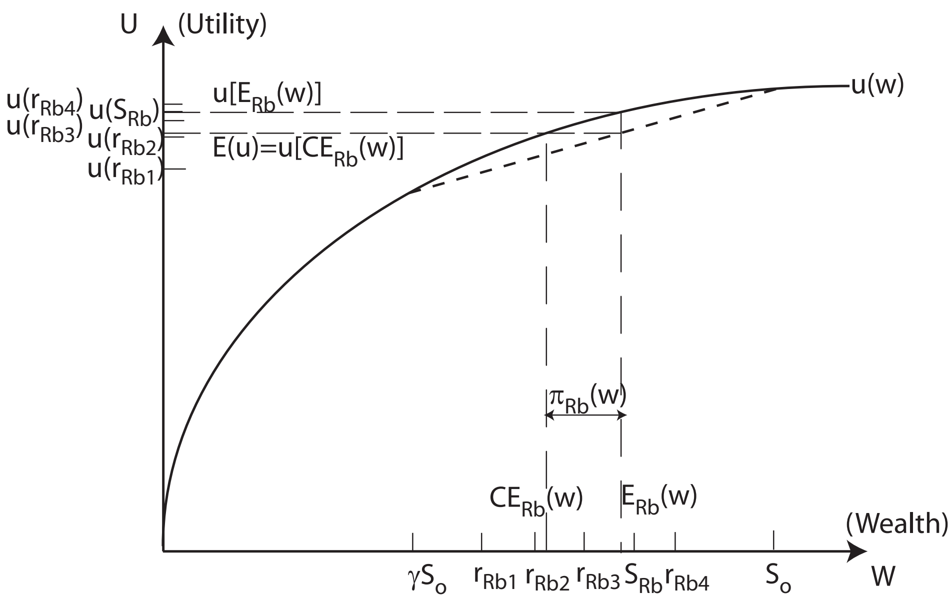

Figure 3.4: Certainty equivalence for the optimal robust solution.

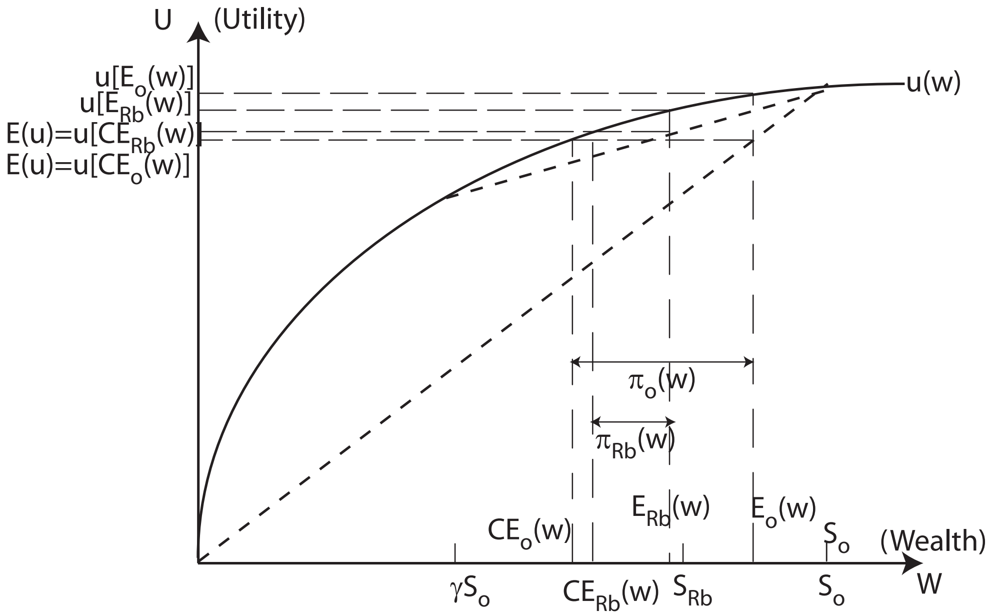

Figure 3.5: Certainty equivalence comparison.

the expected revenue of the optimal *robust* solution that provides a bound on the minimum acceptable revenue for repair solutions.

From the sensitivity analysis, conducted previously in § 3.1.4, of the effects of bid withdrawal, we can estimate the probability distribution function, $f(x)$, of revenues for all possible repair solutions over $[0, S_o]$. Given the limitations imposed by computational feasibility, namely that repairs provide lower bounds only, how much initial revenue should we sacrifice to find a robust solution whose repair revenues are at least γS_o, $\gamma \in (0, 1)$, given a risk averse bid-taker with a concave utility function $u()$?

A necessary condition for the preference of a robust solution over the optimal solution is the following:

$$u[CE_{Rb}(w)] \succ u[CE_o(w)]. \tag{3.4}$$

Figure 3.3 shows an example of a brittle optimal solution and indicates the certainty equivalent, $CE_o(w)$, for this solution given the large variance of revenue for repair solutions $r_{o1} - r_{o4}$. Figure 3.4 shows a robust solution whose repair solutions exhibit lower variance than that of the optimal WDP solution presented previously in Figure 3.3. In this example robust solution all repair solution revenues exceed γS_o. Note that it is possible for a repair solution (*e.g.* r_{Rb4}) to have higher revenue than the pre-withdrawal solution. The CE for the solutions in Figures 3.3 and 3.4 are contrasted in Figure 3.5. We can see that the CE for the robust solution is greater than that for the optimal solution, thus implying that the former is preferable to the latter ($u[CE_{Rb}(w)] \succ u[CE_o(w)]$).

The expected utility of the optimal solution can be determined using the von Neumann-Morgenstern expected utility property when we know the distributions of repair solution, $f_o(x)$, revenues:

$$E_o(w) = \sum_{i=1}^{b_o} \left(p_{o_i} \int_0^{S_o} u(x) f_o(x) dx \right) + (1 - \sum_{i=1}^{b_o} p_{o_i}) u(S_o), \tag{3.5}$$

where b_o is the number of winning bids, p_{o_i} is the probability of the i^{th} winning bid being withdrawn and $f_o(x)$ is the distribution of optimal repair solutions in terms of revenue. Similarly, the expected utility of a robust solution, S_{Rb}, is determined as follows:

$$E_{Rb}(w) = \sum_{i=1}^{b_{Rb}} \left(p_{Rb_i} \int_{\gamma S_o}^{S_o} u(x) f_{Rb}(x) dx \right) + (1 - \sum_{i=1}^{b_{Rb}} p_{Rb_i}) u(S_{Rb}), \tag{3.6}$$

where $f_{Rb}(x)$ is the distribution of repair solution revenues for the optimal robust solution. Also, b_{Rb} is the number of winning bids in the robust solution, and p_{Rb_i} is the probability of the i^{th} winning bid being withdrawn. The probability distribution of repair revenues may be determined empirically using a Monte-Carlo style simulation. The results from the sensitivity analysis provide such information and permit us to estimate expected revenues for user-defined values of γ, but are only indicative in a repeated auction scenario when the items for sale are identical. We also require the likely number of bids in a solution and their respective probabilities of failure. Realistically, it is more likely that γ is chosen as a lower bound on acceptable revenue in the event of a withdrawal because even repeated auctions can be unpredictable. The optimal robust solution is then compared against the overall optimal (discounting possible withdrawals) in terms of expected utility.

The expected utility is the utility of the certainty equivalent. The *risk premium* is defined as the difference between the expected revenue and the certainty equivalent, from Equation 2.2. It may be impossible to achieve a solution without any risk, but a robust solution, S_{Rb}, offers a solution with reduced risk. A robust solution can, therefore, offer added value to a bid-taker.

Definition 3.3.1 (Actuarially Fair Premium (AFP) payment for a robust solution, S_{Rb}.)**.** *The difference in the certainty equivalent of the optimal solution and the certainty equivalent of* S_{Rb}*.*

Therefore, the AFP in Figure 3.5 is $\pi_{Rb}(w) - \pi_o(w)$. The AFP defines the precise *value of a robust solution* over the optimal solution as determined by the WDP. Note that if this premium is negative, the optimal solution is preferable to the robust solution.

3.3.2 ILP Approaches for finding Robust Solutions

Integer linear programs are typically used to formulate the winner determination problem [33], when a specially-tailored algorithm such as CABOB [115] is unavailable. The robustness criterion, however, poses difficulties for ILP formulation.

Bertsimas *et al.* [13, 14] introduce a formulation of integer programming problems that produces robust solutions when both the cost coefficients and the data constraints may be subject to uncertainty. When only the cost coefficients are subject to uncertainty and the problem is a 0-1 discrete optimization problem on n variables; the procedure of solving the robust counterpart by solving $n + 1$ instances of the original problem is described [13, 14]. They show that if the original problem is polynomially solvable, the robust counterpart problem is also polynomially solvable. Therefore, robust versions of such well-known problems as shortest path or bi-partite matching are polynomially solvable when there is uncertainty over cost coefficients or data constraints. We are interested in uncertainty over decision variables rather than cost coefficients or data constraints. In particular, we deal with binary variables that effectively lose 1 as a possible value. In an auction scenario the bid amounts are set unequivocally, whilst there remains uncertainty over the ability or willingness of a bidder to actually pay that amount.

The main difficulty with the ILP approach stems from the fact that repair variables are provided for all submitted bids, even those bids that do not partake in the robust solution. There are more variables initially instantiated than are strictly necessary, because at most m bids may win when there are m items for sale. In real-life combinatorial auctions, $m \ll n$, where n is the number of bids submitted, so ideally only $\mathcal{O}(m+n)$ variables should be present. It may be possible to reformulate the WDP to attain robust solutions more efficiently than using $\mathcal{O}(n^2)$ variables, but earlier work on reformulation to attain robustness is not encouraging. Hebrard *et al.* [59] examined the reformulation of CSPs for finding super solutions and found that it was a very inefficient means of obtaining robust solutions. This does not preclude the existence of a compact model to represent the problem of solution robustness in the WDP. However, our search for an efficient ILP model proved fruitless so we pursued a search-based approach that involved node-level reparability checking. The creation of repair variables dynamically during branch and bound search, as they are required, should improve efficiency by performing only necessary constraint checks. In $\S$ 8.2.2 we discuss possible future work in the area robustness and ILPs.

A technique that is sometimes used in scheduling applications is the introduction of *slackness* to accommodate potential delays by allotting more time than is strictly necessary to activities [31]. This idea could be translated into penalizing bids likely to fail by reducing their bid amount in the calculations accordingly. For instance, in Example 3.2.1 we could execute a WDP where the bid amounts are multiplied by the likelihood of the bid not being withdrawn, thus ignoring its reparability. This is a simplistic and inefficient approach unsuitable for our purposes. It may unnecessarily discriminate against bids that may be easily repaired by others of slightly lesser amounts. For this reason, we do not proceed with this technique.

3.4 Summary

The popularity of first-price auctions remains unexplained in terms of the Revenue Equivalence Theorem. Waehrer *et al.* [128] showed that for the case of a risk averse bid-taker and risk neutral bidders, first-price auctions are favorable to second-price auctions that are, in turn, preferable to English auctions. Risk aversion amongst bid-takers is a significant contributory factor in the choice of auction form.

The possibility of winning-bid withdrawal represents a significant risk for the bid-taker in combinatorial auctions. Unlike the basic auction forms presented in § 2.1.1, an optimal repair solution may be unattainable given the fact that items have been assigned to other bidders. We have called this the *bid-taker's exposure problem*. Reactive approaches to repairing the solution may be impossible, given reasonable constraints on the ability of the bid-taker to involuntarily revoke items from winning bidders. The risk associated with irreparable breaks in a solution must, therefore, be considered in advance of winner determination using a proactive approach. We, therefore, desire robust solutions that can be repaired to form alternative solutions with satisfactory revenue. It is also important that undue disturbance is not visited upon non-reneging winning bidders.

A robust solution for a combinatorial auction that reduces the variance of repair solution revenues was suggested. Unfortunately, finding robust solutions is impractical using conventional ILP formulations that seek to maximize expected revenue. Other approaches including introducing slackness or Bertsimas' robust discrete optimization [13] are also unsuitable. In the following chapter we study a constraint programming framework that provides robust solutions. We shall investigate the feasibility of applying this approach to combinatorial auctions.

Chapter 4

Super Solutions for Combinatorial Auctions

Previously we have seen how the bid-taker's exposure problem presents a risk of significant revenue losses when winning bids may be withdrawn. A pro-active approach is desirable whereby a robust solution is found that is resistant to bid withdrawal. In this Chapter we explore the possibility of using super solutions for constraint programs to attain solution robustness for CAs [58].

We present the super solutions (SSs) framework and show how solution robustness for constraint programs can be achieved [59]. A basic constraint model for CAs is described and we show how a SS can be found using this model. This is the first time that a robust solution has been sought for such auctions. We present an empirical study of the trade-off between solution robustness and other optimization criteria using problems that illustrate the difficulties associated with finding a SS. In particular, we examine the orthogonal notions of finding a SS that maximizes either reparability or revenue. In circumstances where no robust solution exists it may be useful to maximize reparability, whereas revenue can be optimized when more than one SS exists.

There are several different notions of robustness, besides those of SSs [59] and fault tolerance [134] in CP, that can be applied to algorithms or solutions. Branching Constraint Satisfaction [47] is a technique that focuses on finding robust partial solutions to a dynamically changing problem in which new variables may be added after assignments have been made. This may be viewed as a form of online algorithm for a problem that receives new input variables at unspecified times in the future. This is in contrast with super solutions, whose problem variables remain fixed. There are other CP frameworks such as Dynamic CSP [37], Fuzzy CSP [38, 113], Probabilistic CSP [43], Mixed CSP [44], Soft CSP [15] and Stochastic CSP [130] that also deal with uncertainty. However, none of the above frameworks provide robust solutions that bound corrective changes should an assignment be invalidated.

Bent and Van Hentynryck's online stochastic scheduling [10], proposes a consensus-based stochastic approach to dealing with future uncertainty, and give packet scheduling as an example. That work assumes that there is new incoming information and there are time constraints on the time available to determine what packet should be scheduled next. Wallace and Freuder's Stable Solutions for Dynamic CSPs [129] examined the possibility of losing domain values and having constraints added during search. Many temporary changes are anticipated and the intention is to minimize conflicts using a hill-climbing approach that backtracks during search when a change occurs. This is in contrast to super solutions, in which small breaks are anticipated and guarantees are provided on reparability. Yorke-Smith and Gervet [142] adopted another approach to handling uncertainty, in which they focus on uncertainty due to incomplete or erroneous data. They seek a reliable model, in which there is no uncertainty because all approximations are removed. The "full closure", or set of all possible solutions, is then computed. There is no concept of solution robustness following possible changes in this work because the uncertainty is removed in the model to ensure closure over guaranteed solutions.

Weighted CSP associates a cost with tuples in constraints and is a non-idempotent, soft-constraint framework. The optimization problem is typically to minimize the total cost of the solution. A MAC-based search [76] can be used to find solutions for WCSPs. The cost function is defined as the summation of the costs of all constraints. However, none of the above techniques are concerned with preempting uncertainty so that the solution is easily reparable following a failure, as is the case with super solutions [59].

Finally, in this Chapter, we critique the SSs framework and discuss some of

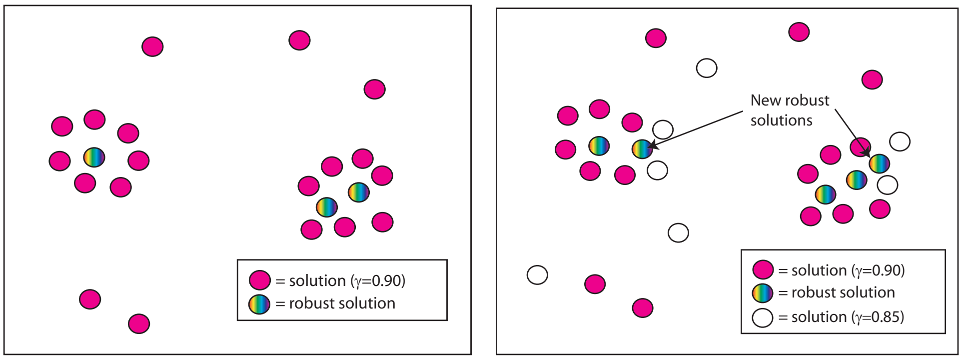

(a) Solutions; $\gamma = 0.90$ (b) Solutions; $\gamma = 0.85$

Figure 4.1: Robust solutions intuition.

its limitations. We propose novel extensions to circumvent these difficulties and increase its flexibility.

4.1 Motivation and Overview

Finding a solution that optimizes expected utility is computationally infeasible, whereas a simple heuristic approach such as discounting bid amounts for unreliable bidders is computationally feasible but unsatisfactory in terms of solution quality. There is a tradeoff between expected utility optimization and computational feasibility.

We propose a parameterization of robustness in combinatorial auctions that will facilitate a search mechanism for finding solutions that exhibit reduced variance in expected utility. Adding a *minimum revenue constraint* so that revenue must exceed a certain fraction, $\gamma \in [0, 1]$, of the optimal revenue S_o reduces the set of possible solutions. Figure 4.1(a) shows an intuitive example where $\gamma = 0.90$. The relative proximity of solutions in this diagram indicates the ease with which one solution can be changed to another should a failure occur. If the revenue constraint is further relaxed to $\gamma = 0.85$, as in Figure 4.1(b), the cardinality of the solution set increases. A robust solution is one from which it is easy to transition to another (repair) solution should any possible break cause it to be invalidated. As the number of available solutions increases in Figure 4.1(b), so too does the number of robust solutions from which to choose.

The cost of repairing a solution is another important parameter in solution robustness. For combinatorial auctions, some items are already awarded to other non-reneging bidders and revocation of these items is undesirable. The metric for establishing the cost of repairing a solution is equally important when finding robust solutions. Super solutions measure repair costs according to the cardinality of the repair set, *i.e.* how many variable assignments have changed.

4.2 Super Solutions for Constraint Programs

In this section we present the SSs framework for constraint programs [52, 58, 59] and show how it can provide solution robustness for combinatorial auctions. In constraint programming (CP) [36] , a constraint satisfaction problem (CSP) is modeled as a set of n variables $X = \{x_1, \ldots, x_n\}$, a set of domains $D = \{D(x_1), \ldots, D(x_n)\}$, where $D(x_i)$ is the set of finite possible values for variable x_i and a set $C = \{C_1, \ldots, C_m\}$ of constraints, each restricting the assignments of some subset of the variables in X.

Constraint satisfaction involves finding values for each of the problem variables such that all constraints are satisfied. Its main advantages are its declarative nature and flexibility in tackling problems with arbitrary side constraints. Constraint optimization seeks to find a solution to a CSP that optimizes some objective function. A common technique for solving constraint optimization problems is to use branch-and-bound techniques that avoid exploring sub-trees that are known not to contain a better solution than the best found so far. An initial bound can be determined by finding a solution that satisfies all constraints in C or by using some heuristic methods.

A classical *super solution* (SS) is a solution to a CSP in which, if a limited number of variables lose their values, repair solutions are guaranteed with a bounded number of changes, thus providing solution robustness [59]. It is a generalization of both fault tolerance in CP [134] and supermodels in propositional satisfiability (SAT) [52]. In the classical SSs framework, an (a, b)-super solution guarantees that if at most a variables lose their assignments, a repair solution can be found by reassigning those a variables and at most b others. Therefore, a SS is a proactive approach for dealing with uncertainty that ensures solution robustness.

Super solutions for constraint programs place an upper bound on the number of assignments that need to be changed when forming a repair solution. Only a particular set of variables in the solution may be subject to change and these are said to be members of the *break-set*. Hebrard *et al.* [59] have also described possible means of restricting break and repair sets. For each combination of brittle assignments in the break-set, a *repair-set* is required that comprises the set of variables whose values must change to provide another solution. The *cardinality* of the repair set is used to measure the cost of repair. Example 4.2.1 gives a simple example of how a SS provides solution robustness for a small CSP.

Example 4.2.1. *Consider a CSP with two variables* X *and* $Y \in \{0, 1\}$*. There are three possible solutions* $\langle 1, 1 \rangle$*,* $\langle 0, 1 \rangle$ *and* $\langle 1, 0 \rangle$*.* $\langle 1, 1 \rangle$ *is a (1,0)-super solution because if either of the assignments breaks we are still left with a solution after making 0 changes to other variables. However, solutions* $\langle 1, 0 \rangle$ *and* $\langle 0, 1 \rangle$ *are (1,1)-super solutions because we can always repair one variable if necessary to form another solution. A (1,0)-super solution may, therefore, be considered more robust than a (1,1)-super solution because it causes less disruption when forming a repair solution.* $\triangle$

4.2.1 (1,b)-Super Solution Search Algorithm

Super solutions are found by using a MAC-based [114] search algorithm† that examines the reparability of each variable assignment at each node of the search tree. MAC is a state-of-the-art hybrid search and inference algorithm used to solve constraint satisfaction problems. Partial solutions are extended after all previous brittle assignments are deemed reparable. Algorithm 1, from [58], shows how the MAC-based search algorithm may be implemented for (1,b)-super solutions.

A repair solution, R_x, is associated with each variable x. This repair solution consists of an array of assignments for all variables should variable x's assignment break. The `backtrack` procedure attempts to extend the partial assignments of $\mathcal{X}$. The `reparable` procedure checks that each previous variable assignment, in $Past$, can still be repaired after another variable in $\mathcal{X}$ has been assigned. This

†It is possible to reconfigure the algorithm to enforce different levels of consistency at nodes of the tree. For example, forward-checking could be used instead of MAC.

involves extending the partial repair solutions in R_x for all $x \in Past$. If at any stage, a partial repair solution cannot be extended without satisfying all constraint then backtracking occurs.

The search for a repair solution, R_x, ensures that the repair value for x is different to that assigned in the SS because R_x is the alternative solution should the assignment for x break.

Backtracking may also occur if the current master search for a (1,b)-super solution may not be able to extend the partial solution in a manner that is consistent with the problem constraints. Hebrard *et al.* [58] have shown that this algorithm terminates, is sound and complete. They have also showed that it is $\mathcal{NP}$-complete when a is fixed.

Algorithm 1: SUPER-SOLVE [58]

```
input  : b, CSP: P={X, D, C}
output: S: a (1,b)-super solution and R:the set of repair solutions
begin
    S ← ∅  // Solution
    R ← ∅ // Set of repair solutions
    Past ← ∅ // Ordered set of assigned variables
    foreach x ∈ X, y ∈ X do
        R_x[y] ← min(D(y))
    AC(P, S) // Perform arc-consistency
    backtrack(P,S,Past,R,b,0)
end
```

It is also possible to impose restrictions on the break and repair sets. For example, some assignments may not be brittle so it may not be necessary to find repair solutions for all variables. Repair set restrictions may impose a constraint on which variables are allowed to change. Also, alternative value restrictions may preclude some domain values from acting as repairs for others.

Although the search for reparable solutions is described as a procedure in [58], it can in fact be implemented as a `reparable` global constraint using an appropriate filtering algorithm that counts the difference in assignments between the SS and the relevant repair. For further details on the algorithm for (1,b)-super solutions, the reader is referred to [58].

Procedure `backtrack(`P,S,$Past$,R,b,lvl`)`:Boolean [58]

```
begin
    if 𝒳 = Past then return true
    choose x ∈ 𝒳 \ Past
    Past[lvl] ← x
    foreach v ∈ 𝒟(x) do
        save 𝒟, R
        S ← S ∪ {(x, v)}
        if AC(P,S) then
            foreach y ∈ Past & brittle(y) do
                if ¬reparable(P,S,Past,R_x,0,b) then break
                if backtrack(P,S,Past,R,b,lvl + 1) then return true
        restore 𝒟, R
        S ← S \ (x, v)
    Past[lvl] ← ∅
    return false
end
```

It is possible to reformulate a CSP so that the only solutions are SSs [52, 59, 134]. A fault tolerant solution is the same as (1,0)-SS. Weigel and Bliek [134] presented a simple reformulation that duplicates variables, domains and constraints. An inequality constraint is also added between all the original variables and their duplicates. The solution to the original variables (if one exists) is then a (1,0)-SS. Hebrard *et al.* proposed a reformulation using the cross-product of domains [59]. However, we focus on a search-based approach because [59] provides evidence of its superior performance over reformulation.

4.2.2 A Pedagogical Example

The search for a SS runs concurrently with searches for repair solutions whereby the search for a SS in the master tree backtracks if any of the repair searches are unable to reach the same node level. Figure 4.2 provides an example of how the search algorithm traverses a master tree while looking for a SS for the auction presented in Table 4.1, The master tree search only progresses to a child node if all previous brittle assignments can have their partial repair solutions extended to the same level of the search tree.

Table 4.1: Combinatorial auction with a (1,0)-SS.

	Items		
Bids	A	B	C
$BidX$	100	0	0
$BidY$	0	100	0
$BidZ$	0	0	100

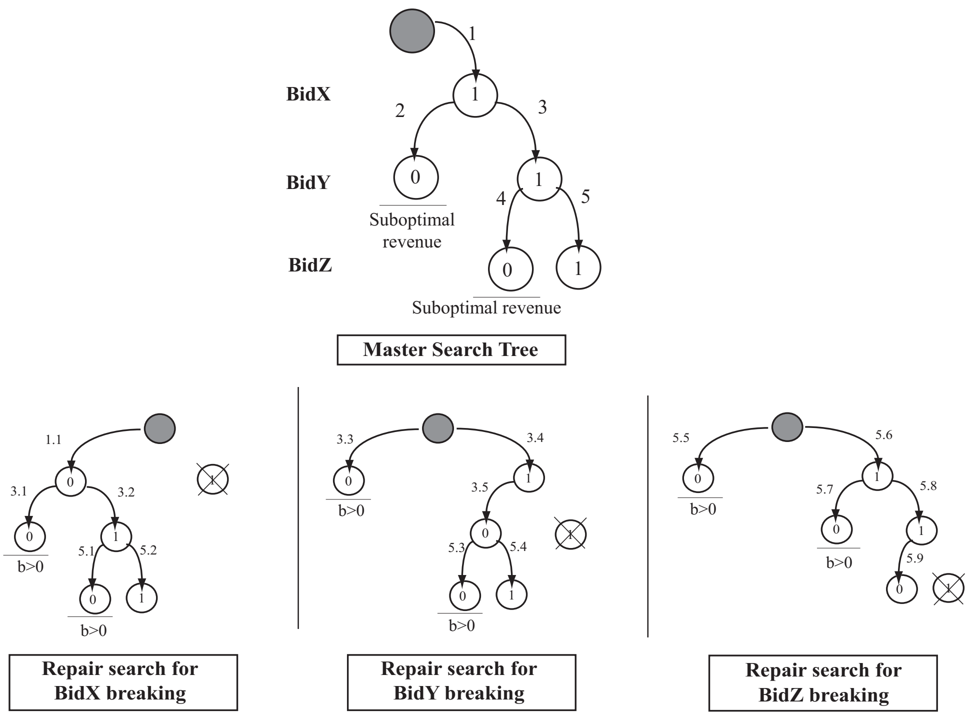

Figure 4.2: Flow of control between master search tree and repair searches.

In this simple example we shall assume that the revenue constraint stipulates that the minimum tolerable revenue for any solution is 100. There are three bids and all bids are for non-overlapping sets of items. Therefore, the set of 7 possible solutions are $\langle 0,0,1\rangle$, $\langle 0,1,0\rangle$, $\langle 1,0,0\rangle$, $\langle 0,1,1\rangle$, $\langle 1,0,1\rangle$, $\langle 1,1,0\rangle$ and $\langle 1,1,1\rangle$. It is easy to see that any solution with two or three winning bids is a (1,0)-SS because a winning-bid withdrawal is reparable without making any other changes to the solution. However, we are seeking a robust solution with the maximum possible revenue.

We follow the search for an optimal (1,0)-SS in the master search tree when $BidX$ is set to 1 (assignment 1). Previously a (1,0)-SS was found with revenue 200 ($\langle 0,1,1\rangle$), so we have this as a lower bound on revenue for an optimal SS. The annotated arrows in Figure 4.2 indicate the chronology of assignments. Following a valid assignment in the master search tree, control switches to the search for partial repair solutions for instantiated variables. The repair search for $BidX$ breaking simulates this bid being withdrawn by assigning its repair value to 0 (assignment 1.1). This does not violate any constraints, so search resumes in the master search tree. When $BidY$ is set to 0, (assignment 2), the current lower revenue bound of 200 cannot be improved upon in the subtree so it is instead assigned to a value of 1 (assignment 3)[†]. The repairs for $BidX$ and $BidY$ are then successfully extended to this level of the search tree (assignments 3.1-3.5). The search continues successfully until $\langle 1,1,1\rangle$ is found, which constitutes an optimal (1,0)-SS with revenue of 300. It is easy to see that any break in the optimal solution is reparable.

4.3 The Need for Super Solutions in Combinatorial Auctions

Chapter 3 presented the bid-taker's exposure problem that is encountered when winning bids are withdrawn in CAs. We then saw how super solutions can offer robustness by restricting the number of required changes to form a repair solution

[†]We shall explain in later chapters how lower bound on revenue in subtrees is determined efficiently.

following a winning-bid withdrawal. The purpose of finding super solutions is that if the solution is perturbed, another solution that satisfies a minimum revenue constraint may be found by changing a limited number of other variables.

We can represent bids as binary variables in our CP model. Winning bids may be withdrawn so that all bid variables are included in the break set. However, only those variables whose value is 1 can be withdrawn, so 0 is considered a robust value that does not fail and, therefore, does not require a repair solution.

A (1,b)-SS allows us to find a solution to a CA such that if any winning bid is withdrawn, the solution can be repaired by making at most b other changes to the allocation. We can limit the maximum number of involuntary withdrawals of items from winning bidders by the bid-taker to b.

However, finding SSs can be computationally expensive [58] so a pure CP approach to the WDP that has an exponential search space does not scale very well. A hybrid approach incorporating OR techniques is required for larger auctions. Section 4.4 concentrates upon a fixed size of problem (20 items and 100 bids) and shows how (1,b)-super solutions are achievable for tighter revenue constraints in auctions with fewer items in each bid. Later, we shall explore the use of operations research (OR) techniques to aid scalability, but for now we focus upon the trade-off between robustness and revenue.

4.4 Empirical Analysis of Winning-bid Withdrawal

We performed an empirical study of the inherent trade-offs involved in finding robust solutions for CAs using the SSs framework. In particular, we examined the following scenarios:

1. Constraint Satisfaction: finding *any* (1,b)-super solutions;
2. Constraint Optimization: optimizing revenue when there are many SSs from which to choose;
3. Constraint Optimization: optimizing reparability when the problem is unsatisfiable. This approach seeks to maximize the number of possible failures that are reparable when it is not possible to guarantee that all assignments can be repaired.

We used the Combinatorial Auction Test Suite (CATS) [78] to generate sample auction problems in which there are 20 items for sale and 100 non-dominated bids† that are represented using 100 binary CSP variables. Any bids that share an item are *overlapping* and a constraint between those variables precludes the success of both bids.

There are a number of ways to accelerate the search for SSs in combinatorial auctions but we shall initially use a basic model to examine the trade-offs between robustness and other optimization criteria. For example, polynomial matching algorithms [22, 75, 99] may be used in auctions whose bids contain very few items, such as those for airport landing/take-off slots. Another technique is to use a linear relaxation of the Set Packing Problem (SPP) to give an upper bound on potential revenue in sub-branches of the search tree. Such additional techniques, some of which are outlined in [33, 54, 115], can aid the scalability of a CP approach. Later, in Chapter 6, we explore some OR techniques for reducing search effort.

We performed our experiments using the EFC constraint solver [9]. Twenty instances of each problem were used to generate average results. We used various bid distribution types that simulate different economically motivated auction scenarios. In the `arbitrary` distribution, complementarities may not be as universal as geographical adjacency, with valuations being normally distributed again. Bidders view the complementarity of items slightly differently. Recall from § 2.3.2 that the `regions` distribution is modeled on a scenario in which items are location dependent and complementarity is a function of the proximity of these items in 2-dimensional space, such as in a spectrum or real-estate auctions, and valuations are distributed normally. The `scheduling` distribution simulates an auction for time slices on a resource in a distributed job-shop scheduling problem. The bids for this distribution type tend to contain fewer items, therefore, there are more possible combinations of successful bids. This increases the difficulty of finding the optimal winner in the WDP but also increases the likelihood of being able to find a robust solution because of the increased availability of repair solutions.

---

†The CATS flags included int_prices with the bid_alpha parameter set to 1000.

4.4.1 Constraint Satisfaction

In these experiments we first solved the WDP optimally using our EFC-based constraint solver. This could also have been done using an ILP solver such as ILOG's CPLEX. We then stipulate a minimum percentage of optimum revenue that is acceptable and the maximum number of variables that can change, b, when repairing a SS. We then use the constraint-based solver to search for a satisfactory super solution. We present the success rate, *i.e.* the fraction of auctions that have a (1,b)-SS, in Figure 4.3.

The contours on the horizontal plane of the graphs indicate the gradient of the surface in the graph. This helps illustrate the rates of fall-off in the running time and the success rates for the different distributions. It is evident from the contours in Figure 4.3(c) that the `scheduling` distribution reaches a very high success rate with $b = 1$ and acceptable revenue of at least 90%, see Figure 4.3(c). The `arbitrary` (Figure 4.3(a)) and `regions` (Figure 4.3(b)) distributions have lower success rates, thus indicating that it is more difficult to find robust solutions for these distributions.

The average running times to prove the existence/non-existence of a (1,b)-super solution are give in Figure 4.4. The contours also help to illustrate where the peak running times are encountered for the various distributions. It is evidently easier to show satisfiability/unsatisfiability of (1,b)-super solutions for `arbitrary` and `regions` auction instances than for `scheduling`. When the reparability and revenue constraints are tight, (b is low and minimal revenue is high), it is easy to detect the non-existence of a SS so running times are reduced. When constraints are very loose it is also easy to find a SS. However, there is a transition area where the existence of a SS is unpredictable and the running times are high. This peak in running time is most clearly visible in Figure 4.4(c). The running times are not tightly correlated with b. This is because an increase in b can serve to make problems trivially satisfiable or alternatively it can make it more difficult to prove unsatisfiability. For example, in Figure 4.4(b) we can see that when the revenue constraint is at its most relaxed (80% of optimum) increases in b lead to reduced running times because it is easy to find SSs. However, when the revenue constraint is tightened, to say 90% of optimum, there are fewer satis-

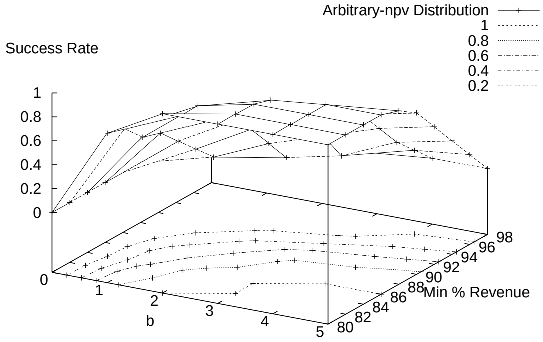

(a) `arbitrary` distribution.

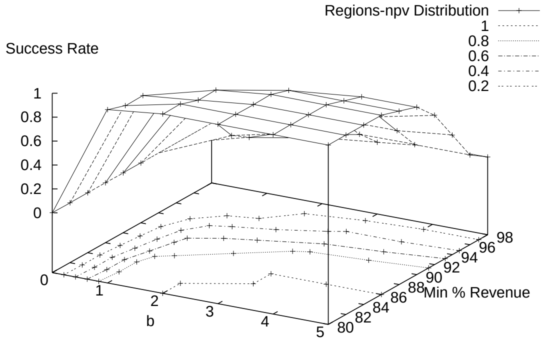

(b) `regions` distribution.

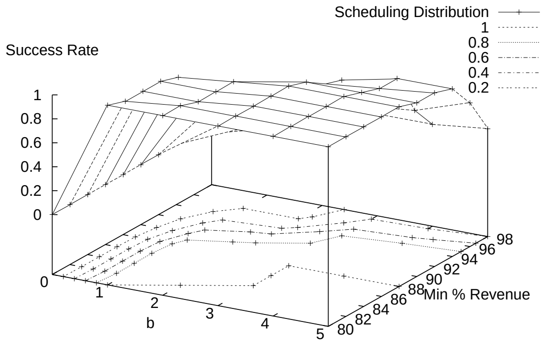

(c) `scheduling` distribution.

Figure 4.3: Success rate.

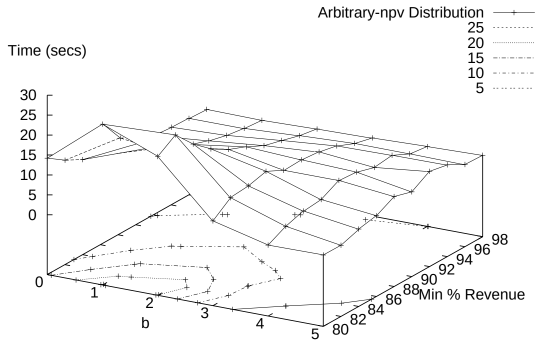

(a) `arbitrary` distribution.

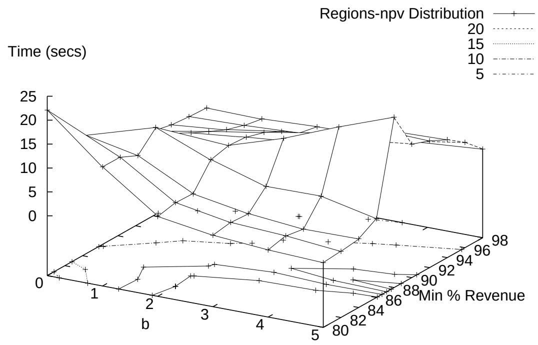

(b) `regions` distribution.

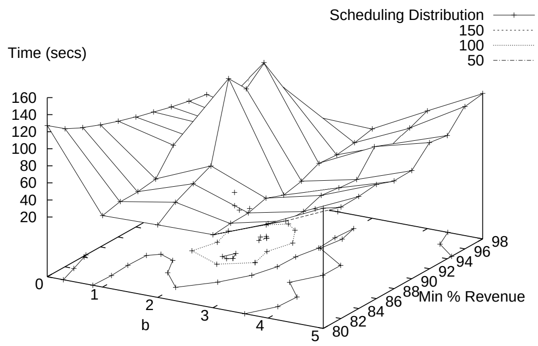

(c) `scheduling` distribution.

Figure 4.4: Running times.

fiable instances (see Figure 4.3(b)) so there is more search time required in these problems because the repair search visits more nodes when b is increased.

We can also estimate from Figures 4.3 and 4.4 that the hardest satisfaction problems for the various distributions occur when the success rate is approximately 75%.

4.4.2 Constraint Optimization

In reality, a bid-taker who seeks a robust solution desires the optimal robust solution in terms of either robustness or revenue. We may then employ a search algorithm to find the optimal SS with the appropriate objective function. We use a branch and bound algorithm that finds SSs and optimizes on either reparability, in the case of an over-constrained problem, or revenue, when there are many SSs that satisfy the given constraints. This could also be used as an anytime algorithm that finds the best possible robust solution in a given time-frame. Our analysis focuses on two forms of optimization:

1. Optimizing Revenue: Intended for use in an under-constrained scenario in which there are many super solutions. We seek the super-solution with maximal revenue.
2. Optimizing Reparability: Intended for use in an over-constrained scenario in which there are no super-solutions. We seek a solution that contains the minimum number of irreparable variables. This is the CA allocation that satisfies the revenue constraint and minimizes the number of possible bid withdrawals that cannot be repaired to satisfy this revenue constraint.

Optimizing Revenue

If there are many (1,b)-super solutions satisfying the revenue constraints then it is desirable to find a revenue-maximizing SS. This is a more difficult problem than simply finding *any* SS given a constraint on revenue, as performed in § 4.4.1.

We have developed a branch and bound algorithm that returns the optimal super-solution in terms of an objective function. In this case we searched for a SS whose revenue is maximized whilst the constraints on the revenue for repair solutions remains unchanged. An alternative approach may be to maximize the minimal revenue on the SS and all repair solutions. We do not explore this avenue because we conjecture that the probability of non-withdrawal is higher than that of a single withdrawal. Hence, optimization of revenue of the SS is a priority over that of optimization of repair revenues.

Figures 4.5(a), 4.5(b) and 4.5(c) show clearly that when we permit more variable changes, the expected increase in revenue of the optimal SS increases significantly. Notice how optimization is far more computationally expensive than constraint satisfaction. This can clearly be seen by comparing the running times in Figure 4.6 with those in Figure 4.4.

We have restricted our analysis of the `scheduling` distribution to revenue greater than 90% because the high solution-density, low-revenue constraints cause the optimization in this uncomplicated constraint model to become very computationally expensive. As the constraint on acceptable revenue for repair solutions is tightened by raising the minimum revenue for repair solutions, in some cases this leads to a SS of reduced optimal revenue.

Optimizing Reparability

A scenario may occur where there exists no (1,b)-super solution that satisfies the minimal revenue criterion. Instead, we may seek a solution that minimizes the number of irreparable variables in a SS, thus compromising on b whilst satisfying the revenue constraint. Hebrard *et al.* [58] developed a (1,b)- SUPER-BRANCH-AND-BOUND algorithm that will find a solution with a minimal number of irreparable variables if no SS exists for the given value of b.

Previously, in § 4.4.1, we learned that finding robust solutions for the `arbitrary` and `regions` bid distributions was less likely. In this section we are examining those auctions where there is no robust solution. Figure 4.7 shows how many variables in the super-solution cannot provide a repair solution when at most b assignments can be changed. This illustrates the severity of irreparability for the various distributions. The `arbitrary` distribution is the least robust distribution (Figure 4.7(a)), closely followed by `regions` (Figure 4.7(b)). For example, when the revenue constraint on an auction with an `arbitrary` distribution type

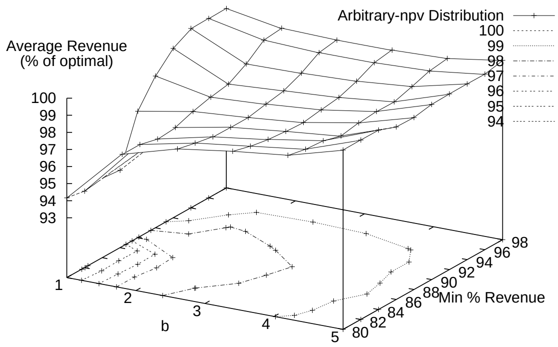

(a) `arbitrary` distribution.

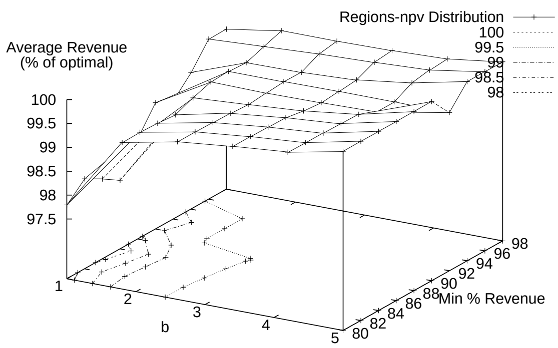

(b) `regions` distribution.

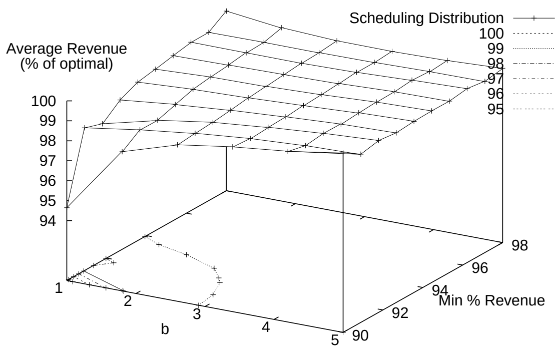

(c) `scheduling` distribution.

Figure 4.5: Average optimal revenue of satisfiable instances.

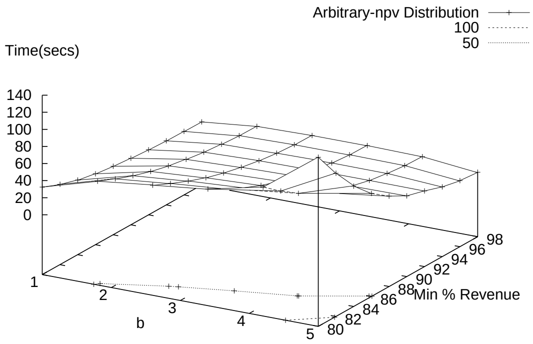

(a) `arbitrary` distribution.

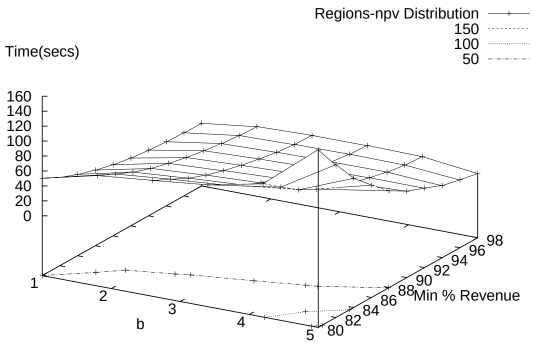

(b) `regions` distribution.

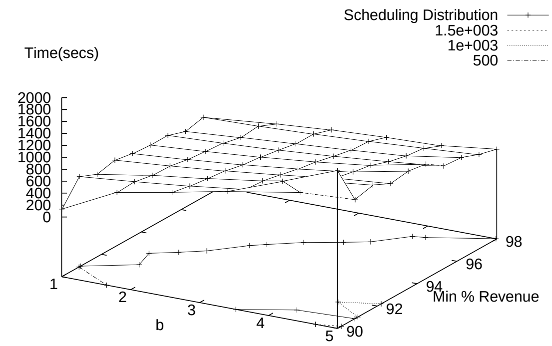

(c) `scheduling` distribution.

Figure 4.6: Average time to find the SS of maximal revenue.

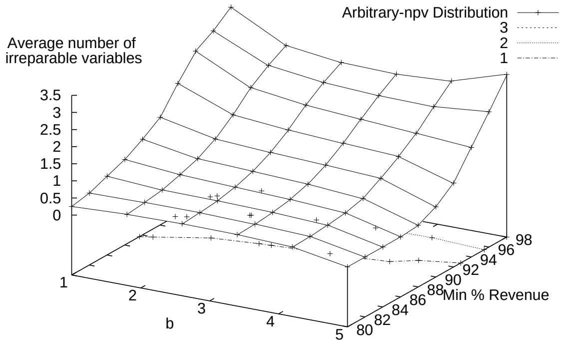

(a) `arbitrary` distribution.

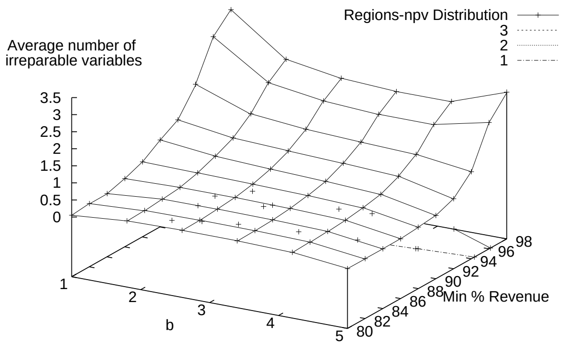

(b) `regions` distribution.

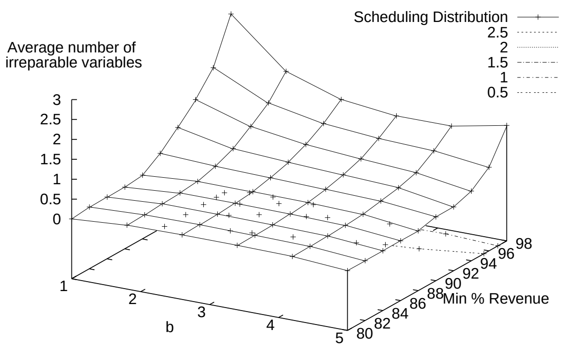

(c) `scheduling` distribution.

Figure 4.7: Maximizing reparability for over-constrained problems.

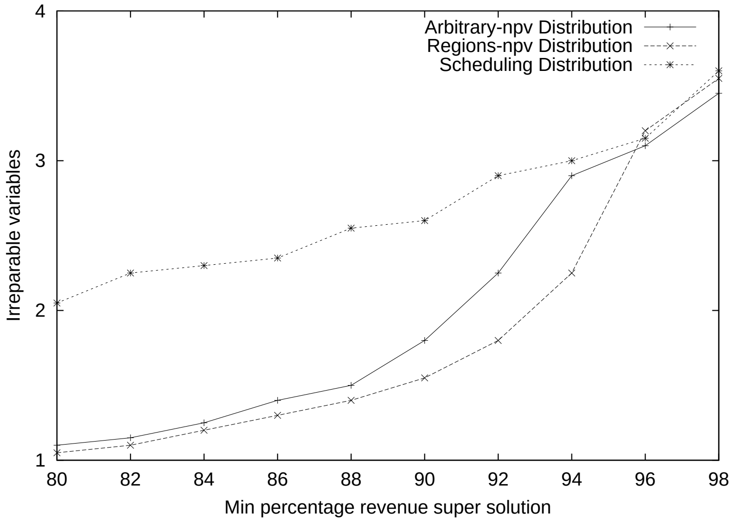

Figure 4.8: Maximizing reparability for over-constrained problems (extreme case: *b*=0).

is tight ($\geq 95\%$) and few changes can be made to the solution ($b \leq 2$) then there are typically 2-3 bids in the auction that will not provide a repair solution given those constraints on revenue and b. The `regions` and `scheduling` auctions have shorter bids and denser solution spaces so it is easier to find super-solutions that support such tight constraints.

Figure 4.8 shows the average number of irreparable variables in the extreme case of $b = 0$, for the three auction distributions. Recall that a (1,0)-super solution remains a valid solution if any single winning bid is withdrawn and no new winning bids are declared. This is only possible when there are many winning bids whose value is less than that of the tolerable loss in potential revenue. This is a very tight constraint that is in fact unsatisfied by any of the sample auctions in our test-set, (see Figure 4.3). However, we can attempt to find a solution that minimizes the number of winning bids that do not satisfy this constraint, or *irreparable* variables. Figure 4.8 shows the increase in the minimum number of such variables as the constraint on minimum revenue is tightened.

There are two principal factors that determine the reparability of an auction allocation, the number of winning bids and the number of possible repair solutions. The `arbitrary` distribution has fewer winning bids and fewer solutions so its reparability degrades rapidly. The `regions` and `scheduling` auctions tend to have more available repair solutions, therefore, degrade more slowly as the revenue increases. Also, there is a denser solution space for these distributions that also contributes towards a slower increase in the number of irreparable variables as the revenue constraint is tightened.

4.5 Required Extensions for Super Solutions

Whilst the SSs framework provides a valuable means for finding robust solutions, it is somewhat inflexible in some respects that are important to real-life applications. In particular, there are some limitations in the approach when applied to CAs. There is an underlying assumption that when a repair solution is created, the incurred cost of changing each variable's value in the repair set is the same and the total cost of repair is the cardinality of the repair set. In a real-world scenario, informing a losing bidder that they have now won because of the withdrawal/disqualification of a winning bid would typically incur a lower cost than informing a winning bidder that they have now lost. The auctioneer may have to break a contract or pay a penalty for such an action. This can be seen as a disadvantage of the SSs framework and militates against its deployment. Computing the cost associated with changing the losing/winning status of any bid is in reality a more complex issue that may depend on several other factors. Determining the legality of a repair solution by measuring the cardinality of the repair set may be overly restrictive in many application domains, besides auctions. Example 4.5.1 illustrates how SSs are an imperfect approach towards establishing robust CA solutions.

Example 4.5.1. *Consider an auction whose bids comprise those listed in Table 3.1 on page 44, with each bid considered to be brittle. Also, let the minimum repair revenue be 180, so* $\gamma = \frac{180}{200} = 0.9$.

We see that $\langle 1, 1, 0 \rangle$ *is a (1,2)-SS because a single winning-bid withdrawal can be repaired by modifying the assignment of the two other bid variables. However,* $\langle 0, 0, 1 \rangle$ *also qualifies as a (1,2)-SS, because if the winning bid fails, the other two losing bids can be reassigned to winning to form a repair solution. However, this* *solution can be repaired to form a solution with revenue of 200 without revoking any items from winning bidders. It is, therefore, easier to repair this assignment because changing a variable from 0 to 1 is less costly than the reverse.* △

In other application domains such as scheduling, the cost associated with changing the value of a variable in a solution may depend on its destination value. Consider a factory scheduling problem where variables represent machines and values correspond to states. The cost of changing the state of any machine depends on both the source and destination states.

The cost of changing a variable may also depend on the variable(s) that caused the break. For example, if a particular agent withdraws a bid from an auction, the auctioneer may favor rejection of the agents' other successful bids rather than disturbing an innocent third party. The cost of a repair solution should, therefore, depend on the value assigned initially, the destination value and the break. We argue that changing some variables in a repair solution incurs less cost than others, thereby motivating the introduction of a different metric for determining the legality of repair sets. We justify this approach by using robust solutions to combinatorial auctions as an example application domain but other domains may reap benefits from these extensions.

Hebrard *et al.* [58] also described how some variables may fail (such as machines in a job-shop problem) and others may not. If we generalize this approach so that there is a probability of failure associated with each variable value, we can then alter the criteria for repair solutions according to the likelihood of potential breaks. This extension to SSs is motivated by the maintenance of robust solutions for repeated combinatorial auctions. For example, a business may hold auctions on a regular basis with the same or a similar set of bidders that may act as suppliers or customers. These trading partners may have different levels of trustworthiness and an expressive means of describing sets of assignments that may fail according to a probabilistic scale offers the possibility of providing repair solutions for sets of assignments of arbitrary cardinality.

4.6 Summary

In this chapter we applied the super solutions framework to the problem of finding robust solutions for combinatorial auctions. Super solutions offer a promising platform for finding robust solutions for such auctions that can alleviate the bid-taker's exposure problem. We examined the trade-off between solution reparability and the minimum revenue constraint and also against b, the reparability. Our experimental analysis highlighted a number of interesting points that we shall use to focus our development of an improved framework for solution robustness in subsequent chapters:

1. Robust solutions for the `arbitrary` and `regions` bid distributions are less likely to exist than for `scheduling`;
2. Running times peak when there is approximately a 75% chance of finding a robust solution for any distribution;
3. An important positive result is that optimum robust solutions, can attain high revenue;
4. The search for super solutions is computationally expensive. We shall need to augment the CP model with OR techniques to improve scalability;
5. The minimum number of irreparable variables, when optimizing robustness, increases significantly for all distributions when the revenue constraint is tightened to ≥ 90% of optimum.

We also presented some limitations of the approach and proposed some valuable extensions to the framework that may improve its flexibility and accuracy. We shall proceed to extend the framework to accommodate the requirements of combinatorial auctions.

Chapter 5

Weighted Super Solutions for Constraint Programs

In this Chapter we present weighted super solutions which extend the basic framework in two important ways that were motivated by problems encountered in Chapter 4. The objective is to develop a more expressive framework that can reflect the risks involved in combinatorial auctions and measure the costs of solution reparability more accurately.

In the new weighted super solutions (WSSs) framework [64] the set of assignments that may fail is determined using a probabilistic approach enabling us to find repair solutions for assignments that are most likely to fail. The probabilities of variables losing their values can be described using either static probabilities or Weibull probability distribution functions over time. We include a novel metric for reasoning about the cost of repair. These two extensions provide a very expressive framework for finding robust solutions that describe potential failures and repair costs in a flexible manner.

This is a versatile framework that can be applied to many application domains. We have designed the framework in a generic manner that accommodates failure and repair characteristics beyond what is strictly necessary for combinatorial auctions. To highlight this versatility, we present a pedagogical example using the N-queens problem where squares on a chessboard can become unavailable after the initial allocation. We also briefly demonstrate how the WSSs framework may be used for job-shop scheduling problems that seek to minimize makespan. However, as with super solutions, this framework is designed so that small breaks can be repaired with small changes. If the probabilities of failure are high and many changes are anticipated, it may be preferable to adopt an alternative approach such as one based on Markov Decision Processes.

5.1 Introduction

Solutions are *robust* if a repair solution is available should some assignments become invalid (break). Previously in Chapter 4 we saw how a *super solution* to a constraint program guarantees that if the solution breaks, another solution can be found by changing a limited number of other assignments [52, 59].

The weighted super solutions (WSSs) framework allows us to capture failure characteristics, such as those that underly metal fatigue, corrosion and abrasion, which exhibit various probability distribution functions. Accurate failure prediction facilitates contingency planning where it is needed so that repair solutions are easily achieved should an assignment fail. This can be described in two different ways: using *static* probabilities of failure or *dynamic* failure rates. The latter describes the probability of failure over time, such as in factory scheduling problems, where machines or components have probabilistic failure rates.

It is important that repair solutions are available for assignments that are likely to break within a given time-frame. We incorporate a powerful approach to modeling failures, based on the *Weibull distribution* [95] from the field of reliability engineering, to describe the failure rates of assignments in solutions to constraint programs. Using this distribution we can represent many of the most common failure distributions such as normal, lognormal and exponential. Having access to such generality is important for reasoning about robustness in a realistic manner.

It is equally important that we can model the cost of repair. We introduce a novel metric for repair costs. Once a break in a solution occurs, *e.g.* a machine in a factory breaks down, the WSSs framework has a more accurate means of comparing the costs associated with alternative ways of repairing the solution. Changing the values of certain variables may incur a heavier cost than others, *e.g.*

changing the production schedule on a large production line may be easier than on a smaller one.

This chapter will introduce the range of probabilistic failures that can be represented by the WSSs framework. The framework's versatility is highlighted through simple examples that are unrelated to auctions but serve to illustrate its usefulness nonetheless. It is left to the Chapter 6 to provide an in-depth study of robustness for CAs using this framework.

5.2 Probabilistic Failure

Breaks in solutions to constraint programs may or may not be time-dependent. For example, a solution to a CA results in a break that is effectively immediate when a bidder refuses to pay. Therefore, a constant probability of failure can be associated with each variable assignment. A factory scheduling problem, however, may exhibit failure rates over a period of time. We discuss each case in detail below.

Constant probabilities of failure (Static WSS). Hebrard *et al.* [58] described how some values in a solution are not subject to change, referred to as *robust* values. In a WSS we propose that assignments have varying degrees of robustness, in order to differentiate between those that are more or less likely to cause failure. In this manner, repair solutions can be determined for sets of assignments that are likely to fail, whilst more robust assignments may not require any repair solution if their probability of failure is less than some threshold. Probabilistic robustness may be particularly useful in recurring scenarios where historical information pertaining to the reliability of assignments is available.

Probabilistic rates of failure (Dynamic WSS). The failure rates of components of mechanical/electrical or electronic devices are typically defined in terms of probability distributions over time. In general, the shape or type of failure distribution depends upon the component's inherent failure mechanisms.

5.2.1 The Weibull Distribution

The Weibull distribution is commonly used in the field of reliability engineering [95]. It can be used to model a variety of distributions including normal, lognormal, exponential and Rayleigh, which are exhibited in corrosion, diffusion, fatigue, abrasion and many other degradation processes [133]. It is the most widely used distribution in reliability engineering, thus it is ideal for simulating probabilistic failure rates in WSSs, in order to determine the sets of assignments that require repairs.

The probability density function of the 2-parameter Weibull distribution is defined as follows:

$$f(t) = \frac{\gamma}{\eta}(\frac{t}{\eta})^{(\gamma-1)} e^{(-(\frac{t}{\eta})^{\gamma})} \tag{5.1}$$

where γ and $\eta > 0$. γ is the shape parameter of the distribution. If γ is greater than 1, the failure rate is increasing; if γ is less than 1, the failure rate is decreasing; $\gamma = 1$ implies the failure rate is constant. η is the scale parameter and is also known as the characteristic life, where $\frac{1}{\eta}$ is defined as the time at which there is a 0.632 probability that failure will have occurred. Figures 5.1(a) and 5.1(b) illustrate how the various parameters affect the probability and cumulative probability distributions, respectively.

The *cumulative distribution function* (CDF) is the probability of failure before time t. This is relevant for the scenario in which we are interested, i.e. calculating which sets of assignments are likely to fail in a given time-frame. The CDF for the 2-parameter Weibull distribution is as follows:

$$F(t) = 1 - e^{-(\frac{t}{\eta})^{\gamma}}, t \geq 0; \gamma > 0; \eta > 0. \tag{5.2}$$

We use the Weibull distribution to describe the failure rates of assignments whose probability of failure is time-dependent. When this probability is at least some threshold, α, at a given time, τ, then it is deemed brittle and requires a repair solution. Table 5.1 describes the meaning of the different parameters of the Weibull distribution.

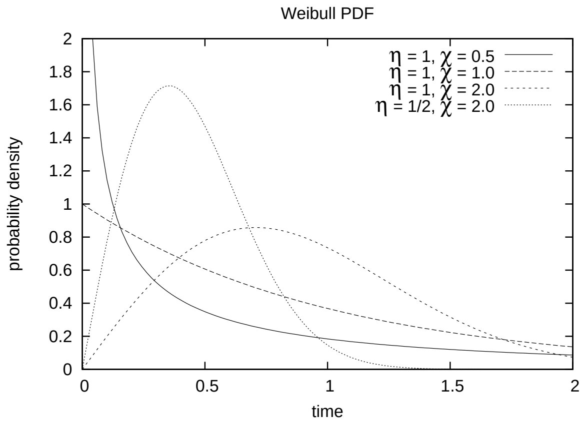

(a) Weibull PDF

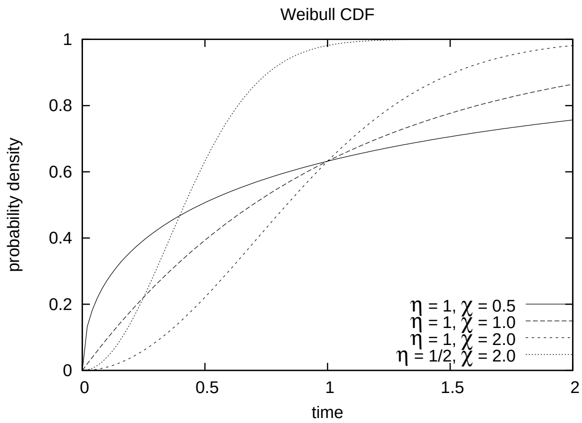

(b) Weibull CDF

Figure 5.1: Weibull distributions.

5.2.2 Other Common Distributions

The Weibull distribution can be used to represent any of the following distributions that are frequently used in failure analysis.

Normal. The normal (or Gaussian) distribution describes equipment failure behavior whose failure rate increases over time. It also models strengths of materials and the lifetime of consumables [95]. It is defined by the mean (μ) and standard deviation (σ). The mean is the expected value of the associated random variable about which it is symmetric.

Table 5.1: Parameters of the Weibull distribution.

Mean (μ)	The mean of the distribution.
Standard Deviation (σ)	The square root of the variance that is a measure the dispersion, of a distribution measured by the second central moment of the distribution.
Scale parameter (α)	$\frac{1}{\alpha}$ is defined as the time at which 63.2% of the device population will have failed.
Shape Factor (γ)	The shape of the hazard curve. When γ = 1, the curve takes on an Exponential distribution. (This is the constant failure rate curve.) When β = 2, the curve takes on a Rayleigh distribution. (This is the linear failure rate curve.) When γ = 3.5, the curve takes on a Normal distribution (approximately).

Lognormal. Lognormal distributions are encountered in metal-fatigue testing, maintainability data (time to repair), chemical-process equipment failures and repairs, crack propagation and other processes whose time to failure has cumulative contributing factors [95]. The lognormal distribution is similar to the normal distribution except that the logarithms of the values of random variables, rather than the values themselves, are normally distributed.

Exponential. For some goods the failure rate remains constant over time. This means that the remaining life of a component is independent of its current age and a used component is assumed to be as good as a new component. Electronic components are a classic example of items that exhibit such failure distributions for most of their useful lives. When variables in a CSP represent actions rather than machine states, the distribution may also be seen as time-independent. Such distributions have a wide range of applications in analyzing the reliability and availability of electronic systems, various queuing networks, and Markov chains [95].

5.3 The Cost of Repair

The SSs framework guarantees the availability of repair solutions in cases where up to a variables may break and b other variables are allowed change. This approach assumes that the cost of changing all variables is the same and that the cardinality of the repair set is a reasonable measure of the total cost of repair. We propose a new metric for measuring the cost of repairing an assignment. This extension to the super solutions framework permits differentiation between the cost of alternative repairs.

5.3.1 Destination and Break Dependencies

The cost of changing an assignment in a solution may have multiple dependencies. We include the original and final assignments as well as the break variable(s) as inputs to a function to determine the cost of making a single change. This provides a more flexible and accurate description of the cost of making changes to a solution. For example, in a job-shop scheduling problem the values associated with variables may represent the various states of a machine, and the cost of alternative state transitions may be different. The cost of a state transition may, furthermore, be influenced by the cause of the break in the solution.

The repair restrictions presented in [58] may be fully expressed using break and destination dependent repair costs by using a value of ∞ as the cost of making changes that are disallowed.

Definition 5.3.1 (Cost of Repair)**.** *The cost of repair describes the cost associated with changing the value of variable x from v_1 to v_2 when A is the set of break variables, $C^{(x,A)}_{v_1 \to v_2} \in \mathbb{R}^+_0$.*

If at most k variables participate in any potential break, there are $\binom{n}{k}$ possible breaks, therefore, the space complexity for storing the repair cost values becomes $O(n^{k+1}d^2)$ when represented extensionally.

5.3.2 Total Cost of Repair

The total cost of repair is computed using a function, $f()$, that combines the cost of reassigning the individual variables that are modified to form the repair solution. Typically, a summation of the repair costs is used to determine the overall cost of repair. However, computing the maximum of the set of the repair costs could also be useful, such as in situations where variables are associated with agents and the imposition of change upon each agent is limited.

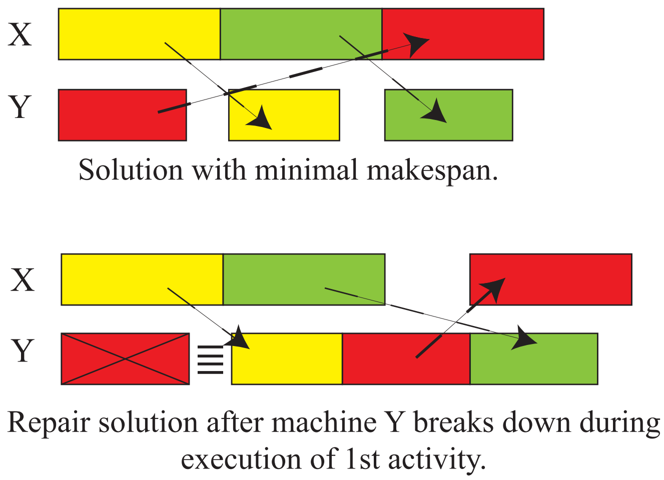

Figure 5.2: An example JSP.

The algorithm that we propose in this chapter requires that $f()$ is monotone non-decreasing in the size of the repair set so that we can terminate the search for a WSS when an assignment is deemed irreparable. If R_1 and R_2 are two repair sets and $R_1 \subset R_2$, then $f(R_1) \leq f(R_2)$. This restriction is necessary because it allows us to cease searching for a repair solution in a branch once a threshold cost, β, has been exceeded. Otherwise, the search for repair solutions would require the computation of a lower bound at each node of the search for a repair. This search would become prohibitively expensive without this constraint on the cost function.

5.3.3 Motivating Example: Job-shop Scheduling

Consider an example of a job-shop scheduling problem (JSP) in which three jobs, each comprising two activities, need to be scheduled. Each activity requires a fixed length of time on machines X and Y. An example solution is presented in Figure 5.2. The arrows indicate precedence constraints between activities. If Machine Y breaks down whilst executing the first activity, there is a period of repair time followed by the execution of a modified schedule so that the interrupted activity is re-scheduled and other activities are also re-scheduled to maintain the precedence constraints. The repair solution clearly has a longer makespan, but another difficulty is that two other activities, besides the one that was interrupted, must be re-scheduled.

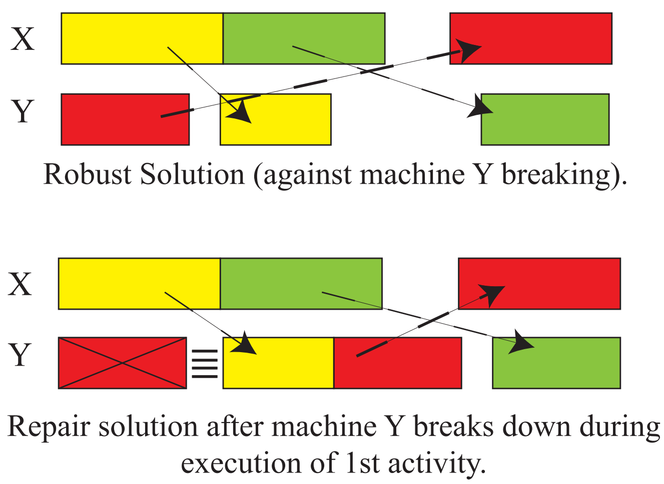

Figure 5.3: An robust solution for JSP example.

A robust solution (see Figure 5.3) allows a repair solution to be constructed without imposing a heavy re-scheduling cost in the event of machine Y breaking down. Super solutions only allows us to reason about the number of changes that are necessary when determining repair costs. In practice, the cost of rescheduling activities may depend on the assignments that failed, and the source and destination assignments for the changed variables. For example, reconfiguring some activities on some machines may require higher labor costs than others. Fixing a machine that breaks may also consume resources so the cause of the failure may also affect costs. The WSSs framework allows us to fully express such cost dependencies. Furthermore, the WSSs framework also permits us to seek repairs for combinations of failures of arbitrary size whose probability of occurrence is greater than a threshold value.

However, in this JSP example we see that we may have to sacrifice optimal makespan in favor of robustness. We assume that machine X is robust and no activity is likely to fail while on this machine. We see that it is possible to delay the start times of the final activity in machine Y so that if the first activity fails on this machine, a repair solution that reschedules only the failed activity is achievable.

5.4 The Weighted Super Solutions Framework

The WSSs framework uses probabilistic failures to determine the sets of variables that may break, and repair costs to accurately measure the difficulty in reassigning variables to form a repair solution. We now define both *static* and *dynamic weighted super solutions* in terms of robustness and cost of repair.

Definition 5.4.1 (Static WSS)**.** *A solution to a CSP is a static weighted super solution, or (α,β)-static WSS, if any set of variables whose probability of their assignments being subsequently invalidated is at least α, can be repaired by reassigning other values to these and other variables with a repair cost of at most β.*

Probabilistic failure rates can also be used to determine the sets of variables that require repairs. The only difference between static and dynamic WSSs is in the selection of assignments for which repairs are necessary. Dynamic-WSS is applicable when variable assignments have probabilistic failure rates over time, such as machine states in factory scheduling. We define *dynamic weighted super solutions* as follows:

Definition 5.4.2 (Dynamic WSS)**.** *A solution to a CSP is a dynamic WSS, or (α,β,τ)-dynamic WSS, if any set of variables whose probability of their assignments being subsequently invalidated is at least α before time τ, can be repaired by reassigning other values to these and other variables with a repair cost of at most β.*

There is one special case of problem worth mentioning, for which there cannot exist a WSS. A constraint program that contains a *backbone variable,* that takes the same *brittle* value in all solutions [87], is necessarily irreparable. If the assigned value of this variable is brittle then there does not exist any robust solution because the failure of this assignment cannot be repaired.

5.4.1 Pedagogical Example: N-Queens

Consider the following pedagogical example of how a WSS may provide solution robustness in a problem where queens must be placed on a board, such that no queen is attacked by any other. In an 8-queens scenario there are many possible solutions, four of which are presented in Figure 5.4. We shall assume, for simplicity, that these are the only possible solutions. The repair cost associated with changing from one solution to the other may differ because the cost of changing the assignment of any variable varies according to its original value. For example, reducing a variable's assigned value may be cheaper than increasing it. Of course, it is possible to stipulate that a variable cannot be changed in a repair solution by setting its cost of repair to ∞.

In this example each queen's assignment is brittle and has a static probability of failure of $\alpha \in [0, 1)$. Only single breaks may occur because the likelihood of a combination of 2 or more breaks occurring is $\alpha^2 < \alpha$. Variables represent queens, and domain values of 1 to 8 represent the column on which to place the queen; therefore, a queen can be placed on any square along a row. We seek a solution that will accommodate any break by finding a repair solution and limiting the movement of queens required to generate a repair solution. We let the repair cost associated with changing the assignment be the difference between the original and final values. This means that the cost of moving a queen m places along a row is m, or more formally $C^{(x,A)}_{v_i \to v_j} = |v_i - v_j|, \forall i, j \in \{1, \ldots, 8\}, i \neq j$.

Consider the solutions for the 8-queens problem presented in Figure 5.4. The costs of changing from one solution to another following a break may vary. For example, given Solution W it would cost $(1 + 1 + 6 + 5 + 0 + 3 + 3 + 1 = 20)$ to move to Solution X. The maximum cost required to move from any one of the four solutions to another for any break is given in Table 5.2(a). Solution Z is the cheapest to repair, by reassigning the queens to the squares in Solution Y. However, there are three queens that occupy the same squares in these two solutions, namely a1, b7 and h3. Table 5.2(c) shows how many possible breaks are reparable by other solutions, thus only five assignments in Solution Z can be repaired by Solution Y. If any of these three assignments break, the solution can instead be repaired by Solution W at a cost of 18. Solution Z, thus, constitutes a (α,18)-WSS.

The total cost of repairing a solution may, alternatively, impose a maximum change for any single variable assignment rather than a bound on the summation of individual costs. This alternative costing is achieved by setting the cost function,

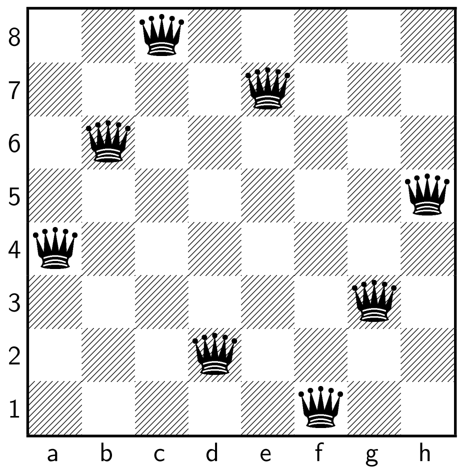

(a) Solution W.

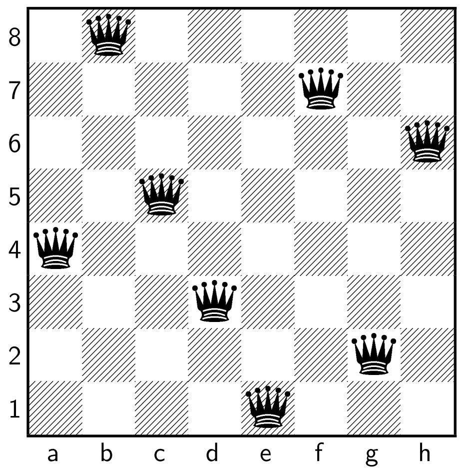

(b) Solution X.

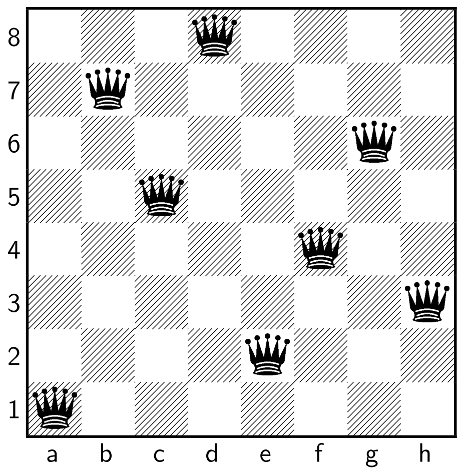

(c) Solution Y.

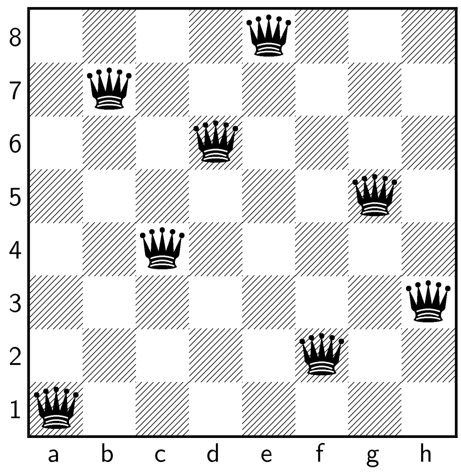

(d) Solution Z.

Figure 5.4: Four possible solutions to the 8-Queens problem.

Table 5.2: 8-Queens solution transition matrices.

(a) Total cost of repair.

	W	X	Y	Z
W	-	20	26	18
X	20	-	22	26
Y	26	22	-	12
Z	18	26	12	-

(b) Max individual repair cost for any single variable.

	W	X	Y	Z
W	-	6	5	5
X	6	-	5	4
Y	5	5	-	4
Z	5	4	4	-

(c) Number of breaks repaired.

	W	X	Y	Z
W	-	7	8	8
X	7	-	7	8
Y	8	7	-	5
Z	8	8	5	-

$f()$, as described in § 5.3.2, to be the `max` operator. Table 5.2(b) shows the cost of repairing solutions under this condition.

5.4.2 WSS Algorithm Description

The WSSs framework requires several modifications to the algorithms used to find classical SSs [58, 59], to support repair solutions for *sets of break variables* of different size and repair sets of *arbitrary cardinality*. WEIGHTED-SUPER-SOLVE (Algorithm 3) can be used to find both static and dynamic WSSs. When τ is not defined, denoted $\emptyset$, the algorithm detects that a static WSS is required in the `failure` procedure and determines an assignment's probability of failure using probabilities, $\alpha_{(x,v)}$, otherwise Weibull parameters $\gamma_{(x,v)}$ and $\eta_{(x,v)}$ are used to compute the probability of failure by time τ. Any combination of assignments whose probability of failure is at least α requires a repair. We examine all combinations of assignments whose size is at most k, the largest possible break.

Algorithm 3: WEIGHTED-SUPER-SOLVE

input : α, β, τ, RepairCosts: C, CSP: P=$\{\mathcal{X}, \mathcal{D}, \mathcal{C}\}$ // Let $\tau = \emptyset$ for static-WSS
output: S: an (α,β,τ)-WSS: R: the set of repair solutions
begin
  $S \leftarrow \emptyset$ // Solution
  $R \leftarrow \emptyset$ // Set of repair solutions
  $Past \leftarrow \emptyset$ // Ordered set of assigned variables
  AC(P, S) // Perform arc-consistency
  backtrack(P,S,$Past$,R,0,α,β,τ,0)
end

A repair solution, R_b, is provided for every possible set of break variables b. The `backtrack` procedure is called from WEIGHTED-SUPER-SOLVE and attempts to extend the current partial assignment by choosing a variable and assigning it a value. Backtracking may occur for one of two reasons: we cannot extend the current partial assignment to satisfy the given constraints, or it cannot be associated with a repair solution whose cost is less than or equal to β for a possible break.

Procedure `backtrack(`P,S,$Past$,R,lvl,α,β,τ,m`)`:Boolean

```
begin
    if 𝒳 = Past then return true
    choose x ∈ 𝒳 \ Past
    b ← ∅ // set of break variables
    Past[lvl] ← x
    foreach v ∈ 𝒟(x) do
        save 𝒟, m and R
        m ← max(m, failure({x}, S, τ))
        k ← ⌊log_m α⌋ // Max size of any break
        S ← S ∪ {(x, v)}
        if AC(P,S) then
            foreach b ∈ 𝒫(Past), | b |≤ k, failure(b, S, τ) ≥ α do
                if ¬reparable(P,S,Past,R_b,0,β) then break
                if backtrack(P,S,Past,R,lvl + 1,α,β,τ,m) then return true
        restore 𝒟, m and R
        S ← S \ (x, v)
    Past[lvl] ← ∅
    return false
end
```

The procedure `reparable` searches for *partial* repair solutions using backtracking and attempts to extend the last repair found, just as in (1,*b*)-super solutions; the differences being that a repair is provided for a set of break variables rather than a single variable and the cost of repair is considered.

The procedure `check-wss-repair` determines the cost of the transition to the repair solution and verifies consistency. A summation operator is used to determine the overall cost of repair in this case, but this may be any other monotone, non-decreasing function. Note that summation is non-decreasing because all costs are non-negative. This procedure can also include auxiliary break and repair restrictions described in [58], if required.

Theorem 5.4.1. WEIGHTED-SUPER-SOLVE *terminates and is sound and complete.*

Proof.
Termination: The algorithm never revisits any partial assignment or repair for a

Procedure `reparable(`P,S,$Past$,R_b,lvl,β`)` :Boolean

begin
  if $lvl = |S|$ **then** return true
  $y \leftarrow Past[lvl]$
  for $v \leftarrow R_b[y]$ *to* $\max_{init} \mathcal{D}(y)$ *// Last value in lex order*
  do
    if $y \notin b$ *or* $S[y] \neq v$ **then** $R_b[y] \leftarrow v$
    if *check-wss-repair(*P,S,$Past$,R_b,lvl,β*)* **then**
      if *reparable(*P,S,$Past$,R_b,$lvl+1$,β*)* **then** return true
  $R_b[y] \leftarrow \min(\mathcal{D}(y))$
  return false
end

Procedure `check-wss-repair(`P,S,$Past$,R_b,lvl,β`)` : Boolean

begin
  $cost \leftarrow 0$
  for $i \leftarrow 0$ *to* lvl **do**
    $y \leftarrow Past[i]$
    if $y \notin b$ *and* $R_b[y] \neq S[y]$ **then**
      $cost \leftarrow cost + C(x, S[y], R_b[y])$
  if $cost > \beta$ **then** return false
  return consistency of the l first values in R_b
end

given break set, of which there are finitely many.

Soundness: $\forall b \in \mathcal{P}(Past)$, R_b is the first repair solution of S, when taken in lexicographical order, for b in the problem restricted to $Past$.

Completeness: MAC [114] is complete, therefore, no partial assignment is omitted before checking for reparability.[†] The check for reparability starts from the last repair found. The cost of repair function is monotone, non-decreasing so no assignment before this last repair in the search tree can be extended to the current variable because each prior partial assignment had a minimum cost of repair that exceeded β. □

We also show that finding a WSS is $\mathcal{NP}$-complete in general for any fixed α

[†]Note that it is possible to implement the algorithm using any complete consistency propagator.

Procedure `failure(`b, S, τ`)` : Real

begin
- $fail \leftarrow 1.0$
- **for** $x \in b$ **do**
  - $v \leftarrow S[x]$
  - **if** $\tau = \emptyset$ **then** $fail \leftarrow fail \times \alpha_{(x,v)}$ // Static WSS
  - **else** $fail \leftarrow fail \times$ CDF($\eta_{(x,v)}$,$\gamma_{(x,v)}$,τ) // Dynamic WSS
- return $fail$

end

(and τ for the case of dynamic WSSs). We define two decision problems: (α,β)-STATIC WEIGHTED SUPER SOLUBILITY and (α,β,τ)-DYNAMIC WEIGHTED SUPER SOLUBILITY as the problems of deciding whether there exists an (α,β)-static WSS or an (α,β,τ)-dynamic WSS to a problem, respectively.

Lemma 5.4.1. *Let $m = \max\left(\bigcup_{(x,v)\in s} \alpha_{(x,v)}\right)$, where $\alpha_{(x,v)}$ is the constant probability of failure associated with the assignment of value v to variable x for static-WSS and the probability of failure at time τ in the dynamic case. If $\lfloor \log_m \alpha \rfloor$ is bounded by a constant k, the number of possible breaks requiring repair solutions is polynomial in k.*

Proof. For each solution S there must be a repair solution for each subset $s \in \mathcal{P}(S)$, the power-set of S, whose probability of failure is greater than or equal to α, i.e. $\prod_{(x,v)\in s} \alpha_{(x,v)} \geq \alpha$. But $\alpha_{(x,v)} \leq m$, so we can say that:

$$\prod_{(x,v)\in s} \alpha_{(x,v)} \leq m^{|s|}.$$

If s requires a repair we have

$$\alpha \leq m^{|s|} \therefore \log \alpha \leq \mid s \mid \cdot \log m.$$

Since $\log \alpha$ and $\log m$ are both negative:

$$\frac{\log \alpha}{\log m} \geq \mid s \mid \therefore \log_m \alpha \geq \mid s \mid.$$

Since $\lfloor \log_m \alpha \rfloor \leq k$ and $\mid s \mid \in \mathbb{Z}$, $k \geq \mid s \mid$.

The size of the largest possible break needing consideration for repair is $\leq k$, therefore, at most $\binom{n}{k}$ repair solutions are necessary, where n is the number of variables. □

Theorem 5.4.2. *(α,β)-*STATIC WEIGHTED SUPER SOLUBILITY *and (α,β,τ)-* DYNAMIC WEIGHTED SUPER SOLUBILITY *are $\mathcal{NP}$-complete when $\lfloor \log_{\max(\alpha_{(x,v)})} \alpha \rfloor$ is bounded by a constant k.*

Proof.
Hardness: To show they are in $\mathcal{NP}$ we need a polynomial witness. This is simply an assignment of the variables that satisfies the constraints in the problem and, for each of the $\mathcal{O}(n^k)$ possible breaks, the set of repair values. This is polynomial for fixed k from Lemma 5.4.1.
Completeness: We present a reduction from binary CSP. Duplicates of each value in the domain of all variables are created. Constraints are added to behave equivalently on the duplicate (primed) values. Additional constraints are added to enforce that a solution can involve only either primed or duplicate values. This problem is satisfiable *iff* the original problem is also satisfiable. If a solution does exist, a WSS does also because any set of k values may be primed to form a repair solution. □

Whilst an upper bound on α is fixed, there is no such constraint on β. If α is unbounded, then (α,β)-STATIC and (α,β,τ)-DYNAMIC WEIGHTED SUPER SOLUBILITY are in PSPACE. Note that in Procedure `backtrack`, we determined the size of the largest possible break to be k.

5.4.3 Optimizing Robustness

It may not always be possible to find a robust solution for given values of α, β and τ. In such situations the problem constraints may be relaxed in several ways so that different trade-off scenarios are considered. A branch and bound search can trivially incorporate optimization over any one of these parameters in the search tree.

1. Minimize α: This effectively decreases α so that we find a WSS that repairs as many combinations of the most brittle assignments as possible.

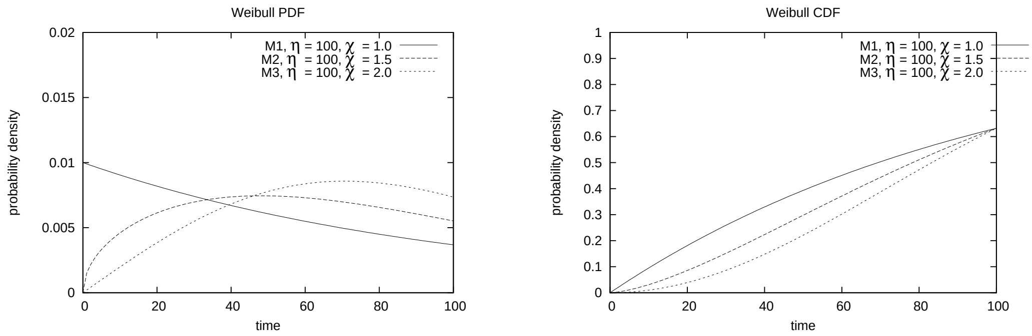

(a) Weibull PDF

(b) Weibull CDF

Figure 5.5: Machine failure rates.

2. Minimize β: This places as small a bound as possible on repair costs for any potential break.

3. Maximize τ (Dynamic-WSS only): The solution is reparable for as long as possible, thus maximizing the mean time to irreparable failure.

The non-existence of a robust solution need not preclude the use of the WSSs framework and optimization of robustness may provide a valuable alternative approach. Previously in Chapter 4, we studied optimization of robustness for CAs when instances were unsatisfiable.

5.5 Application: Job-shop Scheduling

We applied the WSSs framework to job-shop scheduling problems (JSPs) to demonstrate its versatility. JSPs consist of scheduling tasks on machines whose reliability may be time dependent. We modeled failure probabilities of these machines using Weibull probability distributions and repair costs differed for all machines.

Each problem consisted of 3 machines and 4 jobs, each comprising a sequence of 3 activities. Each activity required each machine for a duration chosen over a uniform random distribution [1,5]. The objective was to schedule all activities so that the precedence and resource constraints amongst activities were respected whilst minimizing the overall makespan.

In these experiments machines $M_1 - M_3$ exhibited different failure rates and we assumed Weibull distributions with γ as 1.0, 1.5 and 2.0, respectively, and a random repair period from [1,5] on each machine following a failure. The CDF in Figure 5.5(b) describes the likelihood of each machine failing by time τ. We let the characteristic life $\eta = 100$ so that each machine had a probability of failure of 0.632 by this time. We modeled the activities as variables, whose domain values represented start-times. The duration of each activity allowed us to determine the end-time for each activity on each machine given the start-time. We could, therefore, assign a static probability of failure using the difference in the CDF on the relevant machine between these two times. The brittleness of a variable (activity on a particular machine), therefore, depended upon its assigned value (start-time). Our solver used a dynamic, minimum-degree heuristic over the variables and values were chosen in lexicographical order.

We assumed that the cost of repairing a solution was machine dependent. Given a break in a solution, the costs of changing an activity on $M_1 - M_3$, were 25, 50 and 75, respectively. Rescheduling activities on different machines may vary because of calibration/setup/labor costs, etc. We examined the trade-off between robustness and makespan for 50 randomly generated instances. Problems had 12 variables, with domain sizes containing approximately 15-20 values. The last value in the domain was an upper bound on the latest start-time required for that activity. A solution was found by bounding the makespan to an initial lower bound. If a solution was not found, this makespan was incremented by one and the problem was resolved. This process was repeated until a valid solution was found.

Tables 5.3(a)-5.3(d) show how the minimal makespan, number of nodes visited in the search tree and breaks checked for reparability vary with α and β. It is noticeable that the optimal makespan decreases as β increases and is more pronounced when $\beta > 200$. When β is low, reparability is inhibited because fewer repair solutions can be considered, so the makespan increases. It is also clear that as α decreases there is a larger number of activities that require repair solutions so the makespan increases as some repair solutions require later start-times.

Finding a WSS can become computationally expensive when α is low, increasing the number of combinations of breaks needing consideration. For example,

Table 5.3: Results (4x3 JSP).

(a) $\alpha = 0.01$

β	mksp	nodes×1000	breaks
0	18.82	67,624	303,848
50	18.92	638,739	743,005
100	19.26	533,368	298,325
150	18.88	500,723	234,598
200	18.32	383,725	136,273
250	18.42	577,310	159,820
300	18.20	449,798	227,233
350	18.48	821,485	174,816
400	17.84	525,212	209,574

(b) $\alpha = 0.02$

β	mksp	nodes×1000	breaks
0	18.28	15,878	53,369
50	18.32	354,186	128,888
100	18.32	290,072	78,951
150	18.28	412,006	105,511
200	18.12	415,698	74,169
250	17.32	243,148	45,748
300	17.80	360,334	53,963
350	17.50	281,017	60,642
400	16.98	421,379	63,510

(c) $\alpha = 0.03$

β	mksp	nodes×1000	breaks
0	18.40	18,915	41,493
50	17.68	33,864	10,959
100	17.80	175,650	40,622
150	18.24	182,240	38,803
200	17.74	193,117	24,728
250	18.18	257,233	73,544
300	17.68	290,801	52,235
350	17.70	370,700	44,384
400	18.16	318,035	59,601

(d) $\alpha = 0.04$

β	mksp	nodes×1000	breaks
0	16.96	3,748	5,329
50	16.98	6,005	1,159
100	17.06	473	163
150	17.14	1,517	362
200	17.08	10,853	1,126
250	16.84	3,445	639
300	16.20	66,100	6,260
350	15.96	249	33
400	15.82	8,158	1,123

from Tables 5.3(a)-5.3(d) we can see how the number of nodes in the search tree can grow exponentially. As the number of possible breaks decreases, the number of concurrent searches for separate repair solutions also decreases. The search-effort is greatly reduced when $\alpha = 0.04$ (Table 5.3(d)) because fewer assignments require repair solutions. When α was increased to 0.05, only a negligible number of assignments were deemed irreparable.

The number of nodes visited and breaks checked are not tightly correlated with β for the following reason. When β is low, searches for repair solutions may fail quickly, whereas when it is high, the repair searches visit more nodes but the success rate increases thereby leading to initial solutions more quickly. These two effects counteract one another as β increases.

Tables 5.3(a)-5.3(d) also show the number of breaks checked for reparability before we extend the search tree (column breaks). This corresponds to the number of calls made to the `reparable` procedure. From these results it is clearly important that repair solutions are not sought for assignments that are robust in order to minimize the computational burden. This is why probabilistic failures in the WSSs framework are critical in subtly differentiating between assignments that are brittle and those that are robust.

5.6 Summary

The weighted super solutions (WSS) framework extends the classical super solutions framework [59] in two important ways: firstly, it can model the most likely causes of failure in a solution and secondly, it can reason about the costs of alternative repair solutions.

The framework provides an expressive means of establishing robust solutions when variable assignments are invalidated for some reason. This framework is both versatile and practical in many application domains, such as scheduling [64], where reasoning about uncertainty and the cost of repair is important. We demonstrated its applicability for job-shop scheduling problems. The purpose of developing this new framework was to accommodate uncertainty in combinatorial auctions more flexibly and facilitate a more expressive means of measuring the cost of repairing solutions. The WSSs framework is powerful enough to cater for such an auction mechanism as we shall show in the following chapter.

The following chapter will provide an in-depth analysis of solution robustness for combinatorial auctions and examine the benefits/opportunities it can provide for the bid-taker. We shall utilize the extensions developed in the WSSs framework to full effect by developing an auction mechanism that enhances reparability using probabilistic bid withdrawals and bid-dependent repair costs to facilitate backtracking upon assignments of items. We also describe some ancillary, yet highly significant, features of robust solutions generated by the WSSs framework that provide added value to bid-takers.

Chapter 6

Weighted Super Solutions for Combinatorial Auctions

Thus far, we have studied the bid-taker's exposure problem and demonstrated the significance of the problem caused by winning-bid withdrawal. We explored achieving solution robustness using the SSs framework and described its short-comings in terms of its application to combinatorial auctions. We then proposed a probabilistic framework with a more flexible means of determining reparability requirements and a metric for expressing repair costs more flexibly. The WSSs framework was the culmination of these extensions to the classical SSs frame-work and we shall now demonstrate the clear advantages of using this approach for finding robust solutions for CAs.

Firstly, we apply the WSSs framework to combinatorial auctions and demon-strate its usefulness in different economically motivated settings. We also propose an auction model that enhances reparability by introducing mandatory *mutual bid bonds*, that may be seen as a form of leveled commitment contract [116, 117]. We present an extensive empirical evaluation of the approach presented in this chapter with very encouraging results.

6.1 Combinatorial Auctions and the WSSs Framework

In the following section we describe how the WSSs framework can be applied to combinatorial auctions. We also provide an in-depth example to show how the search for a robust solution to a CA is conducted. Definition 6.1.1 outlines more precisely the requirements for a robust solution in the context of a WSS.

Definition 6.1.1 (Robust Solution for a CA)**.** *A robust solution for a combinatorial auction is one where any subset of successful bids whose probability of withdrawal is at least α can be repaired by reassigning items at a cost of at most β to other previously losing bids, in order to form a repair solution whose revenue is at least a fraction, γ, of optimal revenue.*

We also define a restricted form of robust solution that can withstand a single bid withdrawal only.

Definition 6.1.2 (1-robust Solution for a CA)**.** *A 1-robust solution for a combinatorial auction is one where any* single *bid whose probability of withdrawal is at least α can be repaired by reassigning items at a cost of at most β to other previously losing bids, in order to form a repair solution whose revenue is at least a fraction, γ, of optimal revenue.*

6.1.1 Probabilistic Failure

The choice of failure model is relatively straightforward for CAs. We are assuming an expedient transaction process whereby payment is immediate following an auction and the time lapse is negligible between the announcement of the solution and the break in the solution. The probabilities of failure are, therefore, time-independent so a static-WSS approach (Definition 5.4.1) is appropriate for CAs. A dynamic-WSS is applicable for cases in which assignments have failure rates over time. Recall that, in the case where the transaction process is not quickly completed and bid withdrawals are time-dependent, the value of the items to bidders may change by the time of the withdrawal. For this reason it is unreasonable to assume that the losing bids still stand, and another auction may be necessary.

There is a probability of failure associated with each winning bid and sets of winning bids whose probability of failure is greater than or equal to a threshold value, α, require repair solutions.

6.1.2 Repair Costs for Combinatorial Auctions

The WSSs framework considers the *cost* of repair required, rather than simply the number of assignments modified, to form an alternative solution. For CAs this may be a measure of the compensation penalties paid to winning bidders to break existing agreements. WSSs offer a means of expressing which variable's re-assignments incur a heavy cost and those that are are easily implemented [64]. For example, consider an auction with rules which stipulate that a reneging bidder may have items from other non-withdrawn winning bids disqualified. A WSS could model this situation by setting the repair cost for winning bids by a reneging bidder to be zero, thus facilitating withdrawal of bids from an unreliable bidder at no extra charge.

It may be desirable to reassign items to different bidders in order to find a repair solution of satisfactory revenue. Compensation may have to be paid to bidders who lose items during the formation of a repair solution. The repair cost of a bid reflects the cost of changing its state. For winning bids this may reflect the necessary compensation penalty for the bid-taker to break the agreement (if such breaches are permitted), whereas for previously losing bids this is a free operation. The total amount of compensation payable to bidders may depend upon other factors, such as the cause of the break. There is a limit to how much these overall repair costs should be, and this is given by the value β. This value may not be known in advance and may depend upon the break. Therefore, β may be viewed as the fund used to compensate winning bidders for the unilateral withdrawal of their bids by the bid-taker.

For example, a (0.1,100)-WSS allows any set of bids whose probability of withdrawal is at least 0.1 to be repaired with changes to the original robust solution at a cost of at most 100. This cost stems from compensation payments that need to be paid to bidders that were declared as winners but had those items revoked when the items were reallocated in a repair solution following a bid withdrawal by another party.

6.1.3 WSS Search for Combinatorial Auctions

The depth-first search for a WSS (see Algorithm 3) maintains arc-consistency [114] at each node of the tree. Recall that as search progresses, the reparability of each previous assignment is verified at each node by extending a partial repair solution to the same depth as the current partial solution. This may be thought of as maintaining concurrent search trees for repairs. A repair solution is provided for every possible set of winning-bid withdrawals.

The WEIGHTED-SUPER-SOLVE algorithm attempts to extend the current partial assignment by choosing a variable and assigning it a value. For CAs with binary variables representing bids, 1 signifies a winning bid and 0 signifies losing. Clearly, repair solutions are only required for winning bids so repairs are only sought for winning bids.

Backtracking may occur for one of two reasons. Firstly, the branch and bound search cannot extend the assignment to attain a solution that satisfies the lower bound on revenue. Alternatively, the current partial assignment cannot be associated with a repair solution whose cost of repair is less than β should a break occur. We now step through an example of finding a WSS for a CA to illustrate how the algorithm works.

Example 6.1.1. *We shall step through the example given in Table 3.1 (page 44) when searching for a WSS. Each bid is represented by a single variable with domain values of 0 and 1, the former representing bid-failure and the latter bid-success. The probability of failure of the variables is 0.1 when they are assigned to 1 and 0.0 otherwise. The problem is initially solved using an ILP solver such as* `lp_solve` *[11] or* `CPLEX` *[66], and the optimal revenue is found to be 200. A fixed percentage of this revenue can be used as a threshold value for a robust solution and its repairs. The bid-taker wishes to have a robust solution so that if a single winning bid is withdrawn, a repair solution can be formed without withdrawing items from any other winning bidder. This example may be seen as searching for a (0.1,0)-weighted super solution, β is 0 because no funds are avail-*

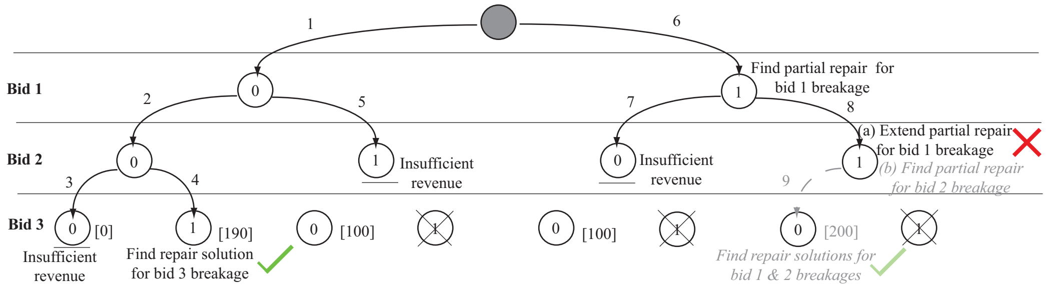

(a) Search for WSS.

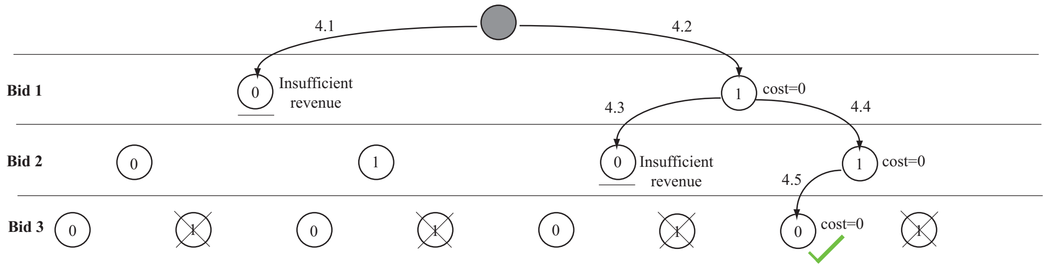

(b) Search for a repair for bid 3 breakage.

Figure 6.1: Search tree for a WSS without item withdrawal.

able to compensate the withdrawal of items from winning bidders. The bid-taker is willing to compromise on revenue, but only by 5%, say, of the optimal value.

Bids 1 and 3 cannot both succeed, since they both require item A, so a constraint is added precluding the assignment in which both variables take the value 1. Similarly, bids 2 and 3 cannot both win so another constraint is added between these two variables. Therefore, in this example the set of CSP variables is $V = \{x_1, x_2, x_3\}$, whose domains are all $\{0, 1\}$. The constraints are $x_1 + x_3 \leq 1$, $x_2 + x_3 \leq 1$ and $\sum_{x_i \in V} v_i x_i \geq 190$, where v_i reflects the relevant bid-amounts for the respective bid variables. In order to find a robust solution of optimal revenue we seek to maximize the sum of these amounts, $max \sum_{x_i \in V} v_i x_i$.

When all variables are set to 0 (see Figure 6.1(a) branch 3), this is not a solution because the minimum revenue of 190 has not been met, so we try assigning bid 3 to 1 (branch 4). This is a valid solution but this variable is brittle because there is a 10% chance that this bid may be withdrawn (see Table 3.1). Therefore, we need to determine if a repair can be formed should it break. The search for a repair begins at the first node, see Figure 6.1(b). Notice that value 1 has been *removed from bid 3 because this search tree is simulating the withdrawal of this bid. When bid 1 is set to 0 (branch 4.1), the maximum revenue solution in the remaining subtree has revenue of only 100, therefore, search is discontinued at that node of the tree. Bid 1 and bid 2 are both assigned to 1 (branches 4.2 and 4.4) and the total cost of both these changes is still 0 because no compensation needs to be paid for bids that change from losing to winning. With bid 3 now losing (branch 4.5), this gives a repair solution of 200. Hence* $\langle 0, 0, 1\rangle$ *is reparable and, therefore, a WSS. We continue our search in Figure 6.1(a) however, because we are seeking a robust solution of* optimal *revenue.*

When bid 1 is assigned to 1 (branch 6) we seek a partial repair for this variable breaking (branch 5 is not considered since it offers insufficient revenue). The repair search sets bid 1 to 0 in a separate search tree, (not shown), and control is returned to the search for a WSS. Bid 2 is set to 0 (branch 7), but this solution would not produce sufficient revenue so bid 2 is then set to 1 (branch 8). We then attempt to extend the repair for bid 1 (not shown). This fails because the repair for bid 1 cannot assign bid 2 to 0 because the cost of repairing such an assignment would be ∞*, given that the auction rules do not permit the withdrawal of items from winning bids. A repair for bid 1 breaking is, therefore, not possible because items have already been awarded to bid 2. A repair solution with bid 2 assigned to 1 does not produce sufficient revenue when bid 1 is assigned to 0. The inability to withdraw items from winning bids implies that* $\langle 1, 1, 0\rangle$ *is an irreparable solution when the minimum tolerable revenue is greater than 100. The italicized comments and dashed line in Figure 6.1(a) illustrate the search path for a WSS if both of these bids were deemed reparable.* $\triangle$

§ 6.2 introduces an alternative auction model that allows the bid-taker to receive compensation for breaks and, in turn, use this payment to compensate other bidders for withdrawal of items from winning bids. This will enable the reallocation of items and permit the establishment of $\langle 1, 1, 0\rangle$ as a second WSS for this example.

6.2 Mutual Bid Bonds: A Backtracking Mechanism

Some auction solutions are inherently brittle and it may be impossible to find a robust solution. If we can alter the rules of an auction so that the bid-taker can retract items from winning bidders, then the reparability of solutions to such auctions may be improved. In this section we propose an auction model that permits bid and item withdrawal by the bidders and bid-taker, respectively.

6.2.1 Proposed Model and Justification

We propose a model that incorporates *mutual bid bonds* to enable solution reparability for the bid-taker, a form of insurance against the winner's curse for the bidder whilst also compensating bidders in the case of item withdrawal from winning bids. Motivated by work on leveled commitment contracts [117], we propose that such "*winner's curse and bid-taker's exposure*" insurance comprise a fixed percentage, $\kappa \in (0, 100]$, of the bid amount for all bids. Money is held in escrow so that the bid-taker may be compensated for any withdrawals by a bidder, and likewise, winning bidders may be compensated for item revocation by the bid-taker.

Such mutual bid bonds are mandatory for each bid in our model†. The conditions attached to the bid bonds are that the bid-taker be allowed to annul winning bids (item withdrawal) when repairing breaks elsewhere in the solution. In the interests of fairness, compensation is paid to bidders from whom items are withdrawn and is equivalent to the penalty that would have been imposed on the bidder should he have withdrawn the bid.

Combinatorial auctions impose a heavy computational burden on the bidder so it is important that the hedging of risk should be a simple and transparent operation for the bidder so as not to further increase this burden unnecessarily. We also contend that it is imperative that the bidder knows the potential penalty for withdrawal in advance of bid submission. This information is essential for bidders when determining how aggressive they should be in their bidding strategy. Bid

†Making the insurance optional may be beneficial in some instances. If a bidder does not agree to the insurance, it may be inferred that he may have accurately determined the valuation for the items and, therefore, less likely to fall victim to the winner's curse. The probability of such a bid being withdrawn may be less, so a repair solution may be deemed unnecessary for this bid. On the other hand, it decreases the reparability of solutions.

bonds are commonplace in procurement for construction projects. Usually they are mandatory for all bids, are a fixed percentage, κ, of the bid amount and are unidirectional in that item withdrawal by the bid-taker is not permitted. Mutual bid bonds may be seen as a form of leveled commitment contract in which both parties may break the contract for the same fixed penalty. Such contracts permit unilateral decommitment for pre-specified penalties. Sandholm *et al.* showed that this can increase the expected payoffs of all parties and enables deals that would be impossible under full commitment [116, 118, 120].

In practice a bid bond typically ranges between 5 and 20% of the bid amount [67, 98]. If the decommitment penalties are the same for both parties in all bids, κ does not influence the reparability of a given set of bids. It merely influences the levels of penalties and compensation transacted by agents. Low values of κ incur low winning-bid withdrawal penalties and simulate a dictatorial bid-taker who does not adequately compensate bidders for item withdrawal. Andersson and Sandholm [3] found that myopic agents reach a higher social welfare quicker if they act selfishly rather than cooperatively when penalties in leveled commitment contracts are low. Increased levels of bid withdrawal are likely when the penalties are low also.

High values of κ tend towards full-commitment and reduce the advantages of such "winner's curse & bid-taker's exposure" insurance. The penalties paid are used to fund a reassignment of items to form a repair solution of sufficient revenue by compensating previously successful bidders for withdrawal of the items from them.

6.2.2 Example with backtracking

We shall now step through the previous example assuming that there are mutual bid bond agreements between all bidders and the bid-taker.

Example 6.2.1. *Consider the example given in Table 3.1 (page 44) once more, where the bids also comprise a mutual bid bond of 5% of the bid amount. If a bid is withdrawn, the bidder forfeits this amount and the bid-taker can then compensate winning bidders whose items are withdrawn when trying to form a*

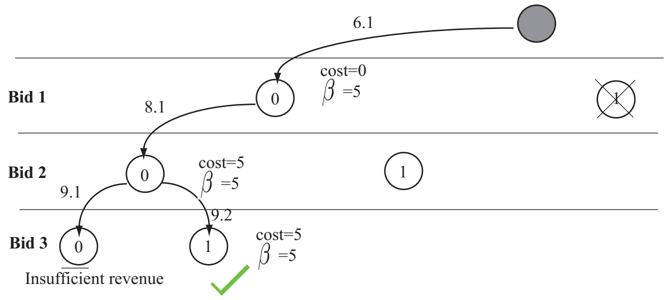

(a) Search for a repair for bid 1 breakage.

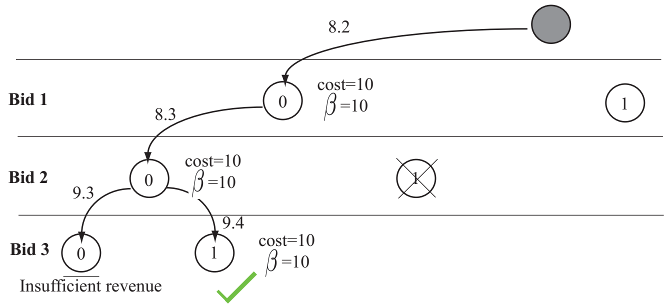

(b) Search for a repair for bid 2 breakage.

Figure 6.2: Repair search tree for breaks 1 and 2, $\kappa = 0.05$.

repair solution later. The search for repair solutions for breaks to bid 1 and bid 2 appear in Figures 6.2(a) and 6.2(b), respectively[†]*.*

When bid 1 breaks, there is a compensation penalty paid to the bid-taker equal to 5 that can be used to fund a reassignment of the items. We, therefore, set β *to 5 and this becomes the maximum expenditure allowed to withdraw items from winning bidders.* β *may also be viewed as the size of the fund available to facilitate backtracking by the bid-taker. When we extend the partial repair for bid 1 so that bid 2 loses an item (branch 8.1), the overall cost of repair increases to 5, due to this item withdrawal by the bid-taker, and is just within the limit given by* β*. In Figure 6.1(a) the search path follows the dashed line and sets bid 3 to be 0 (branch 9). The repair solutions for bids 1 and 2 can be extended further by assigning bid 3 to 1 (branches 9.2 and 9.4). Therefore,* $\langle 1, 1, 0 \rangle$ *may be considered*

[†]The actual implementation of WEIGHTED-SUPER-SOLVE checks previous solutions to see if they can repair breaks before searching for a new repair solution. $\langle 0, 0, 1 \rangle$ is a solution that has already been found so the search for a repair in this example is not strictly necessary but is described for pedagogical reasons.

a robust solution. Recall, that previously this was not the case. △

Using mutual bid bonds thus increases reparability and allows a robust solution of revenue 200 as opposed to 190, as was previously the case.

6.3 Experiments

We have used the Combinatorial Auction Test Suite (CATS) [78] to generate sample auction data. We generated 100 instances of problems in which there are 20 items for sale and 100-2000 bids that may be dominated in some instances†. Such dominated bids can participate in repair solutions although they do not feature in optimal solutions. CATS uses economically motivated bidding patterns to generate auction data in various scenarios. We use sensitivity analysis to examine the brittleness of optimal solutions and hence determine the types of auctions most likely to benefit from a robust solution. We then establish robust solutions for CAs using the WSSs framework.

6.3.1 Sensitivity Analysis for the WDP

We previously used sensitivity analysis (see § 3.1.4) to determine the effects of a single winning-bid withdrawal on optimal solutions to combinatorial auctions. The following four distributions were examined: airport take-off/landing slots (`matching`), electronic components (`arbitrary`), property/spectrum-rights (`regions`) and transportation (`paths`). This analysis proved that the "bid-taker's exposure problem" poses a significant danger to risk averse auctioneers. The sensitivity of optimal solutions may also be used as a baseline against which robust solutions may be measured. If the expected utility of the optimal solution is greater than that of the optimal WSS, then the optimal solution may be selected as the winner. These results are repeated here in Figure 6.3 for ease of comparison with the WSS results.

In this chapter we focus upon the `regions` (Figure 6.3(c)) and `arbitrary` (Figure 6.3(d)) distributions because the sensitivity analysis indicated that these

†The CATS flags included int_prices with the bid_alpha parameter set to 1000.

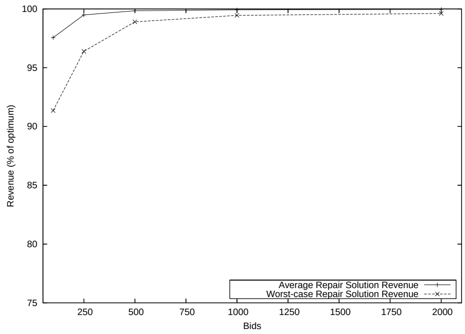

(a) `paths`.

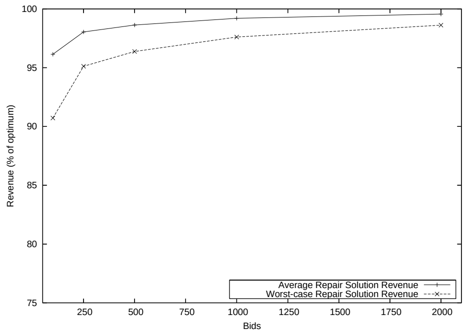

(b) `matching`.

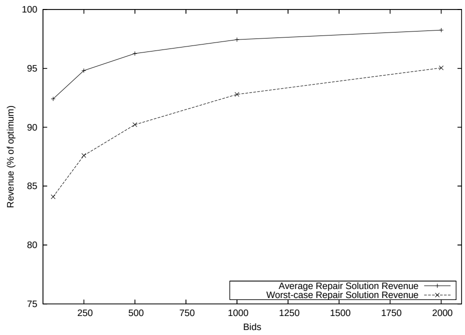

(c) `regions`.

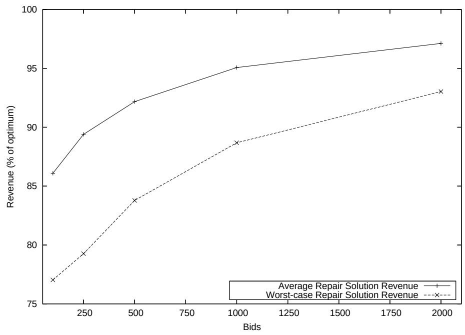

(d) `arbitrary`.

Figure 6.3: Sensitivity of bid distributions to single winning-bid withdrawal.

types of auctions produce optimal solutions that tend to be most brittle and, therefore, stand to benefit most from solution robustness. We ignore the auctions with 2000 bids because the sensitivity analysis has indicated that these auctions are inherently robust with a very low average drop in revenue following a winning-bid withdrawal. They would also be very computationally expensive, given the extra complexity of finding robust solutions. The search for robust solutions exhibits decreasing relative returns for increased search effort as the auction size increases. Robust solutions are, therefore, more applicable to smaller auctions.

6.3.2 Hybrid CP Model for Combinatorial Auctions

A pure CP approach needs to be augmented with global constraints that incorporate operations research techniques to increase pruning sufficiently so that thousands of bids may be examined. Global constraints exploit special-purpose filtering algorithms to improve performance [107]. There are a number of ways to speed up the search for a WSS in a CA, although this is not the main focus of our work. Polynomial matching algorithms may be used in auctions whose bid length is short, such as those for airport landing/take-off slots for example. The integer programming formulation of the WDP stipulates that a bid either loses or wins. If we relax this constraint so that bids can partially win, this corresponds to the linear relaxation of the problem and is solvable in polynomial time. At each node of the search tree we can quickly solve the linear relaxation of the remaining problem in the subtree below the current node to establish an upper bound on remaining revenue. If this upper bound plus revenue in the parent tree is less than the current lower bound on revenue, search at that node can cease. The (continuous) LP relaxation thus provides a vital speed-up in the search for WSSs, which we have exploited in our implementation. The LP formulation is as follows:

$$max \sum_{x_i \in V} v_i x_i$$

$$s.t. \sum_{j|i \in S_j} x_j \leq 1,\ \forall i \in \{1 \dots m\},\ \ x_j \geq 0,\ \ x_j \in \mathbb{R}.$$

Additional techniques, that are outlined in [115], can aid the scalability of a CP approach but our main aim in these experiments is to examine the robustness of various auction distributions and consider the tradeoff between robustness and revenue. The WSS solver we have developed is an extension of the SS solver presented in [58, 59]. This solver is, in turn, based upon the EFC constraint solver [9].

The dual of the LP formulation allows us to establish *shadow prices* on items that give an upper bound on the price of each item. We have not implemented such pruning using dual solutions, however. We found that such efficiency improvements were not necessary to find WSSs of optimal revenue for auctions with 1000 bids. The dual problem is:

$$min \sum_{i=1}^{m} y_i$$

$$s.t \sum_{i \in S_j} y_i \geq p_j, \; \forall j \in \{1 \ldots n\}, y_i \geq 0, \;\; y_i \in \mathbb{R}.$$

Other techniques such as the use of cutting planes [92, 141] may further reduce search times. These are constraints that do not affect the solution of the integer program but do constrain the LP polytope. The extent to which cuts, if any, are worth adding requires experimentation. It depends not only on the problem, but also on the specific instance at hand. The more cuts that are inserted, the fewer nodes required in the search takes due to enhanced upper bounding. However, a trade-off occurs because more time is spent at each node. Again, we found that the LP relaxation was sufficient for our needs so we did not require the use of cutting planes.

Combinatorial auctions are easily modeled as a constraint optimization problems. We have chosen the branch-on-bids formulation because in tests it worked faster than a branch-on-items formulation for the `arbitrary` and `regions` distributions. All variables are binary and our search mechanism uses a reverse-lexicographic, value ordering heuristic. This complements our dynamic variable ordering heuristic that selects the most promising unassigned variable as the next one in the search tree. We use the product of the solution of the LP relaxation and the degree of a variable to determine the likelihood of its participation in a robust solution. High values in the LP solution are a strong indication of variables most likely to form a high revenue solution whilst a variable's degree reflects the number of other bids that overlap in terms of desired items. Bids for large numbers of items tend to be more robust, which is why we weight our robust solution search in this manner. We found this heuristic to be slightly more effective than the LP solution alone. As the number of bids in the auction increases, however, there is an increase in the inherent robustness of solutions so the degree of a variable loses significance as the auction size increases.

6.3.3 Results

Our experiments simulate three different levels of solution reparability tightness. The first is that no winning bids are withdrawn by the bid-taker and a repair solution must return a revenue of at least 90% of the optimal overall solution. Secondly, we relaxed the revenue constraint to 85% of optimum. Thirdly, we allowed backtracking by the bid-taker on winning bids using mutual bid bonds but maintaining the revenue constraint at 90% of optimum. We let κ be any percentage (greater than 0) to enable backtracking. Note that the value of κ does not affect solution reparability because penalties by withdrawing bidders are directly proportional to compensatory payments made by the bid-taker for item revocation.

Prior to finding a robust solution we solved the WDP optimally using the `lp_solve` ILP solver [11]. We then set the minimum tolerable revenue for a solution to be 90% (then 85%) of the revenue of this optimal solution. We assumed that all bids were brittle, thus a repair solution is required for every bid in the solution. Initially, we assumed that no backtracking was permitted on assignments of items to other winning bids given a withdrawal elsewhere in the solution.

Optimal Solutions that are Inherently Robust

Table 6.1 shows the percentage of optimal solutions that are robust for minimum revenue constraints for repair solutions of 90% and 85% of optimal revenue. Relaxing the revenue constraint on repair solutions to 85% of the optimum revenue greatly increases the number of optimal solutions that are robust. We also conducted experiments on the same auctions in which backtracking by the bid-taker is permitted using mutual bid bonds. This significantly improves the reparability of optimal solutions whilst still maintaining repair solutions whose revenue is at least 90% of optimum. An interesting feature of the `arbitrary` distribution is that optimal solutions can become more brittle as the number of bids increases. The reason for this is that optimal solutions for larger auctions have more winning bids. Some of the optimal solutions for the smallest auctions with 100 bids have only one winning bidder. If this bid is withdrawn it is usually easy to find a new repair solution within 90% of the previous optimal revenue. Also, repair solutions for bids that contain a small number of items may be made difficult by the fact that a reduced number of bids cover only a subset of those items. A mitigating factor is that such bids form a smaller percentage of the revenue of the optimal solution on average.

We also implemented a rule stipulating that any losing bids from a withdraw-

Table 6.1: Optimal solutions that are inherently robust (%).

	#Bids				
Min Revenue	100	250	500	1000	2000
`arbitrary`					
repair $\geq 90\%$	21	5	3	37	93
repair $\geq 85\%$	26	15	40	87	100
MBB & repair $\geq 90\%$	41	35	60	94	≥ 93
`regions`					
repair $\geq 90\%$	30	33	61	91	98
repair $\geq 85\%$	50	71	95	100	100
MBB & repair $\geq 90\%$	60	78	96	99	≥ 98

ing bidder cannot participate in a repair solution. This acts as a disincentive for strategic withdrawal and was also used previously in the sensitivity analysis.

Occurrence of Robust Solutions

In some auctions, a robust solution may not exist. Table 6.2 shows the percentage of auctions that support robust solutions for the `arbitrary` and `regions` distributions. It is clear that finding robust solutions for the former distribution is particularly difficult for auctions with 250 and 500 bids when revenue constraints are 90% of optimum. This difficulty was previously alluded to by the low percentage of optimal solutions that were robust for these auctions. Relaxing the revenue constraint helps increase the percentage of auctions in which robust solutions are achievable to 88% and 94%, respectively. This improves the reparability of all solutions thereby increasing the average revenue of the optimal robust solution. It is somewhat counter-intuitive to expect a reduction in reparability of auction solutions as the number of bids increases because there tends to be an increased number of solutions above a revenue threshold in larger auctions. The MBB auction model performs very well however, and ensures that robust solutions are achievable for such inherently brittle auctions without sacrificing over 10% of optimal revenue to achieve repair solutions.

Table 6.2: Occurrence of robust solutions (%).

	#Bids			
Min Revenue	100	250	500	1000
`arbitrary`				
repair $\geq 90\%$	58	39	51	98
repair $\geq 85\%$	86	88	94	99
MBB & repair $\geq 90\%$	78	86	98	100
`regions`				
repair $\geq 90\%$	61	70	97	100
repair $\geq 85\%$	89	99	99	100
MBB & repair $\geq 90\%$	83	96	100	100

Revenue of Robust Solutions

Figure 6.4 shows the average revenue of the optimal robust solution as a percentage of the overall optimum. Repair solutions found for a WSS provide a lower bound on possible revenue following a winning-bid withdrawal. Note that in some instances it is possible for a repair solution to have higher revenue than the original solution. When backtracking on winning bids by the bid-taker is disallowed, this can only happen when the repair solution includes two or more bids that were not in the original. Otherwise the repair bids would participate in the optimal robust solution in place of the bid that was withdrawn. A WSS guarantees minimum levels of revenue for repair solutions but this is not to say that repair solutions cannot be improved upon. It is possible to use an incremental algorithm to determine an optimal repair solution following a break, whilst safe in the knowledge that in advance of any possible bid withdrawal we can establish a lower bound on the revenue of a repair. Kastner *et al.* have provided such an incremental ILP formulation [70]. The reader is referred to § 3.1.1 for further details.

Mutual bid bonds facilitate backtracking by the bid-taker on already assigned items. This improves the reparability of all possible solutions thus increasing the revenue of the optimal robust solution on average. Figure 6.4 shows the increase in revenue of robust solutions in such instances. The revenues of repair solutions are bounded by at least 90% of the optimum in our experiments, thereby, allowing

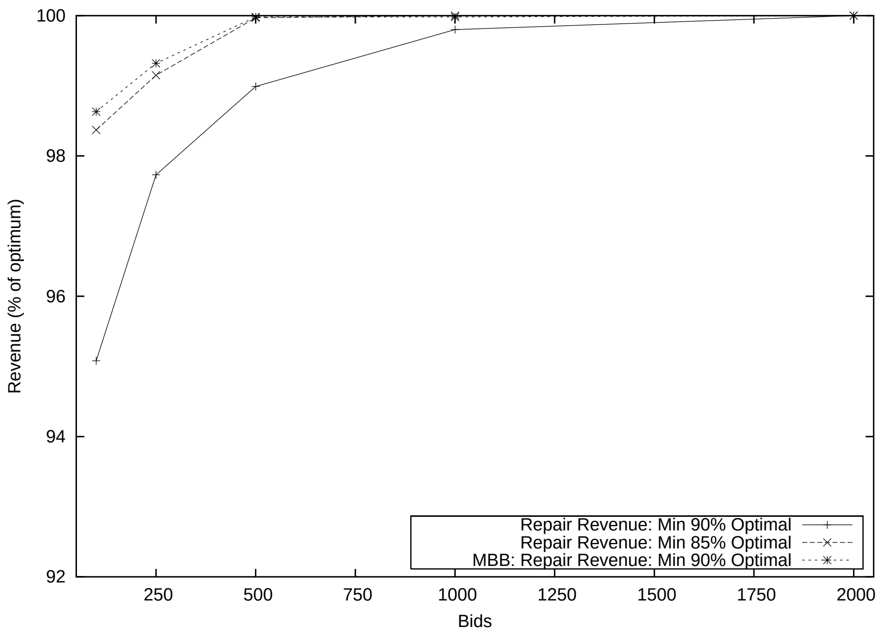

(a) `regions`

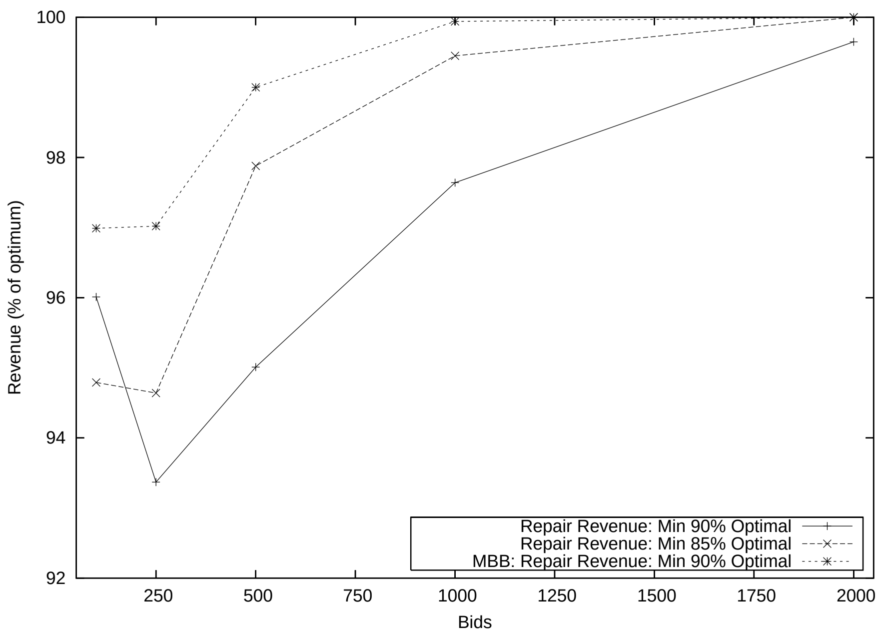

(b) `arbitrary`

Figure 6.4: Revenue of optimal robust solutions.

Table 6.3: Number of winning bids.

Solution	#Bids 100	250	500	1000	2000
`arbitrary`					
Optimal	3.31	5.60	7.17	9.31	10.63
Repair $\geq 90\%$	1.40	2.18	6.10	9.03	(≈ 10.63)
Repair $\geq 85\%$	1.65	3.81	6.78	9.31	(10.63)
MBB ($\geq 90\%$)	2.33	5.49	7.33	9.34	(≈ 10.63)
`regions`					
Optimal	4.34	7.05	9.10	10.67	12.76
Repair $\geq 90\%$	3.03	5.76	8.67	10.63	(≈ 12.76)
Repair $\geq 85\%$	3.45	6.75	9.07	(10.67)	(12.76)
MBB ($\geq 90\%$)	3.90	6.86	9.10	10.68	(≈ 12.76)

a direct comparison with robust solutions already found using the same revenue constraint but not providing for backtracking. It is immediately obvious that such a mechanism can significantly increase revenue whilst still maintaining solution robustness.

Solution Characteristics

Table 6.3 shows the number of winning bids participating in optimal and optimal robust solutions given the three different constraints on repairing solutions listed at the beginning of this section. As the number of bids increases, more of the optimal overall solutions are robust. This leads to a convergence in the number of winning bids. The numbers in brackets are derived from the sensitivity analysis of optimal solutions that reveals the fact that almost all optimal solutions for auctions of 2000 bids are robust. We can, therefore, infer that the average number of winning bids in revenue-maximizing robust solutions converges towards that of the optimal overall solutions.

A notable side-effect of robust solutions is that fewer bids participate in the solutions. It can be clearly seen from Table 6.3 that when revenue constraints on repair solutions are tight, there are fewer winning bids in the optimal robust solution on average. This is particularly pronounced for smaller auctions in both distributions. Although MBBs aid solution reparability, the number of bids in the solutions increases on average. This is to be expected because a greater fraction of these solutions are in fact optimal, as we saw in Table 6.1. Fewer winning bidders leads to benefits for the bid-taker such as reduced overheads in dealing with fewer suppliers. In procurement scenarios, by decreasing the supply base, a company can also leverage its spend with fewer suppliers to attain better overall prices. A smaller number of winning bidders also facilitates improved relationship management with key suppliers [24, 124]. Giunipero *et al*. [53] state the following, "*Being truly agile and adaptable means that we need fewer suppliers. The key is to select the right ones. Globally, we must reduce the number of relationships to manage.*" Robust solutions, therefore, provide a significant benefit in terms of improved relationship management as well as risk management.

In the case where no robust solutions are found, it is possible to optimize reparability, instead of revenue, by finding a solution of at least a given revenue that minimizes the probability of an irreparable break. In this manner the *least brittle* solution of adequate revenue may be chosen, as described in § 5.4.3.

6.4 Risk Management

Following determination of a robust solution, a comparison of the utility of the optimal solution and the WSS is possible using the von Neumann-Morgenstern expected utility property [127], as presented in § 2.4. The risk attitude of the bid-taker is instrumental in determining the chosen solution that maximizes expected revenue. Should the WSS be preferable to S_o, the optimal solution to the WDP, it is also possible to ascertain a value that reflects the premium a bid-taker is willing to pay to secure a robust solution in place of the optimal one.

6.4.1 Utility Comparison

A risk averse bid-taker has a concave utility function. For example, this function, $u()$, may take the form of $u(w) = w^{1-\epsilon}$ where $\epsilon \in [0, 1)$. A risk neutral bid-taker is represented with $\epsilon = 0$ and low values of ϵ signify a modest degree of risk aversion. As $\epsilon \to 1$ the bid-taker becomes more risk averse and, therefore, more likely to prefer a WSS over a brittle optimal solution. Example 6.4.1 demonstrates the choice facing a bid-taker after finding a WSS and performing a sensitivity analysis on that solution so that it can be compared against the optimal solution and the expected utility can be maximized.

Example 6.4.1. *Consider an optimal solution generating a revenue of 1000 and repair solutions of revenue* $\{700, 800, 900\}$*, denoted* $\langle 1000, \{700, 800, 900\}\rangle$*. A WSS has also been found for single bid withdrawals,* $\langle 970, \{940\}\rangle$*, and the bid-taker must choose between the two solutions. We let the Bernoulli utility function of the bid-taker be* $u(w) = (w^{1-\epsilon})$*, where* $\epsilon \in [0, 1)$*,and the independent probability of each single winning-bid withdrawal is 0.05. Recall that we bound the number of possible withdrawals so that there is a polynomial number of possible outcomes in order to maximize expected utility. Therefore, the expected utilities of the respective solutions using Equation 2.1 are:*

$$u_{opt}(w) = (0.85 \times 1000^{1-\epsilon}) + (0.05 \times 700^{1-\epsilon}) + (0.05 \times 800^{1-\epsilon}) + (0.05 \times 900^{1-\epsilon}),$$

$$u_{WSS}(w) = (0.95 \times 970^{1-\epsilon}) + (0.05 \times 940^{1-\epsilon}).$$

The level of risk aversion of the bid-taker, hence ϵ*, determines the utility maximizing option. In the case of a risk neutral bid-taker,* $\epsilon = 0$*, the expected utility of the optimal solution* $u_{opt}(w) = 970$*, whereas the utility of the WSS* $u_{WSS}(w) = 968.5$*. However, as the bid-taker becomes more risk averse, the more robust WSS becomes more attractive. In this instance we find that when* $\epsilon \gtrsim 0.43$*, the bid-taker is sufficiently risk averse to prefer the WSS over the optimal solution.* $\triangle$

Sometimes the bid-taker is uncertain of their precise Bernoulli utility function, in which case some ambiguity may occur regarding the utility maximizing solution.

6.4.2 Risk Premium Calculation

The risk premium is the amount that a risk averse agent is willing to forego in order to obtain the certainty equivalent (CE), an amount without any risk. It may

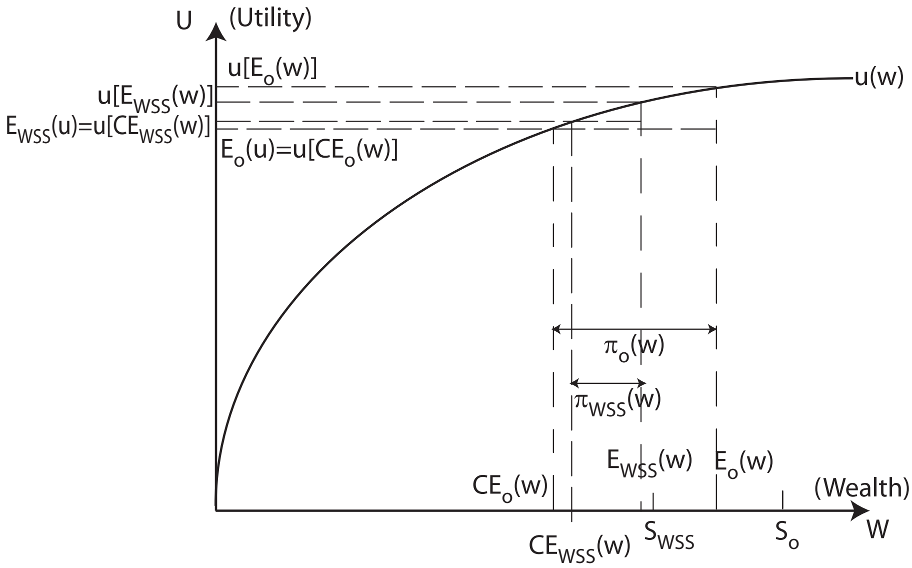

Figure 6.5: Actuarially fair premium determination.

be impossible to achieve a solution without any risk, but a WSS offers a solution with reduced risk. A WSS, therefore, is of value to a bid-taker and the premium in this instance is the added value a bid-taker places on a WSS over the optimal solution.

Figure 6.5 shows a schematic diagram of the solutions presented in Example 6.4.1 for a risk averse bid-taker with $\epsilon > 0.43$. Notice how the CE for the WSS, CE_{WSS}, is greater than the CE for the optimal solution, CE_O. The actuarially fair premium that the bid-taker is willing to pay is given by:

$$AFP = CE_{WSS}(w) - CE_O. \tag{6.1}$$

This premium increases with risk aversion and is a measure of the value of a robust solution over the S_o. When $\epsilon \approx 0.43$, the bid-taker is indifferent to the two solutions but as ϵ increases so too does the value of the WSS over the optimal solution. For example, setting $\epsilon = 0.6$, we get $E_{WSS}(u) = 15.647$ and $E_O(u) = 15.643$. Therefore, $CE_O(w) = e^{\frac{\log 15.643}{1-0.6}} = 967.14$ and $CE_{WSS}(w) = e^{\frac{\log 15.647}{1-0.6}} = 967.76$. The actuarially fair premium in this case is, therefore, 967.76-967.14=0.62. This provides a measure of the surplus utility of the WSS compared to S_o. For example, if a bid-taker in a CA is only capable of determining S_o, then a third party can make a profit by selling a robust solution to this bid-taker. The bid-taker, in this case, is willing to pay up to 0.62 for the WSS. Naturally, this premium increases when the likelihood of bid withdrawal is higher, the repair revenues for the optimal solution are lower, or when the bid-taker is more averse to risk.

6.5 Summary

We successfully applied the WSSs framework to the problem of finding robust solutions for CAs. This is an important result and is the first mechanism for proactive risk management in such auctions when faced with the possibility of bid withdrawal. It allows us to select an allocation that is resistant to bid withdrawals and reflects the fact that involuntary item revocation from bidders is disallowed. We proposed an auction mechanism, mutual bid bonds, to aid reparability by permitting involuntary item withdrawals by the bid-taker in exchange for a compensatory payment. Also, bid withdrawal incurs a penalty payment by the bidder that funds a reallocation of items to establish a repair solution. We showed that this significantly improves reparability and leads to higher revenue robust solutions that are very close, on average, to the optimal WDP revenue.

We also illustrated how such solutions tend to have fewer winning bids than overall optimal solutions, thereby reducing any overheads associated with dealing with more bidders. In procurement scenarios, by decreasing the supply base, a company can also leverage its spend with fewer suppliers to attain better overall prices. A smaller number of winning bidders also facilitates improved relationship management with key suppliers [24, 124]. Giunipero *et al*. [53] state the following, "*Being truly agile and adaptable means that we need fewer suppliers. The key is to select the right ones. Globally, we must reduce the number of relationships to manage.*" Robust solutions, therefore, provide a significant benefit in terms of improved relationship management as well as risk management.

Furthermore, we analyzed the added value of WSSs for risk averse bid-takers and showed how an actuarially fair premium can be calculated to precisely measure this surplus utility over the optimal solution to the WDP.

One possible drawback of robust solutions is that calculation of an optimal bidding strategy is a difficult problem in such a non-incentive compatible mechanism. A truthful mechanism would ease this problem by incentivizing truthful bidding, thereby trivializing the bidding strategy for bidders. This allows bidders to concentrate their efforts upon computing valuations for items more accurately and also facilitates the optimization of a social objective, as described by the Revelation Principle in § 2.2. The following chapter explores truthful mechanism design for robust solutions in detail.

Chapter 7

Truthfulness and Robust Solutions

Thus far, we have developed a means of establishing robust solutions for CAs. This auction mechanism is non-incentive-compatible since bidders may strategize by trying to predict the behavior of others in order to maximize expected utility.

It is possible to incentivize truthful bidding in auctions through careful design of the allocation and payment mechanism. The Revelation Principle, presented in § 2.2, implies that in a wide variety of settings, only "truthful revelation mechanisms" in which agents truthfully announce their types need to be considered when the maximization of a social objective is required. In other words, there are no manipulable mechanisms, in which agents strategically report their types, that attain superior outcomes than any non-manipulable mechanism. This principle allows us to focus our attention upon truthful mechanisms only, when the objective is to optimize some social objective.

There has been considerable research effort in this field by the computer science community in recent years because of the computational challenges posed by incentive compatibility [93, 94]. In this chapter we attempt to develop a truthful mechanism that attains a robust allocation.

7.1 Advantages and Disadvantages of Truthful Mechanisms

A truthful bidding mechanism that is consistent with finding robust solutions is beneficial because it relieves bidders from the computational burden of deciding upon an optimal bidding strategy. Non-truthful auction mechanisms imply that bidders need to strategize about other agents' valuations and behaviors when trying to determine how they should bid in order to maximize their expected surplus. This is a difficult computational problem for bidders in combinatorial auctions, and complicated further by non-optimal winner determination that is used when finding robust solutions. When truthful bidding is the dominant strategy the bidder's problem of determining an optimal bidding strategy becomes trivial. This allows bidders to focus all their computational resources upon estimating their true valuations for sets of items. A truthful mechanism for robust solutions, therefore, facilitates more accurate valuations.

Bidder behavior is also more predictable because truth-telling is the dominant strategy for rational bidders, and this is particularly desirable for a risk averse bid-taker. A possible disadvantage is that bidders reveal their true costs and this is information which they may not wish to reveal. Consequently, there has been considerable research in the field of minimal preference elicitation [26, 27, 65].

The design of truthful mechanisms for a computationally hard problem such as combinatorial auctions is made difficult by the fact that efficiently computable heuristics for winner determination cannot be employed in VCG-based (see § 2.2.2) mechanisms whilst still maintaining their non-manipulability characteristics. Recall from Chapter 2 that VCG mechanisms [25, 55, 126] guarantee that each agent's dominant strategy is to tell the truth, but it requires solving $m + 1$ optimization problems, where the overall optimal solutions involves m agents, and non-optimal solutions compromise truthfulness [93]. Various polynomial-time approximation algorithms can provide good or near optimal solutions very quickly. However, Nisan and Ronen [94] showed, constructively, that a non-optimal solution can in fact result in payments that can be arbitrarily far from the optimum. If an auctioneer seeks to approximate optimal solutions in a GVA using polynomial-time algorithms the results may not be reliable and agents may have an incentive to lie.

This negative result encourages examination of ways in which the problem may be restricted in some respect so that truthfulness can be achieved in a computationally feasible manner, using a polynomial approximation algorithm.

7.2 Truthful Approximation Algorithms

A seminal positive result, due to Lehmann *et al.* [77], showed that by restricting the set of preferences of agents to what they termed *single-minded*, i.e. agents are only interested in a single bundle of items, it is possible to develop greedy mechanisms that are both truthful and computationally efficient. They also showed that the VCG-based mechanisms remain computationally infeasible for single-minded bidders but that it is possible to create a greedy truthful mechanism, that has provable approximation properties.

There are a number of important requirements that an allocation algorithm must fulfil if it is to be incorporated within a truthful mechanism. Firstly, an allocation algorithm must be monotone [77].

Definition 7.2.1 (Monotone allocation algorithm)**.** *An allocation algorithm is monotone if whenever a bid v_j wins, an increase in the bid to v_j' still wins, assuming all other bids are fixed.*

The construction of monotone, approximate allocation algorithms that are computationally feasible is an important line of research for truthful mechanisms. Lehmann *et al.* [77] showed, axiomatically, that a mechanism is truthful if it fulfils the following requirements:

Allocation monotonicity. See Definition 7.2.1;

Exactness. The bidder is either awarded his set of desired items (and no more) or nothing at all;

Critical value. A winning bidder pays the lowest value he could have declared to win the items;

Payment monotonicity. A bidder never pays less for a superset of items in his bid;

Participation constraints. Losing bidders pay zero.

All of these requirements must be fulfilled by any mechanism that seeks to achieve robust solutions for CAs by incentivizing truthful bidding. The monotonicity of the allocation algorithm and the critical value payment scheme are particularly important constraints because they are more difficult to satisfy. The other requirements are easier to fulfil. For example, the participation constraint is satisfied by normalizing the payment scheme. An agent's payment for a set of items S, $p_j(S)$, is reduced by the payment that would make if they were to receive no items, *i.e.* $p_j^{'} = (p_j(S) - p_j(\emptyset))$. Note that this normalization ensures that for any agent who receives no items, their payment is always zero. Exactness is necessarily satisfied by the non-overlapping constraints of the set packing problem constraints.

The family of algorithms developed by Lehmann *et al.* [77] is based on the greedy allocation scheme described in TRUTH-APPROX-ALLOCATION (Algorithm 8). The criterion for ranking bids must have the property of *bid-monotonicity*, so that decreasing the number of desired items in a bid or increasing its bid amount can only improve its ranking after the sorting phase. For example, a valid criterion may be the average bid amount per item.

Algorithm 8: TRUTH-APPROX-ALLOCATION (from [77])

```
input  : Bids, 𝓑; Criterion, C.
output: Allocation, 𝓦.
begin
    𝓦 ← ∅  // Set of winning bids
    𝓛 ← Rank(𝓑, C)  // Rank bids in decreasing order according to C
    foreach i = 1 ... |𝓛| do
        w ← true // Signal that 𝓑[i] is a winning bid
        foreach j = 1 ... (i − 1) do
            if items(𝓑[j]) ∩ items(𝓑[i]) ≠ ∅ AND 𝓑[j] ∈ 𝓦 then
                w ← false
        if w = true then
            𝓦 ← 𝓦 ∪ {𝓑[i]}
end
```

Algorithm 8 satisfies allocation monotonicity because a bidder can only increase his likelihood of winning by increasing his bid amount or decreasing the number of desired items in a bid. This stems from the bid monotonicity of the criterion, C. This is not sufficient on its own to ensure truthful bidding because a suitable payment scheme is also necessary to satisfy payment-monotonicity and participation constraints.

Algorithm 9, TRUTH-APPROX-PAYMENT, describes a greedy payment scheme proposed by Lehmann *et al.* [77]. Note that $n(\mathcal{L}[j])$ is the first bid following $\mathcal{L}[j]$ that is a losing declaration but would have won if $\mathcal{L}[j]$ had not been declared. The critical payment c is the minimum declaration for bid $\mathcal{L}[j]$ so that it would still be prioritized over $n(\mathcal{L}[j])$ according to criterion $\mathcal{C}$. This "threshold" payment scheme is present in several other truthful mechanisms [5, 6, 46]. We assume that the ties are treated in a consistent manner.

Algorithm 9: TRUTH-APPROX-PAYMENT (from [77])

```
input : Ranked Bids, 𝓛; Winning Bids, 𝓦; Number of bids, n.
output: Payments, P.
begin
    P // Payments
    foreach j = 1 ... n do
        if 𝓛[j] ∉ 𝓦 then
            P[𝓛[j]] ← 0
        else
            // n(𝓛[j]): the next bid that would have won if 𝓛[j] was not
              declared
            n(𝓛[j]) ← min{i|i ∈
            {j + 1, ..., n}, items(𝓛[j]) ∩ items(𝓛[i]) ≠ ∅, ∀l < i, l ≠
            j, 𝓛[l] ∈ 𝓦 ⇒ items(𝓛[l]) ∩ items(𝓛[i]) = ∅}
            // c is the critical payment value
            if n(𝓛[j]) ≠ ∅ then
                P[𝓛[j]] ← c
            else
                P[𝓛[j]] ← 0
end
```

Mu'alem and Nisan extended this work by further restricting the agents to *known single-minded bidders* to obtain a richer class of algorithmically efficient truthful mechanisms [88]. They also presented a set of general tools that permit the creation of such mechanisms. We shall adopt their known single-minded bidders model and explore truthfulness for robust solutions.

7.3 A Truthful Mechanism for Robust Solutions

We present a restricted model for combinatorial auctions that permits truthful mechanisms for CAs, thus enabling maximization of social welfare. We investigate the possibility of embedding the WSSs framework as an allocation algorithm within a truthful mechanism. We discuss the difficulties encountered with this approach and, finally, we then present an important positive result: a *truthful approximation scheme that approximately maximizes expected utility* for single bid withdrawals. Recall from § 3.2.2 that we formulated an integer program that models such a utility maximizing approach.

7.3.1 The Model

We adopt a similar model to [88] to investigate possible mechanisms that incentivize truthful bidding over a set U of m items. Each bidder j, of n bidders in total, has a positive valuation function v_j for a subset $S \subseteq U$ of items. We assume that agents' valuation models are private, as is commonplace within mechanism design, *e.g.* [4, 104]. This removes the possibility of the winner's curse effect altering ex-post bidder valuations, thus eliminating the possibility of bid withdrawals for strategic purposes. However, this restriction also reduces the likelihood of bid withdrawal, thereby inhibiting the possible benefits of robust solutions to situations with exogenous probabilities of bid withdrawal.

Definition 7.3.1 (Single-minded bidder [77]). *Bidder j is single-minded if there is a set of goods $S_j \subseteq U$ and a value $v_j^* \geq 0$ such that* $v_j(S) = \begin{cases} v_j^* & S \supseteq S_j \\ 0 & otherwise. \end{cases}$

Informally, single-minded bidders are willing to pay v_j^*, a privately known valuation, on condition that they minimally receive a set of items, S_j. This condition is necessary to satisfy the individual rationality constraints presented in § 2.2. Mu'alem and Nisan [88] introduced known single-minded bidders, whereby the subsets S_j are known to the mechanism. This mechanism is composed of an allocation algorithm, $A(v)$, whose inputs are the bid declarations $\{v_1 \ldots v_n\}$, the respective desired items $\{S_1 \ldots S_n\}$ and a payment rule $p(v)$. The output of $A(v)$ is a subset of pair-wise disjoint winning bids.

In summary, our model adopts the following assumptions/restrictions to help us attain a truthful mechanism:

Private values. Bidders are indifferent to others' valuations. There is no winner's curse so bidders do not discount bid amounts;

Exogenous probabilities of withdrawal. Withdrawals are caused by random events. Bidders have no control over the reasons for bid withdrawal;

Known single-minded bidders. This is a restriction proposed in [88] to allow polynomial algorithms incentivize truthful bidding;

Rational bidders. This assumption is common to most mechanism design literature [82, 93]. If bidders act irrationally there is no means of incentivizing any form of desired behavior.

7.3.2 Monotonicity and the WSS Algorithm

If we can achieve an approximation scheme that is monotone, we can also achieve a payment scheme that incentivizes truthful bidding [4, 88, 109]. We investigate the monotonicity of the WEIGHTED-SUPER-SOLVE algorithm (see § 5.4.2). We find that an increase in a bid amount may introduce a repair for a previously irreparable bid, thereby enlarging the set of bids that may be considered when determining an optimal robust solution. We shall see how this can occur in practice in Example 7.3.1.

Lemma 7.3.1 provides an impossibility result regarding the monotonicity of WSSs when finding an optimal 1-robust solution (Definition 6.1.2, page 99). The non-existence of a monotone allocation algorithm implies that for some bidder a, if v_a is a winning declaration then for some higher declaration $v_a^{'} \geq v_a$ it becomes a losing declaration. The proof utilizes the fact that if v_a becomes a losing bid, then it must form part of a repair for another previously irreparable bid that partakes in the subsequent winning allocation.

Lemma 7.3.1. *The WSS algorithm for finding an optimal 1-robust solution to a CA with irrevocable assignments of items is non-monotone.*

Proof. We prove by construction as follows: consider three feasible allocations, A_1 A_2 and A_3 with revenues of $rev(A_1)$, $rev(A_2)$ and $rev(A_3)$ respectively. We let $rev(A_1) > rev(A_2) \geq R > rev(A_3)$ with R being a lower bound on acceptable revenue. Let these allocations contain the following winning bids $A_1 = \{v_x\} \cup L$, $A_2 = \{v_y\} \cup M$ and $A_3 = \{v_y\} \cup L \cup N$, where L, M and N are mutually exclusive sets of bids. We can say the following:

$$rev(\{v_x\} \cup L) > rev(\{v_y\} \cup M) \geq R > rev(\{v_y\} \cup L \cup N).$$

We let v_x be the only brittle bid. Its withdrawal from A_1 has an optimal repair allocation of A_3, but $R > rev(A_3)$. Therefore, A_1 is an irreparable allocation because v_x cannot be repaired, so A_2 is the optimal WSS and v_y is a winning bid.

Let an increase in v_y to $v_y^{'} \Rightarrow rev(A_3^{'}) \geq R$. The allocation A_1 becomes reparable following the increase in v_y to $v_y^{'}$ but should not become the winning allocation if the WSS algorithm were monotone. However, by choosing v_x so that

$$rev(\{v_y\} \cup M) > rev(\{v_x\}) > (rev(\{v_y^{'}\} \cup M) - rev(L)), \qquad (7.1)$$

we see that $A_2^{'}$ still has a lower revenue than A_1 following the increased declaration. This means that A_1 becomes the new winning allocation. The condition, $rev(\{v_y\} \cup M) > rev(\{v_x\})$, is necessary so that v_x is not an optimal WSS on its own before the bid increase.

Therefore, the WSS allocation algorithm is non-monotone because v_y becomes a losing declaration following an increase to $v_y^{'}$. Note that as a consequence of the assumption in Equation 7.1, $rev(\{v_y^{'}\}) - rev(\{v_y\}) < rev(L)$. Otherwise, the increase would be so large that v_y would remain a winning bid, although the set of other winning bids would change. □

Lemma 7.3.1 demonstrated that increasing a winning bid has the possible side-effect of creating new robust solutions, by forming a repair solution for a previously irreparable allocation of higher revenue. Example 7.3.1 provides a concrete illustration of the non-monotonicity of the WSS algorithm by showing how a winning declaration can become a losing declaration following an increase in its bid amount.

Example 7.3.1. *Consider bids $v_1\langle 13, A\rangle$, $v_2\langle 12, B\rangle$, $v_3\langle 8, C\rangle$, $v_4\langle 33, AB\rangle$ and $v_5\langle 24, BC\rangle$. All bids are robust,* i.e. *not likely to be withdrawn, except for v_4. The minimum acceptable revenue, R, for a solution is 34. The optimal robust solution is, therefore, $\{v_1, v_5\}$ with revenue of 37, because $\{v_3, v_4\}$ cannot be repaired to form a solution with revenue of at least 34 should v_4 be withdrawn. The optimal repair solution, $\{v_1, v_2, v_3\}$, attains a revenue of only 33.*

Let us suppose that v_1 became $v_1^{'}\langle 16, A\rangle$, so that the declaration has increased by 3. A monotone allocation algorithm requires that v_1 remains a winning bid, see Definition 7.2.1. However, allocation $\{v_3, v_4\}$ can now be repaired to form a solution with revenue of 36 should v_4 be withdrawn, therefore, the new optimal WSS becomes $\{v_3, v_4\}$. This implies that $v_1^{'}$ becomes a losing declaration following the increase from v_1, thus contravening the allocation monotonicity requirement. △

Theorem 7.3.1. *A normalized truthful mechanism is impossible for 1-robust solutions with known single-minded bidders where item assignments are irrevocable, when using the WSS framework.*

Proof. A normalized mechanism is truthful if and only if its allocation algorithm is monotone and its payment scheme is based on critical value [88]. Lemma 7.3.1 shows that a monotone allocation algorithm is impossible to achieve for 1-robust solutions where item assignments are irrevocable. □

Theorem 7.3.1 shows that it is impossible to incentivize truthful bidding to achieve robust solutions in the restricted case of combinatorial auctions with known single-minded bidders. Monotonicity is hindered by the fact that irreparable bids are discounted from consideration when determining the winners.

The lower bound on tolerable revenue was introduced to permit satisfiability testing for repairs in order to improve computational feasibility. However, this constraint is acting as a barrier to truthful mechanisms so we need to circumvent this constraint in order to find a truthful mechanism for robust solutions.

The WSS algorithm is $\mathcal{NP}$-complete, when a fixed bound is placed on the size of any break, so it is not feasible for it to be applied in a VCG-based mechanism. Recall that such a mechanism would need to be solved optimally many times for slightly different instances to determine the payments. Lehmann *et al.* showed the infeasibility of the approach, even when bidders are single-minded [77].

7.3.3 Robust Solutions Approximation Algorithm

Recall from § 3.2.2 that the integer programming formulation for determination of maximum expected revenue given the possibility of single bid-withdrawal, is the following:

$$\max \left(\sum_{j=1}^{n} (1 - p_j) u(v_j) x_j + \sum_{k=1}^{n} \sum_{j=1}^{n} p_k x_k u(v_j) r_{kj} \right) \tag{7.2}$$

subject to the following constraints:

$$\sum_{j|i \in S_j}^{m} x_j \leq 1, \forall i \in \{1, \ldots, m\},\ x_j \in \{0, 1\}, \tag{7.3a}$$

$$\sum_{j|i \in S_j, k \neq j}^{m} x_j + r_{kj} \leq 1, \tag{7.3b}$$

$$\forall i \in \{1, \ldots, m\}, \forall k \in \{1, \ldots, n\},\ x_j, r_{kj(k \neq j)} \in \{0, 1\}, r_{kk} \in \{0\},$$

where p_j is the probability of bid j being withdrawn, x_j the decision variable for bid j, v_j the amount of bid j and r_{kj} is the repair value of bid j given a withdrawal by bid k. We assume conditional probabilities of failure so that the likelihood of multiple bid withdrawals is zero. Notwithstanding this assumption, the formulation still remains computationally infeasible because of the size of the input, $\mathcal{O}(n^2)$ variables. A branch and bound search to find an optimal solution quickly becomes prohibitively expensive in the size of n. The LP relaxation for Equation 7.2 relaxes the integrality constraints on x_j and r_{kj} so that bids, in effect, can partially win ($y_j, r_{kj} \in [0, 1]$). Although this problem also has $O(n^2)$ input variables, the problem can be solved in polynomial time. The LP formulation, denoted $LP(v)$, is as follows:

$$\max\left(\sum_{j=1}^{n}(1-p_j)u(v_j)y_j + \sum_{k=1}^{n}\sum_{j=1}^{n}p_k y_k u(v_j) r_{kj}\right) \tag{7.4}$$

subject to the following constraints:

$$\sum_{j|i\in S_j}^{m} y_j \leq 1, \forall i \in \{1,\ldots,m\},\ y_j \in [0,1], \tag{7.5a}$$

$$\sum_{j|i\in S_j,k\neq j}^{m} y_j + r_{kj} \leq 1, \tag{7.5b}$$

$$\forall i \in \{1,\ldots,m\}, \forall k \in \{1,\ldots,n\},\ y_j, r_{kj(k\neq j)} \in [0,1], r_{kk} \in \{0\},$$

Notice that there are no minimum revenue constraints in this model. These were originally introduced in Chapter 4 as a compromise on the original problem of utility maximization so that we could improve computational feasibility whilst still maintaining the integrality constraints on winning bids. By removing these constraints and optimizing expected utility using the LP relaxation, we need to then determine the allocation from the results of this non-integer solution. It is natural to assume that y_j, representing the status of bid j in the optimal allocation, would provide a good heuristic for guiding the choice of an approximate solution because it reflects a fraction of how much that bid should win in an optimal solution. In general, however, such heuristics are not necessarily truthful. We develop an approximation algorithm similar to the LP-BASED allocation algorithm of Mu'alem and Nisan [88] and show that it is monotone.

Informally, LP-BASED ROBUST (Algorithm 10) assigns any bid whose value in the corresponding LP solution is greater than $\frac{1}{2}$ to be a winning bid. When $y_j > \frac{1}{2}$, the bid gets a larger fraction of any of its desired items than any other conflicting bid, y_c. We can infer that $y_c < \frac{1}{2}$ and is, therefore, a losing bid. This ensures that the final solution is consistent with our resource constraints.

Recall from § 7.3.1 that each bidder j, of n bidders in total, has a positive valuation function v_j for a subset of items. In order to prove the monotonicity of Algorithm 10 it is necessary to show that for any v_{-j}, where $v_{-j} = \{v \setminus \{v_j\}\}$, y_j

Algorithm 10: LP-BASED ROBUST

```
input  : Bids V
output: Allocation X
begin
    Y ← LP(V)  // Optimal LP solution
    for j ← 1 ... n do
        // Y[j] ∈ [0,1] indicates fractional bid success.
        if Y[j] > 1/2 then
            X[j] ← 1 // X[j] is a binary variable indicating bid success.
        else
            X[j] ← 0
end
```

is a non-decreasing function of v_j.

Lemma 7.3.2. *y_j is a non-decreasing function of v_j for any fixed v_{-j}.*

Proof. The proof follows that of the LP-BASED algorithm monotonicity in [88]. Consider optimal feasible solutions $\mathcal{Y}$ and $\mathcal{Y}'$ to the linear program of Equation 7.4, $LP(v)$ and $LP(v')$, respectively, where $v_j \leq v_j'$ and $\Delta = v_j' - v_j$. $\mathcal{Y}'$ is a feasible, though non-optimal, solution to $LP(v)$ and so $\sum_{i=1}^{n} y_i' v_i \leq \sum_{i=1}^{n} y_i v_i$. $\mathcal{Y}$ is also a feasible solution to $LP(v')$ and so $y_j\Delta + \sum_{i=1}^{n} y_i v_i \leq y_j'\Delta + \sum_{i=1}^{n} y_i' v_i$. Therefore, $0 \leq \sum_{i}^{n} (y_i - y_i')v_i \leq (y_i' - y_i)\Delta$, hence $y_j \leq y_j'$, thus proving that y_j is a non-decreasing function of v_j. □

Lemma 7.3.3. *Algorithm 10 (*LP-BASED ROBUST*) is monotone.*

Proof. v_j is a winning declaration if and only if $\frac{1}{2} < y_j$. From Lemma 7.3.2 we know that $\frac{1}{2} < y_j \leq y_j'$ so any $v_j' > v_j$ is also a winning declaration. □

Theorem 7.3.2. *A normalized mechanism that assigns items to bidders according to Algorithm 10 (*LP-BASED ROBUST*) is truthful when its payment scheme is based upon the critical value.*

Proof. We know from [88] that for any v_{-j} and monotone allocation algorithm there exists a single critical value below which v_j is always a losing declaration and above which it is always winning. A normalized mechanism is truthful when its payment scheme is based upon the critical value, and the allocation algorithm is monotone [88]. From Lemma 7.3.3 we know that LP-BASED ROBUST is monotone. □

Theorem 7.3.2 provides a significant positive result. It implies that truthful bidding can be incentivized when 1-robust solutions that can withstand exogenous probabilities of withdrawal are desired by the bid-taker.

7.3.4 Non-monotonicity of the WSS Algorithm with MBB's

Previously, in § 7.3.2, we saw how the WSS algorithm for 1-robust solutions for CAs is non-monotone. In this Section we show that the WSS allocation algorithm remains non-monotone when mutual bid bonds are introduced to aid reparability.

Lemma 7.3.4. *The WSS algorithm for finding an optimal 1-robust solution to a CA with mutual bid bonds is non-monotone.*

Proof. We prove by construction in a similar manner to that of Lemma 7.3.1. Consider once again three feasible allocations, A_1 A_2 and A_3 with revenues of $rev(A_1)$, $rev(A_2)$ and $rev(A_3)$ respectively. We let $rev(A_1) > rev(A_2) \geq R > rev(A_3)$ with R being a lower bound on acceptable revenue. Let these allocations contain the following winning bids $A_1 = \{v_x\} \cup L$, $A_2 = \{v_y\} \cup M$ and $A_3 = \{v_y\} \cup L \cup N$, where L, M and N are mutually exclusive sets of bids. We let v_x be the only brittle bid and its withdrawal from A_1 has an optimal repair allocation of A_3, but $R > rev(A_3)$. Therefore, A_1 is an irreparable allocation because v_x cannot be repaired.

Let an increase in v_y to $v_y^{'} \Rightarrow rev(A_3^{'}) \geq R$. The allocation A_1 becomes reparable following the increase in v_y to $v_y^{'}$ but should not become the winning allocation if the WSS algorithm were monotone. However, by choosing v_x so that

$$rev(\{v_y\} \cup M) > rev(\{v_x\}) > (rev(\{v_y^{'}\} \cup M) - rev(L)),$$

we see that $A_2^{'}$ still has a lower revenue than A_1 following the increased declaration. Because MBBs are in place, we require an extra restriction so that $rev(\{v_y\}) < rev(\{v_i\}), \forall v_i \in M$, thus ensuring that no winning bids can have their items revoked to form an alternative repair solution. This can be achieved by letting $rev(\{v_y\}) < rev(M)$, where M contains only one bid. This means that A_1 becomes the new winning allocation following the increase in v_y to $v_y^{'}$. Therefore, the WSS allocation algorithm is non-monotone, even when MBBs are considered. □

We present an example to highlight how an increase in a winning declaration can result in it becoming a losing declaration because it forms part of a repair solution for a previously irreparable solution. Note that Example 7.3.2 has different bid amounts for v_3 and v_5 than in Example 7.3.1. These were altered so that the withdrawal of v_4 would provide insufficient funds to repair $\langle v_3, v_4 \rangle$ by revoking items from v_3. Mutual bid bonds only permit revocation of items from a set of winning bids whose summation of declarations is less than that of the withdrawn bid.

Example 7.3.2. *Consider bids* $v_1\langle 13, A\rangle$, $v_2\langle 12, B\rangle$, $v_3\langle 38, C\rangle$, $v_4\langle 33, AB\rangle$ *and* $v_5\langle 54, BC\rangle$. *All bids are robust,* i.e. *not likely to be withdrawn, except for* v_4. *The minimum acceptable revenue for a solution is 64. The bid-taker would like to be able to award the items to* $\{v_3, v_4\}$ *but* v_4 *is a brittle bid that cannot be repaired satisfactorily. The optimal robust solution is, therefore,* $\{v_1, v_5\}$ *with revenue of 67.*

Suppose that v_1 *became* $v_1^{'}\langle 16, A\rangle$, *so that the signal has increased by 3. The mutual bid bond provides a fund of* $\kappa\%\times$ *33 if* v_4 *is withdrawn, but this is insufficient to allow for revocation of items from* v_3 *who requires a compensatory payment of* $\kappa\% \times 38$. *A monotone allocation algorithm requires that* v_1 *remains a winning bid, from Definition 7.2.1. However,* v_4 *can now be repaired to form a solution,* $\langle v_1, v_2, v_3\rangle$, *of revenue 66 should the bid be withdrawn, therefore, the new optimal WSS following the increase in* v_1 *is* $\{v_3, v_4\}$. *This implies that* $v_1^{'}$ *becomes a losing declaration following the increase from* v_1, *thus contravening the allocation monotonicity requirement.* △

Theorem 7.3.3. *A normalized truthful mechanism is impossible for 1-robust solutions with known single-minded bidders and mutual bid bonds, when using the WSS algorithm.*

Proof. A normalized mechanism is truthful if and only if its allocation algorithm is monotone and its payment scheme is based on critical value [88]. Lemma 7.3.4 shows that a monotone allocation algorithm is impossible to achieve for 1-robust solutions with mutual bid bonds. □

Theorem 7.3.3 states that a truthful mechanism is impossible using the WSSs framework and mutual bid bonds. Previously we have circumvented this problem by relaxing the constraints on minimum revenue to find a polynomial approximation algorithm that permits truthful bidding. The formulation of a linear program that enables repair solutions that can backtrack upon item assignments in the original solution is not straightforward. We leave the formulation of a truthful mechanism that incorporates mutual bid bonds for future work.

7.4 Summary

Truthful mechanisms pose difficulties from a computational feasibility perspective. A key requirement of these algorithms is that the allocation algorithm is monotone. In this chapter we proved a key negative result, that the WSS algorithm is non-monotone, for the case in which no items can be revoked from winning bidders. We also showed that it is also non-monotone when mutual bid bonds are present.

A restricted form of combinatorial auction whereby bidders are only interested in a particular bundle of items, *i.e.* single-minded, permits the creation of an approximation algorithm that can incentivize truthful bidding [77]. A further restriction so that those bundles are known by the mechanism permits a wider range of truthful approximation algorithms [88].

We succeeded in circumventing the impossibility results associated with the WSSs framework, for the case in which no item retraction by the bid-taker is permitted, by removing the minimal repair solution constraint. When the conditional probability of multiple withdrawals is set to zero we can use an LP-based approximation algorithm to find approximately utility-maximizing solutions. Therefore, we proved that it is possible to create a computationally efficient truthful approximation scheme for robust solutions using this allocation algorithm when the payment scheme is based upon the critical value.

Chapter 8

Conclusion and Future Work

In this dissertation, we have demonstrated a significant problem for bid-takers in CAs, illustrated the inadequacies of currently available technologies to tackle this problem and proposed a novel solution. We also addressed the situation in which the bid-taker wishes to have a robust solution whilst optimizing a social objective via the use of a truthful mechanism. This work also highlighted some interesting directions for possible future work, which are presented in this Chapter.

Most notably, we discuss the implications of solution robustness for bidder aggressiveness and we also present an alternative approach for risk management for the bid-taker's exposure problem. This technique may offer an opportunity for improved revenue in robust solutions.

Risk management for CAs is a rich field of research that has many open questions that deserve further attention because of their practical implications for electronic commerce. We conclude by demonstrating that our work has significantly contributed towards managing risk in CAs and forms a significant initial platform for future work in this field.

8.1 Conclusion

The bid-taker's exposure problem constitutes a serious risk for the bid-taker in combinatorial auctions. We performed the first in-depth analysis of this problem and showed that it can cause large revenue losses for various economically motivated bidding distributions.

The root cause of failure in an auction solution is oftentimes outside the control of the bidder who instigates a bid withdrawal, *e.g.* natural disasters or competition regulator intervention. Therefore, it may not be practicable for the bid-taker to enforce non-withdrawal rules. Auction rules regarding non-withdrawal of bids may constitute a fallacy because of the impracticability of absolute enforcement. Winning-bid withdrawal is a serious threat to solution stability.

A robust solution to a combinatorial auction can withstand winning-bid withdrawal by reassigning items to form a repair solution without causing undue disturbance to other non-reneging winning bidders. Such solutions are particularly desirable for risk averse bid-takers. The problem of finding robust solutions that maximize expected utility is computationally challenging. We proposed a relaxation of the problem so that a lower bound on the revenue of repair solutions is guaranteed. The search for a repair solution thus becomes a satisfiability test rather than an optimization problem. We adopted the super solutions framework to find a robust solution that limits the number of changes required to the solution so that a repair solution can be guaranteed [52, 59]. This approach highlighted some limitations to the expressiveness of the framework and motivated some valuable extensions.

We developed the weighted super solutions framework to enhance the flexibility and the accuracy of robust solutions. This is a framework whereby assignments have either static probabilities of failure or dynamic failure rates over time using the Weibull probability distribution. This distribution can model a wide range of failure types including metal fatigue, corrosion or abrasion. A set of assignments is deemed brittle if their combined probability of failure exceeds a threshold value. The WSSs framework also facilitates an expressive measure for the cost of repair. The reassignment of variables following a break has an associated cost that can depend upon the source and destination assignments as well as the cause of the break. This enables an expressive and versatile approach to finding robust solutions for constraint programs.

We applied the WSSs framework to the problem of finding robust solutions to combinatorial auctions. We analyzed the implications of winning-bid withdrawal in several economically motivated combinatorial auctions and found a large variance in the levels of revenue loss for different bid distributions. We concentrated upon finding robust solutions to those auction types that were most sensitive to solution uncertainty. We showed that robust solutions can attain a high percentage of optimal revenue.

We demonstrated important results concerning a reduction in the number of winning bids in robust solutions. This is a significant attribute in real-world auctions. By decreasing the supply base, a company can leverage its spend with fewer suppliers to attain better overall prices. A smaller number of winning bidders also facilitates improved relationship management with key suppliers [24, 124]. Giunipero *et al*. [53] state the following, "*Being truly agile and adaptable means that we need fewer suppliers. The key is to select the right ones. Globally, we must reduce the number of relationships to manage.*" Robust solutions, therefore, provide a substantial benefit in terms of improved relationship management as well as risk management.

We introduced an auction mechanism for improving solution-reparability called mutual bid bonds. The WSSs framework facilitates backtracking upon agreements and we showed that this can provide robust solutions with near optimal revenues. We also demonstrated how a risk averse bid-taker can choose between the optimal (non-robust) solution and a robust solution using the von-Neumann Morgenstern expected utility property [127] to maximize expected utility.

One disadvantage of robust solutions is the complex nature of the problem faced by bidders in establishing an optimal bidding strategy. This represents a non-incentive-compatible mechanism whereby bidders strategize about the actions of others to identify an equilibrium bidding strategy that optimizes their expected surplus. We proposed a means of addressing this problem by establishing a mechanism that incentivizes truthful bidding as a tradeoff against expected revenue for the bid-taker. This also facilitates the determination of an allocation that maximizes social welfare. Bidders, therefore, need only establish their own true valuations for items and do not need to worry about other agents' behavior. We showed that a truthful mechanism is impossible using the WSSs framework but can instead be achieved by using an LP-based approximation allocation algorithm with payments based upon the critical value.

In summary, solution robustness is an important concept that is often over-looked in auction literature. It is important that there exists sufficient technological support for risk management capabilities to nurture trust in electronic commerce and encourage further deployment of online auctions. This work provides an important means of combatting solution uncertainty in combinatorial auctions, an increasingly popular auction format in B2B commerce.

8.2 Future Work

There are many interesting directions for possible future work that we have categorized in two separate domains, electronic commerce and optimization.

8.2.1 Electronic Commerce

Bidders in auctions can be regarded as selfish rational agents whose behavior can be analyzed using non-cooperative game theory. Robust solutions for combinatorial auctions can complicate the determination of an equilibrium bidding strategy. We consider the implications for bidding aggressiveness when bidders adopt affiliated valuation models whereby the winner's curse can influence bidding behavior.

We mention a further motivation for the use of robust solutions, when the bid-taker wishes to express preferences over suppliers. We also propose the development of a truthful mechanism for robust solutions that incorporate mutual bid bonds as an interesting direction for possible future work.

Bidder Aggressiveness

Recall from § 2.1 that bidders are utility-maximizing and will submit a bid that is lower than their true valuation for an item to increase their surplus utility. There is a trade-off between the decreased probability of winning the item and the increased utility as the bid amount is lowered. Bidders try to estimate other bidders' valuations when determining their bid amount. When allowing for randomized strategies, at least one Nash equilibrium exists in any game with regular payoff functions.

Bidding strategies can become complex in non-incentive-compatible mechanisms where winner determination is no longer necessarily optimal. The robustness of a bid may influence bidding strategy. Bids from trusted/reliable bidders have an advantage over bids that require repair solutions. Non-trusted bidders bid more aggressively (*i.e.* the equilibrium bidding strategy for such bidders is to bid higher) to overcome this hurdle. However, trusted bidders can bid less aggressively because reparability constraints on other bidders can reduce competition.

The perceived reparability of bids deemed to be brittle also influences the equilibrium bidding strategy. If a bidder feels that there are no other bidders interested in the same set or a subset of items in a bid then finding a repair solution may be difficult. Decreasing the bid amount increases the likelihood of finding a repair solution within $\kappa\%$ of optimum revenue, should the bid form part of the optimal WDP allocation. However, it also decreases the likelihood of participation in this allocation.

The possibility of bid withdrawal creates an incentive for more aggressive bidding by providing a form of insurance that caps the potential losses associated with the winner's curse. If a winning bidder's revised valuation for a set of items drops by more than the penalty for withdrawing the bid, then it is in his best interests to forfeit the item(s) and pay the penalty. Should the auction rules state that the bid-taker will refuse to sell the items to any of the remaining bidders in the event of a withdrawal, then insurance against potential losses will stimulate more aggressive bidding [57]. Another factor that increases aggressiveness is the extra utility that can be gained from being compensated by the bid-taker for item revocation following bid withdrawal by another successful bidder.

We seek to repair a solution following a winning-bid withdrawal with the remaining bids. A side-effect of such a policy is to offset the increased aggressiveness by incentivizing reduced valuations in the expectation that a bid is withdrawn. There is also the possibility that a decrease in a bid may cause another bid to become irreparable. If the irreparable bid formed part of an optimal robust allocation then a lower declaration increases the probability of the bid winning, thus providing an incentive for a lower declaration. Previously, in Chapter 7, we saw how this effect causes the non-monotonicity of the WSS algorithm.

Harstad and Rothkopf [57] examined the conditions required to ensure an equilibrium position in which bidding was at least as aggressive as if no bid withdrawal was permitted, given this countervailing incentive to under-estimate a valuation. Three major results arose from their study of single item auctions with bid withdrawal:

1. Equilibrium bidding is more aggressive with sufficiently small probabilities of an award to the second highest bidder in the event of a bid withdrawal;
2. Equilibrium bidding is more aggressive if the number of bidders is large enough;
3. For many distributions of costs and estimates, equilibrium bidding is more aggressive with withdrawal if the variability of the estimating distribution is sufficiently large.

It is important that mutual bid bonds do not result in depressed bidding in equilibrium. An analysis of the resultant behavior of bidders must incorporate the possibility of a bidder winning an item and having it withdrawn in order for the bid-taker to formulate a repair solution after a break elsewhere. Harstad and Rothkopf have analyzed bidder aggressiveness using a strictly game-theoretic model in which the only reason for bid withdrawal is the winner's curse [57]. They assumed all bidders were risk neutral, but concluded that it is entirely possible for the bid-taker to collect a risk premium from risk averse bidders with the offer of such insurance. We leave an in-depth analysis of the sufficient conditions for more aggressive bidding for future work.

Positive Discrimination in Favor of Trusted Bidders

Fairness is often a goal of auctions and its importance is scenario dependent. Concern about equal treatment of competitors (and the appearance of it) is often a key reason for government use of auctions [112]. In an industrial procurement scenario the constraints on fairness may be relaxed because the purchaser can maintain the right to choose winners on a non-price basis. However, in private auctions fairness does influence bidders willingness to participate and should not be completely disregarded.

Fairness is often cited as a reason for choosing the optimal solution solely in terms of revenue [112]. Robust solutions militate against bids deemed to be brittle, therefore, bidders must earn a reputation for being reliable to relax the reparability constraint attached to their bids. This may be regarded as being fair to long-standing business partners whose reliability is unquestioned. Internet-based auctions are often seen as unwelcome price-gouging exercises by suppliers in many sectors [42, 97]. Traditional business partnerships are being severed by increased competition amongst suppliers. Quality of Service can suffer because of the increased focus on short-term profitability to the detriment of the bid-taker in the long-term. Robust solutions can provide a means of selectively discriminating against distrusted bidders in a measured manner. Some possible future work may include an empirical study of how successful this technique may be in auctions when the bid-taker wishes to discriminate with subtlety in favor of certain bidders.

Truthfulness with Mutual Bid Bonds

We presented a truthful mechanism in Chapter 7 that approximates expected utility maximization. In this mechanism it is assumed that the bid-taker cannot withdraw items from bidders after a winning-bid withdrawal. The probabilities of failure were exogenous because the private values model used does not induce the winner's curse and reneging bidders could not be considered when repairing the solution. Bidders were not expected to pay any penalty following withdrawal.

We showed in Chapter 6 that the MBB mechanism improves solution reparability, thus increasing revenue. Reneging bidders pay a penalty following withdrawal. The bid-taker also pays compensation to winning bidders whose items are withdrawn when the solution is repaired following a withdrawal by another bidder. We showed that in Chapter 7 it is possible to create a monotone allocation algorithm using the LP relaxation of the problem that maximizes expected utility when a single bid withdrawal is possible. When the bid-taker can withdraw items from winning bidders using mutual bid bonds, the constraints on reparability are relaxed.

An open question remains whether it is possible to achieve a monotone allocation algorithm for robust solutions so that truthfulness can be achieved with defined approximation ratios whilst remaining flexible enough to handle mutual bid bonds. Archer *et al.* devised a version of randomized rounding that is incentive compatible, giving a truthful mechanism for single-minded bidders [6].

They overcame a non-monotone allocation algorithm by modifying an LP rounding algorithm to ensure monotonicity. An open question remains whether such a monotone allocation algorithm can be developed for robust solutions with MBBs. A *randomized mechanism* can be viewed as a chance selection of a random element that is an input to a deterministic mechanism. The solution concepts for randomized mechanisms have varying degrees of truthfulness with *strongly truthful* implying that the generated deterministic mechanism is always truthful regardless of the randomly generated input. Strong truthfulness is restrictive and *truthfulness in expectation* may be used as a reasonable weakening of the solution concept. Alternatively, *strongly truthful with high probability* has been used by Archer *et al.* as a means of incentivizing truthful bidding for bidders that are not necessarily risk neutral [6]. A reasonable compromise of the solution concept may permit the creation of a truthful mechanism for robust solutions.

8.2.2 Optimization

Robust solutions for combinatorial auctions raise many questions and motivate interesting research questions in the computer science domain. In this Section we present some outstanding issues relating to robustness and optimization that deserve further consideration.

ILP Solution Robustness

Whilst the WSSs framework provides ample flexibility and expressiveness, scalability becomes a problem for larger auctions. Although optimal solutions to large auctions tend to be naturally robust, some bid-takers in such auctions may require improved robustness. Bertsimas *et al.* [13, 14] introduce a formulation of integer programming problems that produces robust solutions when either the cost coefficients or the data constraints may be subject to uncertainty. We are interested in decision variable uncertainty, however, so this approach is not directly applicable.

A possible extension of our work in this dissertation would be to examine the feasibility of reformulating integer linear programs so that the solutions are robust. Hebrard *et al.* [59] examined reformulation of CSPs for finding SSs but found that the approach was inefficient when compared to the MAC-based search algorithm.

But this does not necessarily preclude the formulation of integer programs for robust solutions that are computationally feasible.

An alternative approach for finding robust solutions may be to use a top-down approach by looking at the k-best solutions sequentially, in terms of revenue, and performing sensitivity analysis upon each solution until a robust one is found. In procurement settings the principle of *free disposal* is often discounted and all items must be sold. This reduces the number of potential solutions and the reparability of each solution. The impact of such a constraint on the revenues generated by robust solutions is also left for future work.

Other Applications of the WSSs Framework

The WSSs framework provides a versatile means of establishing robust solutions in many application domains. Chapter 5 demonstrated, in a pedagogical example, how it could be used for job-shop scheduling problems. In this section we briefly discuss some other applications of the WSSs framework.

We propose the use of WSSs for *robust mechanism design*. Porter *et al.* first introduced the notion of *fault tolerant mechanism design* in which agents have private information regarding costs for task completion, but also their probabilities of failure [104]. When the bid-taker has combinatorial valuations for task completions it may be desirable to assign the same task to multiple agents to ensure solution robustness. It is desirable to minimize such potentially redundant task assignments but not to the detriment of completed task valuations. This problem could be modeled using the WSSs framework in a similar manner to that of combinatorial auctions. A minimum set of duplicated tasks may be performed whilst satisfying constraints on reparability in the event of task failure.

When the WSSs framework is applied to optimization problems it may be desirable to express an optimal outcome using a utility function over robustness and the optimization criterion. The framework may be extended to support such utility functions. Solution robustness may be seen as an additional optimization criterion rather than as a constraint, therefore, transforming the problem into a multi-criteria optimization problem. There are inherent trade-offs in such problems that require utility functions to determine an overall optimal solution. Unfortunately, such utility functions are often difficult to articulate precisely and usually involve some uncertainty in themselves [71]. There are means, however, of examining worst-case error in the presence of such uncertainty to guide elicitation of such preferences and suggest refinements of utility information using a notion of *minimax regret* [18], for instance.

The cost of repairing an assignment may consist of a multi-dimensional vector quantity. For example, the cost of repairing a machine after failure involves both time and expense. Total repair costs must, therefore, be measured as a function of such vector quantities and β is the upper bound on the outcome of this function.

An Adaptive Algorithm for Winner Determination

In our auction model, discussed in $\S$ 3.1, the winners are declared simultaneously and the transaction phase involves the simultaneous response of all bidders to this announcement. In some scenarios the transactions process may be of negligible duration and, thus, it may be possible to announce winners sequentially and complete the transaction for each winning bid before declaring the next winner. Sequential declaration of winners increases the knowledge available to the bid-taker when awarding items.

Example 8.2.1. *Consider a combinatorial auction where a horse and cart are for sale where winners are announced sequentially and transactions are completed before the next announcement. Three bids,* $A\langle HorseCart, 900\rangle$*,* $B\langle Horse, 500\rangle$ *and* $C\langle Cart, 500\rangle$*, were entered in the bid submission phase.*

We shall now consider a sequential, partial-winner determination and transaction phase. We can offer either the horse to B *or the cart to* C *initially, and winning-bid withdrawal will allow us to reassess our options after the transaction phase. Let us assume that the the Horse is offered to* B *and the transaction is unsuccessful. The WDP is then reduced to a choice between bids* A *and* C*. The pair of items are then offered to* A *and if this is unsuccessful the Cart is offered to* C*.*

Figure 8.1 shows how a problem may arise in the case where bid B *is transacted successfully. Bid* A *can no longer be considered, therefore, the Cart is offered to* C*. If* C *reneges upon this bid, there is no alternative repair solution*

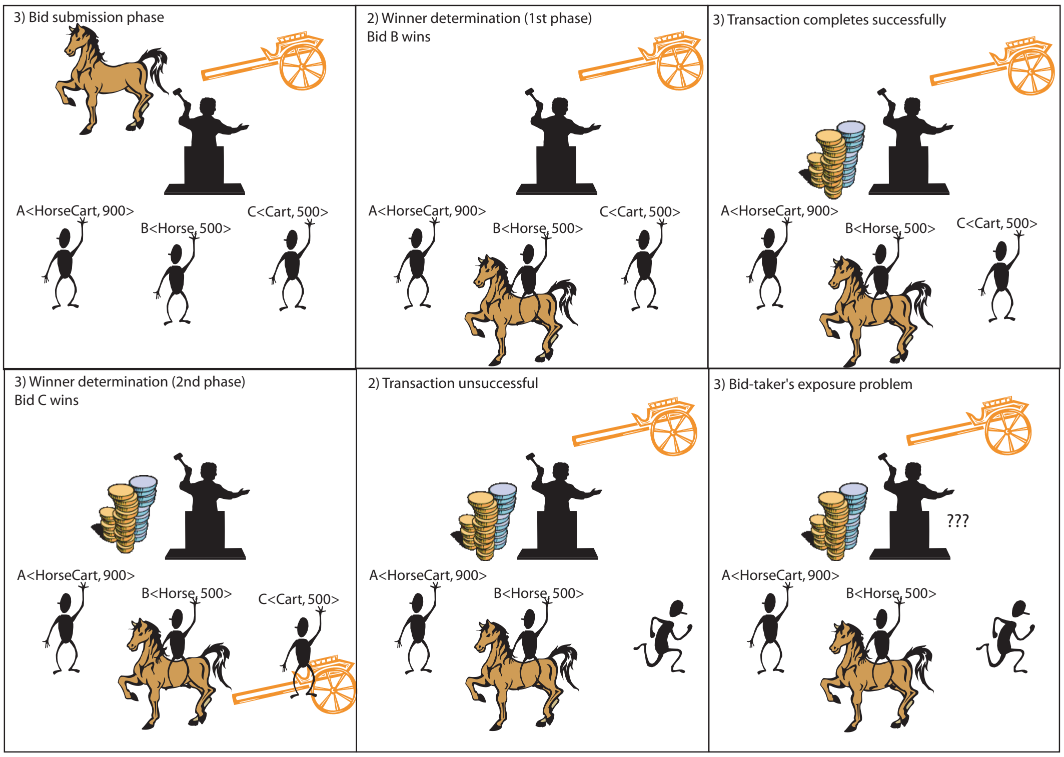

Figure 8.1: Bid-taker's exposure problem for sequential winner determination.

that provides the bid-taker with a revenue greater than 500. The bid-taker's exposure problem is experienced when bid C is withdrawn but it is now reparable if B is withdrawn. $\triangle$

Example 8.2.1 shows how the bid-taker's exposure problem is still present when winner determination and transactions are conducted sequentially. The problem is less severe, however, because bid withdrawal by B is reparable, whereas if the Cart was simultaneously awarded to C it would have been impossible to declare A as winner.

The problem of winner determination in this setting is a stochastic optimization problem. Dean *et al.* [34, 35] describe adaptive algorithms for stochastic packing problems and the benefits of being able to adapt to incoming information. One of the four classes of stochastic packing considered in [35] is stochastic set packing. Recall that the WDP with free disposal is an instance of a set packing problem. They consider a scenario in which there is uncertainty surrounding the set of desired items. However, this differs from the type of uncertainty that causes the bid-taker's exposure problem. We have precise knowledge of desired items but there is a risk that the bid may be withdrawn, in which case the set of desired items is empty and the relevant cost coefficient representing the bid amount is zero. The likelihood of bid withdrawal is represented as a discrete probability distribution.

An algorithm whose decisions depend upon the observations of which previously declared winning bids is called *adaptive*. An allocation algorithm, such as that used in the WSSs framework, that chooses all bids in advance is *non-adaptive*.

We leave the development of an adaptive algorithm for sequential winner determination in the presence of bid withdrawal for future work.

Appendix A

Glossary

Bayes-Nash equilibrium. In normal form games of incomplete information, the players have no possibility to update their prior beliefs about their opponents types. All that a player knows, apart from the game itself (and the priors), is his own type, and the fact that the other players do not know his own type as well. As their best responses, however, depend on the players' actual types, a player must see himself through his opponents' eyes and plan an optimal reaction against the possible strategies of his opponents for each potential type of his own. Thus, a strategy in a Bayesian game of incomplete information must map each possible type of each player into a plan of actions. Then, since the other players' types are unknown, each player forms a best response against the expected strategy of each opponent, where he averages over the (well-specified) reactions of all possible types of an opponent, using his prior probability measure on the type space. Such a profile of type-dependent strategies which are unilaterally unimprovable in expectations over the competing types' strategies forms a Bayes-Nash equilibrium. Source: [123]

Certainty equivalent. The amount of wealth offered for certain, which gives the consumer exactly the same utility as the gamble.

Depth-first search. A search algorithm of a graph that explores the first child of a node before visiting its sibling nodes.

Dominant strategy. A dominant strategy in a game for a player gives a better payoff than any another strategy, regardless of what the other players are doing. It weakly dominates another strategy if it is always at least as good.

Equilibrium bidding strategy. A bidding strategy from which the bidder has no incentive to unilaterally deviate.

Extensive form game. A game in extensive form is one that can be represented as a tree, where each node corresponds to a choice by one of the players. In an extensive form game players make choices sequentially, unlike a normal form game. However, players are not necessarily always aware of the node at which they are positioned (*i.e.*, what moves have been made previously).

Escrow. Money, deeds or a bond that may be held in the custody of a third party for delivery to a grantee following the fulfillment of the conditions specified.

Free disposal. A bid-taker need not necessarily sell all of the items when maximizing revenue, *i.e.* there is no extra cost induced by not selling items.

Game. In general, a game is a formal description of a strategic situation in which players make choices with the intent of optimizing their utility. Formally, a game includes three features:

1. Each player's strategy space that is comprised of a set of pure strategies.
2. Rules for determining an outcome after players' strategies are already decided.
3. A utility function for each player specifying their payoff given each outcome.

Game theory. Game theory is a mathematical theory of strategic interaction where multiple players must make decisions that may affect the interests of other players.

Individual rationality (IR). Each player must expect a positive utility from taking part in a game, otherwise they would not compete. This is also known as the "participation constraint".

Incentive compatibility (IC). IC conditions in a game induce a strategic equilibrium for a given array of types. In particular, they may ensure that it be worthwhile for different types to choose the same action. Yet in most economic problems, incentive compatibility conditions serve to induce a strategic equilibrium which reveals the players' private information by having them choose different *characteristic* equilibrium actions.

Mechanism design. Mechanism design is concerned with specifying the rules of a game where a collection of agents, each holding private information about their preferences over a set of outcomes, interact with each other in their own self-interest.

Mixed strategy. A mixed strategy defines a probability over pure strategies, and randomly chooses a pure strategy based on the distribution given by their mixed strategy. Consequently, every pure strategy is also a mixed strategy with a probability of 1 for the relevant pure strategy.

Nash equilibrium. A profile of strategies such that if other players conform to the equilibrium strategies, no player has an incentive to unilaterally deviate from his equilibrium strategy. A Nash-equilibrium is self-referential and constitutes a profile of strategies that form "optimal reactions" to other agents "optimal reactions". Nash equilibrium is the pure form of the basic concept of strategic equilibrium and as such, it is useful mainly in normal form games with complete information. When allowing for randomized strategies, at least one Nash equilibrium exists in any game with regular payoff functions. Note that a game may possess several Nash equilibria. Source: [123].

Normal form (strategic form). A game that is a compact representation of a game in which players choose strategies simultaneously.

Payoff. The utility of an outcome that reflects its desirability to a player. The expected payoff (in the presence of random outcomes) incorporates the player's risk attitude.

Pure strategy. A pure strategy provides a complete definition for a way a player can play a game. In particular, it defines, for every possible choice a player might have to make, which option the player picks.

Risk attitude. A decision makers risk attitude characterizes his willingness to engage in risky prospects. Focusing on risky prospects with monetary outcomes, a decision maker displays risk aversion if and only if he strictly prefers a certain consequence to any risky prospect whose mathematical expectation of consequences equals that certain amount. Equivalently, a decision maker is said to be risk averse if and only if he strictly refuses to participate in fair games (i.e. games with an expected net outcome of zero). He is said to be a risk preferrer if and only if he strictly prefers the above mentioned risky prospect to its certain consequence. He displays risk neutrality if and only if he is indifferent between the risky prospect and the certain consequence. Source: [123].

Risk premium. The difference between the expected value of the gamble and the certainty equivalent.

Social choice function. A functional representation of how to aggregate individuals preferences into an integrated social preference.

Solution concept. A *solution concept* is used to predict the strategies agents will choose in order to maximize their utility, thus determining an equilibrium position for the game. These concepts assume knowledge about agent preferences, rationality, and shared information about one other. The best known concept is that of a Nash equilibrium.

Strategy. In a normal form game a strategy is one of the given possible actions of a player, whereas in an extensive form game it is a complete plan of actions.

Strategy space. A player's strategy space is the set of pure strategies available to that player.

Winner's curse. The revelation of competing (losing) bidders' valuations infers reduced profitability, or possibly even a loss, for a winning bidder. This problem known as the winner's curse.

Zero-sum game. A game whereby the sum of payoffs to all agents is zero for any outcome.

Bibliography

[1] Maurice Allais. Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine. *Econometrica*, 21:503–546, 1953.

[2] Arne Andersson, Mattias Tenhunen, and Fredrik Ygge. Integer programming for combinatorial auction winner determination. In *ICMAS '00: Proceedings of the Fourth International Conference on MultiAgent Systems (ICMAS-2000)*, page 39, Washington, DC, USA, 2000. IEEE Computer Society.

[3] Martin Andersson and Tuomas Sandholm. Leveled commitment contracts with myopic and strategic agents. *Journal of Economic Dynamics and Control*, 25:615–640, 2001. Special issue on Agent-Based Computational Economics.

[4] Aaron Archer and Éva Tardos. Truthful mechanisms for one-parameter agents. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 482–491, Las Vegas, Nevada, October 2001.

[5] Aaron Archer and Éva Tardos. Frugal path mechanisms. In *SODA '02: Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 991–999, Philadelphia, PA, USA, 2002. Society for Industrial and Applied Mathematics.

[6] Aaron Archer, Christos Papadimitriou, Kunal Talwar, and Éva Tardos. An approximate truthful mechanism for combinatorial auctions with single pa-

rameter agents. In *SODA '03: Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 205–214, Philadelphia, PA, USA, 2003. Society for Industrial and Applied Mathematics.

[7] Kenneth J. Arrow. Aspects of the theory of risk bearing. Helsinki: Trjö Hahnsson Foundation lectures, 1965.

[8] Moshe Babaioff, Noam Nisan, and Elan Pavlov. Mechanisms for a spatially distributed market. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 9–20, New York, NY, USA, 2004. ACM Press.

[9] Fahiem Bacchus and George Katsirelos. EFC constraint solver. `http://www.cs.toronto.edu/~gkatsi/efc/efc.html`.

[10] Russell Bent and Pascal Van Hentenryck. The value of consensus in online stochastic scheduling. In *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 2004)*, 2004.

[11] Michael Berkelaar, Kjell Eikland, and Peter Notebaert. lp_solve version 5.0.10.0. `http://groups.yahoo.com/group/lp_solve/`.

[12] Daniel Bernoulli. Exposition of a new theory on the measurement of risk. *Econometrica (1954)*, 22:23–36, 1738.

[13] Dimitris Bertsimas and Melvyn Sim. Robust discrete optimization and network flows. *Mathematical Programming*, 98:43–71, 2004.

[14] Dimitris Bertsimas and Robert Weismantel. *Optimization Over Integers*. Dynamic Ideas, Belmont, Massachusetts, 2005.

[15] Stefano Bistarelli. *Semirings for Soft Constraint Solving and Programming*, volume 2962 of *Lecture Notes in Computer Science*. Springer, 2004.

[16] Barry Blecherman and Colin F. Camerer. Is there a winners curse in the market for baseball players? mimeograph, Brooklyn Polytechnic University, Brooklyn, NY, 1998.

[17] Dan Borge. *The Book of Risk*. Wiley, 2001.

[18] Craig Boutilier, Tuomas Sandholm, and Rob Shields. Eliciting bid taker non-price preferences in (combinatorial) auctions. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI-04)*, pages 204–211, 2004.

[19] Jeremy Bulow and Paul Klemperer. Prices and the winner's curse. *Rand Journal of Economics*, 33(1):1–21, 2002.

[20] Jeremy Bulow and John D. Roberts. The simple economics of optimal auctions. *The Journal of Political Economy*, 97:1060–1090, 1989.

[21] Edward C. Capen, Robert V. Clapp, and William M. Campbell. Competitive bidding in high-risk situations. *Journal of Petroleum Technology*, 23:641–651, 1971.

[22] Yves Caseau and François Laburthe. Solving various weighted matching problems with constraints. *Constraints*, 5(1/2):141–160, 2000.

[23] James Cassing and Richard W. Douglas. Implications of the auction mechanism in baseballs free agent draft. *Southern Economic Review*, 47:110–121, 1980.

[24] Illinois Manufacturing Extension Center. Supply chain optimization: How to make your company more valuable to your OEM customer. `http://www.imec1.org`, 2005.

[25] E.H. Clarke. Multipart pricing of public goods. *Public Choice*, XI:17–33, 1971.

[26] Wolfram Conen and Tuomas Sandholm. Partial-revelation VCG mechanism for combinatorial auctions. In *Eighteenth national conference on Artificial intelligence*, pages 367–372, Menlo Park, CA, USA, 2002. American Association for Artificial Intelligence.

[27] Wolfram Conen and Tuomas Sandholm. Differential-revelation VCG mechanisms for combinatorial auctions. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 196–197, New York, NY, USA, 2003. ACM Press.

[28] Vincent Conitzer and Tuomas Sandholm. Applications of automated mechanism design. In *The UAI Bayesian Modeling Applications Workshop*, Acapulco, Mexico, 2003.

[29] Vincent Conitzer and Tuomas Sandholm. Automated mechanism design: complexity results stemming from the single-agent setting. In *ICEC '03: Proceedings of the 5th international conference on Electronic commerce*, pages 17–24, New York, NY, USA, 2003. ACM Press.

[30] Vincent Conitzer and Tuomas Sandholm. Computational criticisms of the revelation principle. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 262–263, New York, NY, USA, 2004. ACM Press.

[31] Andrew J. Davenport, Christophe Gefflot, and J. Christopher Beck. Slack-based techniques for robust schedules. In *Proceedings ECP-2001, Sixth European Conference on Planning*, 2001.

[32] Javier David and Nicole Maestri. Google IPO seen as prize catch for NYSE, Nasdaq. Forbes Magazine: `http://www.forbes.com/reuters/newswire/2004/04/28/rtr1351468.html` , April 2004.

[33] Sven de Vries and Rakesh V. Vohra. Combinatorial auctions: A survey. *INFORMS Journal on Computing*, 15(3):284–309, 2003.

[34] Brian C. Dean, Michel X.Goemans, and Jan Vondrák. The benefit of adaptivity: Approximating the stochastic knapsack problem. In *Proceedings of the 45th annual IEEE Symposium on the Foundations of Computer Science (FOCS)*, pages 208–217, 2004.

[35] Brian C. Dean, Michel X.Goemans, and Jan Vondrák. Adaptivity and approximation for stochastic packing problems. In *Proceedings of the 16th annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 395–404, 2005.

[36] Rina Dechter. *Constraint Processing*. Morgan Kaufmann, 2003.

[37] Rina Dechter and Avi Dechter. Belief maintenance in dynamic constraint networks. In *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI-88)*, pages 37–42, St. Paul, MN, 1988.

[38] Didier Dubois, Hélène Fargier Fargier, and Henri Prade. The calculus of fuzzy restrictions as a basis for flexible constraint satisfaction. In *Proceedings of the 2nd IEEE International Conference on Fuzzy Systems*, pages 1131–1136, 1993.

[39] Daniel Ellsberg. Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, November 1961.

[40] Lori Enos. KB Toys snags inventory from bankrupt eToys. `http://www.linuxinsider.com/story/9292.html`, 2001.

[41] Rafael Epstein, Lysette Henríquez, Jaime Catalán, Gabriel Y. Weintraub, and Cristián Martínez. A combinational auction improves school meals in Chile. *Interfaces*, 32(6):1–14, 2002.

[42] Jim Ericson. Reverse auctions: Bad idea. Line 56, September 2001. `http://www.line56.com/articles/default.asp?articleID=2962&TopicID=1`.

[43] Hélène Fargier and Jérôme Lang. Uncertainty in constraint satisfaction problems: a probabilistic approach. In *Proceedings of ECSQARU*, LNCS 747, pages 97–104. Springer Verlag, 1993.

[44] Hélène Fargier, Jérôme Lang, and Thomas Schiex. Mixed constraint satisfaction: A framework for decision problems under incomplete knowledge. In *AAAI/IAAI, Vol. 1*, pages 175–180, 1996.

[45] Joan Feigenbaum and Scott Shenker. Distributed algorithmic mechanism design: Recent results and future directions. In *Proceedings of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, pages 1–13, New York, 2002. ACM Press.

[46] Amos Fiat, Andrew V. Goldberg, Jason D. Hartline, and Anna R. Karlin. Competitive generalized auctions. In *STOC '02: Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 72–81, New York, NY, USA, 2002. ACM Press.

[47] David W. Fowler and K. N. Brown. Branching constraint satisfaction problems for solutions robust under likely changes. In *Proceedings of CP-2000*, LNCS 1894, pages 500–504, September 2000.

[48] Lawrence Friedman. A competitive bidding strategy. *Operations Research*, 4:104–112, 1956.

[49] Milton Friedman and Leonard J. Savage. The utility analysis of choices involving risk. *Journal of Political Economy*, 56:279–304, 1948.

[50] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, August 1991.

[51] Yuzo Fujishima, Kevin Leyton-Brown, and Yoav Shoham. Taming the computational complexity of combinatorial auctions: optimal and approximate approaches. In *Proceedings of the Sixth International Joint conference on Artificial Intelligence, IJCAI-99*, pages 548–553, 1999.

[52] Matthew L. Ginsberg, Andrew J. Parkes, and Amitabha Roy. Supermodels and Robustness. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)*, pages 334–339, Madison, WI, 1998.

[53] Larry Giunipero, R. David Nelson, and Robert A. Kemp. Supply management beyond 2010: What we must know. In *90th Annual International Supply Management Conference*, May 2005.

[54] Rica Gonen and Daniel Lehman. Optimal solutions for multi-unit combinatorial auctions: Branch and bound heuristics. In *Proceedings of The 2nd ACM Conference on Electronic Commerce*, pages 13–20, 2000.

[55] Theodore Groves. Incentives in teams. *Econometrica*, 61:617–631, 1973.

[56] Ron Harstad, John Kagel, and Dan Levin. Comparative static effects of number of bidders and public information on behavior in second-price

common-value auctions. *International Journal of Game Theory*, 24:297–319, 1995.

[57] Ronald M. Harstad and Michael H. Rothkopf. Withdrawable bids as winner's curse insurance. *Operations Research*, 43(6):982–994, November-December 1995.

[58] Emmanuel Hebrard, Brahim Hnich, and Toby Walsh. Robust solutions for constraint satisfaction and optimization. In *Proceedings of European Conference on Artificial Intelligence (ECAI-0)4*, pages 186–190, 2004.

[59] Emmanuel Hebrard, Brahim Hnich, and Toby Walsh. Super solutions in constraint programming. In *Proceedings of CP-AI-OR 2004*, pages 157–172, 2004.

[60] Edward Helmore. Google takes high road to market. Guardian Weekly, May 06 2004.

[61] Gail Hohner, John Rich, Ed Ng, Grant Reid, Andrew J. Davenport, Jayant R. Kalagnanam, Ho Soo Lee, and Chae An. Combinatorial and quantity-discount procurement auctions benefit Mars Incorporated and its suppliers. *Interfaces*, 33(1):23–35, 2003.

[62] Alan Holland and Barry O'Sullivan. Robust solutions for combinatorial auctions. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, pages 183–192, May 2005.

[63] Alan Holland and Barry O'Sullivan. Super solutions for combinatorial auctions. In *Ercim-Colognet Constraints Workshop (CSCLP '04)*, LNAI 3419, pages 187–200. Springer, Lausanne, Switzerland, 2005.

[64] Alan Holland and Barry O'Sullivan. Weighted super solutions for combinatorial auctions. In *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI-05)*, July 2005.

[65] Benoît Hudson and Tuomas Sandholm. Effectiveness of preference elicitation in combinatorial auctions. In *AAMAS '02: Revised Papers from*

the Workshop on Agent Mediated Electronic Commerce on Agent-Mediated Electronic Commerce IV, Designing Mechanisms and Systems, pages 69–86, London, UK, 2002. Springer-Verlag.

[66] ILOG Inc. CPLEX. http://www.ilog.com/products/cplex/.

[67] Selective Insurance. Business insurance. http://www.selectiveinsurance.com/psApps/Business/Ins/bonds.asp?bc=13.16.127.

[68] Niels Erik Jensen. An introduction to bernoullian utility theory, i: Utility functions. *Swedish Journal of Economics*, 69:163–183, 1967.

[69] John H. Kagel and Dan Levin. The winner's curse and public information in common values auctions. *American Economic Review*, 76:894–920, 1986.

[70] Ryan Kastner, Christina Hsieh, Miodrag Potkonjak, and Majid Sarrafzadeh. On the sensitivity of incremental algorithms for combinatorial auctions. In *IEEE International Workshop on Advanced Issues of E-Commerce and Web-Based Information Systems (WECWIS)*, pages 81–88, June 2002.

[71] Ralph L. Keeney and Howard Raiffa. *Decisions with multiple objectives: preferences and value trade-offs*. Wiley, New York, 1976.

[72] Paul Klemperer. *The Economic Theory of Auctions*. Edward Elgar, Cheltenham, United Kingdom, 2000.

[73] Paul Klemperer. Why every economist should learn some auction theory. In Mathias Dewatripont, Lars Peter Hansen, and Stephen Turnovsky, editors, *Advances in Economics and Econometrics: Invited Lectures to 8th World Congress of the Econometric Society*. Cambridge University Press, 2003.

[74] Vijay Krishna. *Auction Theory*. Academic Press, 1 edition, March 2002.

[75] Harold W. Kuhn. The Hungarian Method for the Assignment Problem. *Naval Research Quarterly*, 2:83, 1955.

[76] Javier Larrosa and Thomas Schiex. Solving weighted CSP by maintaining arc consistency. *Artificial Intelligence*, 159(1-2):1–26, November 2004.

[77] Daniel Lehmann, Liadan Ita O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, 2002.

[78] Kevin Leyton-Brown, Mark Pearson, and Yoav Shoham. Towards a universal test suite for combinatorial auction algorithms. In *ACM Conference on Electronic Commerce*, pages 66–76, 2000.

[79] R.D. Luce and H. Raiffa. *Games and decisions*. John Wiley & Sons, New York, 1957.

[80] Mark Machina. "expected utility" analysis without the independence axiom. *Econometrica*, 50(277–324), 1982.

[81] Robert Martin. Stanford encyclopedia of philosophy. `http://plato.stanford.edu/entries/paradox-stpetersburg/`, 2004.

[82] Andreu Mas-Collel, Micheal D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[83] John McMillan. Selling spectrum rights. *Journal of Economic Perspectives*, 8(3):145–62, 1994.

[84] Paul Milgrom. Auctions and bidding: A primer. *Journal of Economic Perspectives*, 3:3–22, 1989.

[85] Paul Milgrom. *Putting Auction Theory to Work*. Cambridge, March 2004.

[86] Paul Milgrom and R. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50:191–215, 1982.

[87] Rémi Monasson, Riccardo Zecchina, Scott Kirkpatrick, Bart Selman, and Lidror Ttroyansky. Determining computational complexity from characteristic 'phase transition'. *Nature*, 400:133–137, 1999.

[88] Ahuva Mu'alem and Noam Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. In *Eighteenth national conference on*

Artificial Intelligence, pages 379–384, Menlo Park, CA, USA, 2002. American Association for Artificial Intelligence. Extended version available at http://www.cs.huji.ac.il/˜ahumu/ .

[89] Milton Mueller. New Zealand's revolution in spectrum management. *Information Economics and Policy*, 5(2):159–177, July 1993.

[90] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, pages 58–73, 1981.

[91] John F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.

[92] George Nemhauser and Laurence Wolsey. *Integer and Combinatorial Optimization*. John Wiley and Sons, 1999.

[93] Noam Nisan and Amir Ronen. Algorithmic mechanism design. In *Proceedings of the 31st Annual ACM Symposium on Theory Computing*, pages 129–140, 1999. Extended Abstract.

[94] Noam Nisan and Amir Ronen. Computationally feasible VCG mechanisms. In *ACM Conference on Electronic Commerce*, pages 242–252, 2000.

[95] NIST/SEMATECH e-handbook of statistical methods. http://www.itl.nist.gov/div898/handbook/ , 2004.

[96] Axel Ockenfels and Alvin E. Roth. The timing of bids in internet auctions: market design, bidder behavior, and artificial agents. *AI Magazine*, Fall:79–87, 2002.

[97] Associated General Contractors of America. Associated general contractors of America white paper on reverse auctions for procurement of construction. http://www.agc.org/content/public/pdf/Member_Resources/ReverseAuctionWhitePaper.pdf , 2003.

[98] National Society of Professional Engineers. A basic guide to surety bonds. http://www.nspe.org/pracdiv/76-02surebond.asp .

[99] Christos H. Papadimitriou and Kenneth Steiglitz. *Combinatorial Optimisation*. Prenitce-Hall, 1982.

[100] David C. Parkes. Auction design with costly preference elicitation. *Annals of Mathematics and AI*, 2004. Special Issue on the Foundations of Electronic Commerce, Forthcoming.

[101] David C. Parkes, Ruggiero Cavallo, Nick Elprin, Adam Juda, Sebastien Lahaie, Benjamin Lubin, Loizos Michael, Jeffrey Shneidman, and Hassan Sultan. ICE: An iterative combinatorial exchange. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, 2005.

[102] Martin Pesendorfer and Estelle Cantillon. Combination bidding in multi-unit auctions. Harvard Business School Working Draft, 2003.

[103] David P. Porter. The effect of bid withdrawal in a multi-object auction. *Review of Economic Design*, 4(1):73–97, 1999.

[104] Ryan Porter, Amir Ronen, Yoav Shoham, and Moshe Tennenholtz. Mechanism design with execution uncertainty. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*, pages 414–421, 2002.

[105] Ryan Porter and Yoav Shoham. On cheating in sealed-bid auctions. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 76–84, New York, NY, USA, 2003. ACM Press.

[106] John W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32:122–136, 1964.

[107] Jean-Charles Régin. Global constraints and filtering algorithms. In Michela Milano, editor, *Constraint and Integer Programming – Towards a Unified Methodology*, chapter 4, pages 89–129. Kluwer Academic Publishers, 2004.

[108] John G. Riley and William F. Samuelson. Optimal auctions. *American Economic Review*, 71:381–392, 1981.

[109] Amir Ronen. Algorithms for rational agents. In *Conference on Current Trends in Theory and Practice of Informatics*, pages 56–70, 2000.

[110] M.H. Rothkopf, A. Pekeč, and R.M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.

[111] Michael H. Rothkopf. A model of rational competitive bidding. *Management Science*, 15:362–373, 1969.

[112] Michael H. Rothkopf and Aleksandar Pekeč. Combinatorial auction design. *Management Science*, 4(11):1485–1503, November 2003.

[113] Zsófia Ruttkay. Fuzzy constraint satisfaction. In *Proceedings of the 3rd IEEE International Conference on Fuzzy Systems*, pages 542–547, Orlando, FL, 1994.

[114] Daniel Sabin and Eugene C. Freuder. Contradicting conventional wisdom in constraint satisfaction. In A. Cohn, editor, *Proceedings of the 11th European Conference on Artificial Intelligence*, pages 125–129, 1994.

[115] Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135(1-2):1–54, 2002.

[116] Tuomas Sandholm and Victor Lesser. Leveled Commitment Contracts and Strategic Breach. *Games and Economic Behavior*, 35:212–270, January 2001.

[117] Tuomas Sandholm and Victor Lesser. Leveled commitment contracting: A backtracking instrument for multiagent systems. *AI Magazine*, 23(3):89–100, 2002.

[118] Tuomas Sandholm, Sandeep Sikka, and Samphel Norden. Algorithms for optimizing leveled commitment contracts. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pages 535–541. Morgan Kaufmann Publishers Inc., 1999.

[119] Tuomas Sandholm, Subhash Suri, Andrew Gilpin, and David Levine. CABOB: A fast optimal algorithm for combinatorial auctions. In *IJCAI*, pages 1102–1108, 2001.

[120] Tuomas Sandholm and Yunhong Zhou. Surplus equivalence of leveled commitment contracts. *Artificial Intelligence*, 142:239–264, 2002.

[121] Vernon L. Smith. An experimental study of competitive market behavior. *Journal of Political Economy*, 70:111–137, 1962.

[122] Vernon L. Smith. *Bargaining and Market Behavior: Essays in Experimental Economics*. Cambridge University Press, 2000.

[123] Sonderforschungsbereich 504. `http://www.sfb504.uni-mannheim.de/glossary/`.

[124] Supply Chain Today. Supply chain refresher. `http://www.supplychaintoday.com/reflections.htm`, February 2005.

[125] Theodore L. Turocy and Bernhard von Stengel. Game theory. CDAM Research Report: Number LSE-CDAM-2001-09, Centre for Discrete and Applicable Mathematics, London School of Economics, October 2001.

[126] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(32):8–37, 1961.

[127] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

[128] Keith Waehrer, Ronald M. Harstad, and Michael H. Rothkopf. Auction form preferences of risk-averse bidtakers. *Rand Journal of Economics*, 29:179–192, 1998.

[129] Richard J. Wallace and Eugene C. Freuder. Stable solutions for dynamic constraint satisfaction problems. In *Principles and Practice of Constraint Programming - CP98*, volume 1520, pages 447–461. Springer, 1998.

[130] Toby Walsh. Stochastic constraint programming. In *Proceeding of the European Conference on Artificial Intelligence (ECAI-98)*, pages 111–115, 2002.

[131] William E. Walsh, Michael P. Wellman, and Fredrik Ygge. Combinatorial auctions for supply chain formation. In *ACM Conference on Electronic Commerce*, pages 260–269, 2000.

[132] Richard Weber. Mathematics for operational research. `http://www.statslab.cam.ac.uk/~rrw1/mor/s.pdf` , 2003.

[133] Waloddi Weibull. A statistical distribution function of wide applicability. *Journal of Applied Mechanics*, pages 293–297, 1951.

[134] Rainier Weigel and Christian Bliek. On reformulation of constraint satisfaction problems. In *Proceedings of the European Conference on Artificial Intelligence (ECAI-98)*, pages 254–258, 1998.

[135] Michael P. Wellman, William E. Walsh, Peter R. Wurman, and Jeffrey K. MacKie-Mason. Auction protocols for decentralized scheduling. *Games and Economic Behavior*, 35(1):271–303, 2001.

[136] Ben White. Aiming to auction its way to a more 'inclusive' IPO: Complex scheme could confuse small investors. Washington Post, April 2004.

[137] Margaret W. Wiener. FCC release: Access spectrum bid withdrawal. `http://wireless.fcc.gov/auctions/33/releases/da011719.pdf` , July 2001.

[138] Robert B. Wilson. Competitive bidding with asymmetric information. *Management Science*, 13:816–820, 1967.

[139] Robert B. Wilson. Competitive bidding with disparate information. *Management Science*, 15:446–448, 1969.

[140] Robert B. Wilson. A bidding model of perfect competition. *Review of Economic Studies*, 44(3):511–18, 1977.

[141] Laurence Wolsey. *Integer Programming*. John Wiley and Sons, 1998.

[142] Neil Yorke-Smith and Carmen Gervet. Certainty closure: A framework for reliable constraint reasoning with uncertainty. In *Principles and Practice of Constraint Programming - CP03*, 2003.

Index