CALCULATION AND ANALYSIS OF NASH EQUILIBRIA OF VICKREY-BASED PAYMENT RULES FOR COMBINATORIAL EXCHANGES

A Thesis presented

by

David Krych

to

Computer Science and Economics

in partial fulfillment of the honors requirements

for the degree of

Bachelor of Arts

Harvard College

Cambridge, Massachusetts

April 2003

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation

The classic story of efficient markets assumes that the demand and supply for various goods are sufficiently uncorrelated that one may consider the supply/demand problem for every individual good separately. And as long as one has many buyers and sellers, and as long as goods are not complements with each other to any very strong degree, this does not present a problem.

However, these assumptions do not hold in certain applications. For example, the FCC is considering a secondary market for radio spectrum licenses [KW02]. While many of these licenses used to be given away rather freely (e.g. via lottery), the scarcity of remaining usable radio spectrum as a result of the explosion of new wireless technologies in recent years has led these licenses to become more and more valuable as a scarce commodity. While there have been government auctions of radio spectrum, the vast majority of the spectrum has now already been licensed. As new wireless technologies emerge that require radio spectrum to operate, it will be necessary to reallocate currently licensed spectrum, rather than putting unused portions of the spectrum up for auction.

Spectrum licenses are generally specified to give radio transmission rights for a given portion of the radio spectrum in a given geographic location of the country.

For example, a wireless company may have the rights to communicate on a certain frequency in the Los Angeles area according to a particular license. To be allowed to broadcast elsewhere requires a different license. In such a market, complementarities quickly arise.

For example, suppose that AT&T wishes to be able to offer a new service throughout California, without forcing its users to switch networks. Such service would make a good selling point to prospective AT&T customers in California, and would likely save AT&T a substantial amount of money–it could support one network infrastructure for the whole of the state, rather than creating two separate infrastructures designed to operate in two different areas of spectrum. As a result, the two licenses together, as a bundle, are worth substantially more to AT&T than two other equivalent licenses. The logical solution to this dilemma is to allow bids and asks on groups of licenses within the market, rather than a single license. We allow one company to bid for two licenses as a unit while allowing other companies to sell their individual licenses or groups of licenses. This is known as a combinatorial double auction, or combinatorial exchange.

We continue a previous analysis of a design for a combinatorial exchange provided in Parkes et al. [PKE01b] (see also longer version [PKE01a]; hereafter referred to as PKE). First, we demonstrate a search-based algorithm to find a symmetric Bayesian-Nash equilibrium (BNE) of the mechanisms initially studied in PKE. Calculating the BNE of the exchanges is computationally problematic. The space of possible bidding strategies which we consider is continuous, and thus many algorithms for calculating Nash equilibria are of little use. We could discretize the space to circumvent the continuous-strategy-space problem, but determining the payoff to an agent of a strategy is computationally difficult. Determining the allocation of a combinatorial exchange is an $\mathcal{NP}$-hard problem [RH95], and so calculating a normal-form expected-payoff matrix on a reasonably discretized strategy space is computationally intractable in any case. Thus the problem becomes one of creating an algorithm which selectively calculates certain elements of the payoff matrix in an attempt to find a Nash equilibrium. We demonstrate the effectiveness of a gradient-descent-based search algorithm for finding Nash equilibria consistently and far more rapidly than the calculation of a payoff matrix. Later in the thesis, we compare the running time of our algorithm to that of a simple enumeration approach. Second, having found the BNEs, we seek to judge the expected allocative efficiency of the various exchange designs. PKE proposed a family of designs that distribute surplus in different ways and have different incentive properties. All the designs have the property that if agents bid truthfully, the designs are efficient, so the designs that minimize the degree of strategic manipulation of values are the most efficient. We empirically test these designs to compare the equilibrium strategies and allocative efficiency of each design.

1.2 Foundation

In this thesis, we analyze the combinatorial exchange under various configurations. In earlier work, PKE presented these rules as a means of creating a reasonable combinatorial exchange. Parkes presented his findings and suggestions at an FCC workshop on the design of an exchange [Fed01], but there is still no agreement on the design.

PKE takes an unorthodox approach to mechanism design. Typical mechanism design tends to focus on the designing of incentive-compatible mechanisms, implementing inefficiencies or other limitations as results of the mechanism. In the paper, the authors create a family of non-incentive compatible mechanisms, in which winners are chosen to maximize their revealed surplus through their bids. The goal, then, is to determine which of these mechanisms cause agents to bid close to their true values. These should lead to higher efficiency.

While PKE does not calculate an equilibrium, it gives an analysis of the payment rules and estimates their efficiency. In particular, PKE concludes with a partial ordering of the rules based on efficiency: Threshold, Large $\succcurlyeq$ Fraction $\succcurlyeq$ Reverse $\succcurlyeq$ Equal, Small. The shortcoming found in the Parkes et al. paper is that the exact

	1	2	3
1	0, 0	0, 0	0, 0
2	0, 0	20, 20	0, 30
3	0, 0	30, 0	10, 10

Figure 1.1: A simple game

BNEs of the markets were not computed. Parkes calculated the expected gains from trade by a single agent, given that all other agents behaved identically–this is typical of a Nash equilibrium calculation. However, PKE gave the agent just one option–it could either play the same strategy as the rest, or it could bid truthfully. It was then assumed that all players would play the strategy that maximized the gain over truthful bidding. This, however, is not necessarily the case with self-interested agents. Essentially, what was unaccounted for was the possibility of a prisoner's dilemma-type scenario. For example, consider the following vastly simplified game:

In PKE, the authors only considered the diagonal elements and the first column–bidding the same as everyone else and bidding truthfully. Their calculation would assume that strategy 2 is chosen, since it maximizes the gain over truthfulness: 20 vs. 10 for strategy 3. However, it is clear that the above game is a classic prisoner's dilemma, and that strategy 2 is not a Nash equilibrium–if all other agents play strategy 2, then it is in any given agent's best interest to play strategy 3. By contrast, if all other agents play strategy 3, it is then in any given agent's best interest to play 3 as well, and we have a stable solution. Thus the agents, acting strategically, do not pick the solution which maximizes the gain over truthfulness. In this paper we expand on the previous work by calculating the actual BNE of the games and performing the efficiency analysis on the Nash equilibrium strategies.

1.3 Main Results

We show that search is an effective means of finding restricted Nash equilibria of the markets, in spite of difficulties imposed by the nature of the combinatorial exchange, and provide an algorithm to calculate the equilibria of the market under various rules for dividing the surplus. While our algorithms are still not what one would call fast, there has been little work on calculating BNE in infinite strategy spaces before. Further, although requiring two to six hours of CPU time to run on a Pentium IV, this is nonetheless an extreme improvement over discretization and manual solving.

On the economic side, our experiments show that the "Large-First" rule is the closest to incentive-compatibility and efficiency, with equilibrium behavior very near optimal. However, it is the only rule suggested which has incentives for over-bidding and under-asking, thus meaning that the rule is not strictly *ex post* individually rational. In particular, the BNE of the rule has bidders overbidding for bundles in hopes of receiving Vickrey discounts if they outbid their opponents by a sufficient amount. This means that if agents play the equilibrium strategy, while on average they will gain over truthfulness, there is the potential that in any given game they will lose money. Of the remaining rules, the threshold rule comes closest to the optimum, and retains *ex post* individual rationality at the cost of some small measurable loss in *ex ante* efficiency of the market compared to that obtained with the large-first rule. Which rule is more desirable from the mechanism design perspective depends largely on whether one is willing to trade efficiency for a guarantee of *ex post*, rather than merely *ex ante*, individual rationality in the selected mechanism.

1.4 Related Work

The combinatorial exchange (double combinatorial auction) problem is a relatively new problem to study, though restricted attempts are beginning to be made. However, the combinatorial exchange is a generalization of problems which have been studied extensively, in particular the bargaining game and the one-sided combinatorial auction. The bargaining game, in which two agents (one buyer and one seller) simultaneously state a bid and an ask for a single item, is a special case of the combinatorial exchange with one buyer, one seller, and one item. The combinatorial auction, heavily studied recently, is a combinatorial exchange with only one seller.

1.4.1 Bargaining Game

The single-buyer, single-seller, single-item bargaining game is small enough to be studied analytically and has been studied quite thoroughly. Much of the work has assumed that, assuming the buyer bids b and the seller s, the price at which the transaction takes place is $\frac{b+s}{2}$, while other work has assumed the existence of a broker who takes the surplus (and thus the buyer pays b, while the seller receives s, and the broker receives surplus $b - s$). For this paper, each of our payment rules assume that all the surplus is divided among the agents, and not given to a broker, except under the "No Discount" rule, which assumes that a broker takes all the surplus. Several of the rules ("Threshold," "Reverse Threshold," "Fraction," and "Equal") reduce to this rule when applied to the bargaining game.

Chatterjee and Samuelson [CS83] took the bargaining game and assumed no broker, surplus divided evenly between the two agents, and values distributed uniformly on the interval $[0, 1]$. They then found the optimal Nash equilibrium of the game. Denoting the buyer's value as v_b and the seller's value as v_a, they concluded that the optimal Nash equilibrium (i.e., the Nash equilibrium which generated the greatest possible expected gains from trade) occurred when the buyer bid $\frac{2}{3}v_b + \frac{1}{4}$, and the seller asked $\frac{2}{3}v_s + \frac{1}{12}$. Leininger, Linhart and Radner [LLR89] expanded on this finding, showing that there were in fact an uncountably infinite number of asymmetric Nash equilibria of the negotiation game, and confirming Chatterjee and Samuelson's finding that the above "linear" equilibrium leads to the maximum possible gains from trade given the above definition of the rules of the game. Leininger, Linhart, and

Radner [LLR89] also show that discretizing the space of possible bids does not significantly affect strategies–the step-wise equilibria which they found follow the contours of the linear equilibrium, becoming closer and closer as more steps were allowed until eventually, in the limit, the equilibrium became the same as that in the continuous strategy space.

In a seminal paper, Myerson and Satterthwaite [MS83] went on to study the negotiation game and posited what is now called the Myerson-Satterthwaite Impossibility Theorem. The conclusion of this theorem is that market efficiency in equilibrium, *ex post* individual rationality, and budget balance cannot coexist in any exchange where sellers have a reservation price. Instead, a trade-off must be made. It is interesting as an aside that while this is known as the impossibility theorem, it was actually a corollary to their paper's other, positive result– calculating the conditions under which an incentive-compatible and individually-rational mechanism could exist. Though proven originally for the negotiation game, the Myerson-Satterthwaite finding generalizes to the combinatorial exchanges studied in this paper, and in fact in any case so long as the seller can set a reservation price in the auction. This motivates our choice of non-incentive compatible, inefficient payment rules and our desire to determine how close to efficiency they may or may not be, rather than seeking a payment rule which actually gives the optimal allocation, given the values of the agents.

1.4.2 Calculation of Large- and Continuous-Space-Game Nash equilibria

There has also been substantial work in the field of solving efficiently for the Nash equilibria of simultaneous-play normal-form games through geometric methods. For two-person normal-form games, the Lemke-Howson algorithm [LJ64] provides an efficient way of finding mixed-strategy Nash equilibria, as the program reduces to a linear complementarity problem. In the case of zero-sum two-player games, the Lemke-Howson algorithm reduces even further to the well-studied problem of linear programming [MM96].

For the more general case of n-person normal-form games where $n > 2$, the Lemke-Howson algorithm is insufficient. Instead, simplical subdivision algorithms such as that of Scarf [Sca73] must be used. Scarf's algorithm recalls the Nash Existence Theorem, and solves the problem of finding a Brouwer fixed point based on the evaluation of function values. However, Hirsch, Papadimitriou, and Vavasis [HPV89] prove that any such algorithm is exponential in both the dimension of the function (in the game theoretic case, the number of agents playing the game) and the number of digits of accuracy.

Last month, separately from work on this paper, Reeves and Wellman [RW03] studied the problem of finding symmetric Nash equilibria in games where the set of strategies is a continuous compact subset of $\mathbb{R}$. Their algorithm, like ours, involves repeated calculation of best-response strategies as a means to search within the continuous space. Their algorithm differs from ours, however. The Reeves-Wellman algorithm finds the global best-response to each strategy (through analytic means for a restricted class of games, or by Monte Carlo simulation), and worked by moving directly to each new best-response in an iterated manner, checking each possible response down to a specified level of granularity. Reeves and Wellman find, however, that their current Monte Carlo simulator is "practically intractable," and are considering modifying their mechanism in future work to automatically adjust the granularity of their search. Because of the extreme time complexity involved in calculating payoffs of the combinatorial exchange, calculating global best-response strategies was not feasible in the current study. At the end of a run of our algorithm, we run a check to ensure that our payoff is a best-response– and this one check takes from 6-50 minutes of CPU time on a Pentium IV, depending on the payment rule. As a result, for our work, we have already adopted an adaptive strategy for finding Nash equilibria. We calculate a small, evenly distributed sample of the payoffs, and move our search part-way in the direction of that strategy in the space, and then use that as the basis for the following iteration.

1.4.3 Combinatorial Auctions

While not the subject of the current paper, the computational characteristics of the winner determination problem in the combinatorial auction setting has been much studied. In terms of winner determination, the combinatorial exchange is the same as the combinatorial (single) auction. Asks input by sellers are equivalent to bids; if a seller is allocated a bundle he offered, then in the most efficient allocation no bids were offered for his bundle that would have produced higher utility. Thus the complexity result that applies for combinatorial auctions also applies to combinatorial exchanges-the selection of the allocation which maximizes the total surplus to the bidders is $\mathcal{NP}$-complete by reduction to the maximum-weight set-packing problem [RH95][BdVSV01]. Thus there is no computationally tractable method to determine the optimum winner of the auction. Nisan and Ronen [NR00] attempt to tackle this problem by creating markets based on the VCG system which could be solved in polynomial time. However, they note that as a result the markets rewarded deviation from bidding one's true value for an item– the exchange lost incentive compatibility and thus its efficiency. Sandholm [SSGL01] attacks the problem by using a heuristic search algorithm to determine the most efficient winning allocation, taking advantage of bid sparseness to define a tree on which to search for the optimal allocation. Others have chosen the approach of working with well-known algorithms for the maximum-weight set-packing problem, whether they knew it or not [ATY00]fujishima99. In particular, Andersson et al. have solved the problem by formulating the winner determination problem as the following mixed integer program [ATY00].

1.5 Thesis Outline

In Chapter 2, we give the basic definitions, notations and results that are used throughout this thesis, which are staples of the mechanism design literature. In Chapter 3, we give our algorithm and the results of experiments on its run time and its ability to find Nash equilibria. In Chapter 4, we show how we apply the algorithm to compare the economic properties of proposed payment rules for combinatorial exchanges. Finally, in Chapter 5 we discuss the main result, give some future directions, and conclude. At the end, in the appendix, we document a few paths that were taken in this research that were not useful.

Chapter 2

The Combinatorial Exchange Problem

An exchange is a means of allocating goods among individuals. It does this by matching bids (offers to buy a good at a given price), and asks (offers to sell a good at a given price). In the case of a combinatorial exchange, we allow bids and asks to be made on "bundles" of goods, rather than just the goods themselves. A bundle of goods is a set of goods (possibly multiple goods of a given type) on which an agent may post bids or asks. Thus each bid and ask in a combinatorial exchange is defined by a set of goods and a price. In the case of a bid, the price is the maximum that the agent will allow himself to pay through the mechanism. In the case of an ask, the ask is the lowest price the agent would be willing to receive for the bundle, before it makes more sense for the agent to keep the item rather than selling it. Let i denote an agent, and G a set of goods, each with multiple units.

Definition 1 (Trade) *A trade $\lambda_i = (\lambda_i(1), ..., \lambda_i(|G|))$, defines a transfer of $\lambda_i(g)$ units of good $g \in G$ to agent i. $\lambda_i(g) < 0$ implies that i gives up that quantity of the given good.*

The exchange takes each agent's bids and asks on various bundles, and returns a trade and a price for each agent, defining the new allocations. In particular, we

sell A	100
buy B	500
sell A and buy B	0

assume an exclusive-or (XOR) bidding language, that allows bidders to express their value for any possible bundle individually, and so is completely expressive. Thus an agent participating in the mechanism can give a value for every possible bundle of goods that he might buy or sell. The bids and asks of an agent are induced by its *valuation function*.

Definition 2 (Valuation Function) *An agent's valuation function, $v_i(\lambda_i)$, gives the agent's value for trade λ_i.*

An example of a simple valuation function for an agent in a combinatorial exchange might be the following, if it has one license, A, but needs another, B, to stay competitive in the business:

Definition 3 (Payments) *A payment $p(i)$, is the amount of money paid by agent i to the mechanism. If $p(i) < 0$, then the mechanism pays money to the agent.*

The job of the designer of a combinatorial exchange is to find a mechanism which takes the bids and asks of all the agents participating in the exchange, and returns the trades and the payments at which they are to be carried out. In creating such a mechanism, one may have several goals. The most common, however, are the following three:

Definition 4 (*Ex Post* Individual Rationality) *A mechanism is individually rational if and only if rational bidding by agents does not lead to any agent being worse off after trade.*

In particular, suppose an agent i participates in mechanism M and M returns price p_i and trade λ_i. Then M is individually rational if and only if for all $i \in I$, $v_i(\lambda_i) > p_i$.

Definition 5 (Allocative Efficiency) *A mechanism M is efficient if M returns the trade which maximizes the total utility gain of all agents when all agents play at a Nash equilibrium strategy. Formally, we seek to maximize*

$$\sum_{i \in I} \lambda_i v_i(\lambda_i)$$

in the space of feasible trades.

Definition 6 (Budget Balance) *The mechanism as a whole does not lose money. Formally, we require $\sum_{i \in I} p(i) \geq 0$ in equilibrium.*

Budget balance may be either *ex post* or *ex ante*–either each run of the mechanism is guaranteed not to lose money, or else the expected payment to the mechanism over many runs must be at least 0. In this paper we will use "budget balance" to refer to "*ex post* budget balance" unless stated otherwise.

While the standard one-item ascending first-price auction with no reservation price has all these properties, in the case of any exchange with both bids and asks, we have the following theorem, noted earlier in section 1.4.1:

Theorem 1 (Myerson-Satterthwaite Impossibility Theorem) *In the combinatorial exchange,* ex post *efficiency,* ex post *budget balance, and* ex post *individual rationality cannot coexist. An exchange may have up to any two of these properties, but never the third. [MS83]*

This has led to the study of the pros and cons of various theoretical mechanisms for the combinatorial situation. Most prominent and commonly cited is the Vickrey-Clarke-Groves (VCG) mechanism [Vic61, Cla71, Gro73]. VCG satisfies the efficiency and rationality properties, making it optimal from the view of the participant in the auction (and the benevolent social planner), but is not budget-balanced in the context of an exchange. As a result, it is impractical to use VCG for many allocation problems.

For practical purposes, while the goal of any market system is to allocate resources efficiently, the other two criteria-budget balance and individual rationality-are perhaps more necessary even if they decrease efficiency. Budget balance is required so that the system will not need to be heavily subsidized by the government or any other entity just to run properly. Individual rationality is also a benefit, because if companies must risk losing money when entering the mechanism, they may choose not to participate in the market at all. We accept these as hard constraints, leaving our goal to find a mechanism which leads to as efficient an allocation of resources as possible.

2.1 The Winner Determination Problem

The first part of defining a mechanism is assigning trades to each individual. We choose to maximize the *reported* surplus of the agents. Let $\lambda = (\lambda_1, \lambda_2, ..., \lambda_{|I|})$ denote a set of trades across all agents in I. We wish to find λ^* such that

$$\lambda^* = \text{arg max}_{\lambda \in \Lambda} \sum_{i \in I} v_i(\lambda)$$

where Λ denotes the set of all feasible trades across all agents. A trade is considered feasible for our purposes if for every positive allocation of a good, there is also a negative allocation– i.e. every good that is allocated to an agent comes from some other agent.

The difficulty with the above calculation is that it is $\mathcal{NP}$-hard to solve exactly for a combinatorial exchange, because the problem is equivalent to solving maximum-weight set-packing. As noted above 1.4.3, there are several algorithms for solving the exchange, and we use the method of solving for the maximum of the revealed surplus by using commercial MIP solvers.

2.2 The Payment Problem

The second part of defining a mechanism is determining how much the winners pay and receive. This in particular is the problem on which we focus.

Definition 7 (True Value/Truthful bidding) *Agent i's* true value $v_i(\lambda_i)$ *for a trade λ_i is the amount of money which causes i to be exactly indifferent between having the payment $v_i(\lambda_i)$ and participating in trade λ_i. We say that i* bids (asks) truthfully *if, given the option of bidding on a bundle λ_i, he chooses to bid his true value $v_i(\lambda_i)$*

We adopt a strong notion of incentive compatibility. In this paper, we may discuss the "extent" to which an exchange is incentive compatible, by which we mean how close to truthful bidding agents will be in equilibrium.

Definition 8 (Incentive Compatibility) *A mechanism is incentive compatible if and only if for every agent i, it is a weakly dominant strategy for i to bid truthfully, regardless of the actions of other agents.*

Note that VCG is a central family of mechanisms, with this strong incentive compatibility property.

Definition 9 (Side Payment/Discount) *A side payment or discount is the difference between the amount an individual bids and the amount he pays, or the amount a seller asks and the amount he receives.*

Groves mechanisms, of which the VCG mechanism is that most central to our work, provide discounts to agents to encourage them to bid truthfully.

The VCG discount is of particular interest as it is provably the smallest discount for which efficiency and *ex post* individual rationality can occur, and because the VCG discount is used as a parameter for several of our payment rules and is the standard to which we compare them.

Definition 10 (Vickrey-Clarke-Groves Discount (Vickrey Discount, VCG discount)) *The discount an agent i would receive in a combinatorial exchange mechanism if all the bids to an exchange were the same, but the VCG mechanism were used. It is intuitively described as the marginal utility which an agent contributes to a system, and calculated as:*

$$\delta_{vick,i} = V^* - V^*_{-i}$$

where V^* is the sum of all gains from trade under the optimal allocation of goods given the agents' revealed bids and asks, and where V^*_{-i} is the sum of all gains from trade under the optimal allocation of goods given all the agents' bids and asks with i's bids and asks removed from the mechanism.

The problem of determining what payments the exchange's participants pay or receive, and how these impact the efficiency of the exchange, is the central question of this thesis. Under the VCG system, discounts are given that are at the bare minimum to maintain incentive compatibility [KP98]. With the rules we study here, the side payments are made so as to keep the system budget-balanced, but as a result the mechanisms are *not* incentive compatible. These payment rules give smaller (or at least, no larger) discounts to agents than they would receive under the VCG system. Thus, the best possible way for a given agent to participate in the mechanism may be to not bid as high as he is willing or ask as low as he is willing. When we determine the winner, in all the mechanisms, we take maximize the efficiency of the outcome given the reported values of the players. If players are do not bid their true values for items, the winner determination algorithm will therefore choose a suboptimal allocation. Thus the exchange is no longer guaranteed to be efficient.

In most of the mechanism design literature, incentive compatible mechanisms have been studied. The mechanisms defined in this thesis are *not* incentive compatible, unlike most of those in the mechanism design literature. Rather, we take mechanisms known to not be incentive compatible and compare, among other things, how close they are to incentive compatibility. However, since the allocation is the same, given the same set of bids and asks, as VCG, no one can, after the fact, see an improvement in the result, unlike with directly inefficient mechanisms. The inefficiencies arise not from the approximately-efficient nature of the mechanism itself, but from the alterations to their true values for the bundles that the strategic agents make in placing their bids.

Because none of the payment rules are incentive compatible, we cannot guarantee perfect *ex post* efficiency for any payment rule. This in fact can be made a guarantee based on the Myerson-Satterthwaite impossibility result that the mechanism will *not* be efficient. As a result, it is our goal to find the most efficient payment rule available. To determine the efficiency of each payment rule, we use the following steps:

1. Determine the best-response decision of an agent going into the game, and use such decisions to simulate a market. More details on the approximations and restrictions on strategies which we used are given below.

2. Given an equilibrium strategy, assume that the agents implement it. Calculate the resulting outcome of the market for many sample markets and value distributions.

3. Determine the average efficiency of the market results. We count the efficiency of the market as being ratio of the total of the agents' benefits from trade to the benefits as they would have been had all agents bid their true values.

Whichever rule, on average, leads to the most efficient market outcomes is the one which we would want to choose. However, this leads to several complications. In particular, determining the general best-response strategy of an agent is very difficult. However, at this point the general best-response strategy appears to be beyond calculation for the combinatorial exchange scenario. We choose here to restrict the space of possible strategies for an agent to that of a constant multiplier times the agent's value for a bundle of items.

PKE defines several payment rules that can be used in a combinatorial exchange setting which all satisfy *ex post* budget balance and individual rationality (except in

Rule	$p(i)$	Parameter
Vickrey	$\Delta_{\text{vick},i}$	N/A
No-Discount	0	N/A
Large	$\Delta_{\text{vick},i}$ if $\Delta_{\text{vick},i} > C_l^*$, 0 otherwise	$min(C_l^*)\|BB'$
Small	$\Delta_{\text{vick},i}$ if $\Delta_{\text{vick},i} < C_s^*$, 0 otherwise	$max(C_t^*)\|BB'$
Threshold	$max(0, \Delta_{\text{vick},i} - C_t^*)$	$min(C_t^*)\|BB'$
Reverse Threshold	$min(\Delta_{\text{vick},i}, C_r^*)$	$max(C_r^*)\|BB'$
Fractional	$\mu^* \Delta_{\text{vick},i}$	$V^*/\sum_{i \in I} \Delta_{\text{vick},i}$
Equal	$V^*/\|I^*\|$	N/A

Table 2.1: The Vickrey-based payment rules.

one case where *ex post* individual rationality is replaced by *ex ante* individual rationality). All the rules which PKE posits are based on the agents' assigned Vickrey payments. These payment rules are each based upon the VCG mechanism–winner determination takes place in the same way as with VCG. However, the payoffs to players are not the Vickrey payments, which are perfectly market-efficient but lead to a non-budget balanced mechanism. Before we define the rules, let us define some notation: $\Delta_{\text{vick},i}$ denotes the discount which agent i would receive in the VCG mechanism. $I^* = \{i \in I | (\exists \lambda | x_i(\lambda) = 1)\}$, where $x_i(\lambda)$ is an indicator function that equals 1 if and only if agent i receives bundle λ.

The payment rules are listed in 2.1. The rules are designed to discourage manipulation by individuals. As a baseline, Regular denotes the system where a seller receives his ask and no more; a buyer pays his bid and no less. The mechanism keeps the surplus from trade, rather than distributing it among the participating agents. This is the procedure which, for example, is used by Nasdaq in trading stocks. As seen below, this rule is not very efficient when there are not many buyers and sellers in competition. At the other end of the spectrum, the "Vickrey" rule gives the non-budget balanced, standard VCG payments, and so under this rule it is always best to report a value truthfully. The efficiency and concept of the other payment rules fall in between "regular" and "Vickrey" and provide exact budget balance. "Equal" divides the surplus equally between the bids and asks among all agents involved in transactions.

The remaining payment rules divide the surplus based on the discounts given in the VCG mechanism. "Small" gives VCG payments to as many individuals as possible, giving priority to those with the smallest payments according to the VCG mechanism. "Large" works on the same principle but gives the priority to those parties who would ordinarily receive the largest payments. "Threshold" and "Reverse Threshold" are akin to "Large" and "Small" but are less radical. "Threshold" limits the amount deducted from a VCG discount: bidders with discounts lower than the threshold receive no discount; those with discounts above the threshold receive their discount less the threshold value. "Reverse Threshold" gives every agent discounts up to a certain threshold: those with VCG discounts below the threshold receive their full discount, and those above the threshold receive the threshold value as a discount. Lastly, "Fraction" gives each player a fraction of his VCG discount.

It should be noted that with the exception of the Vickrey rule, all of the mechanisms are defined so that they guarantee *ex post* budget balance. As long as no agent bids above his value for a bundle, or asks below his value, these payment rules are also *ex post* individual-rational, as none of the payment rules ever assign payments larger in absolute value than an agent's bid or ask. However, incentives are such that with the Large payment rule, we find that the agents have an incentive to bid above and ask below their true values, thus losing the mechanism's guarantee of *ex post* individual rationality. We return to this topic and its implications later.

Chapter 3

Search Algorithms for Finding Nash Equilibria

3.1 The Algorithm

We present an adaptive search algorithm for finding the Nash equilibria of the game defined in Chapter 2, based on the gradient descent approach. We begin with a coarse grid to start the search and then make it finer and finer (but also smaller) until we converge to a problem so small that we are within a reasonable level of error. The algorithm produces reasonable, reproducible results starting from different locations in the search space. Further, we verify the outcomes manually after it is finished, and it has always been successful.

Calculation of Nash equilibria in this game is extremely difficult from a computational standpoint. Each player has exponentially many possible strategies as the number of items and the possible number of bundles increase. This only compounds the fact that the winner-determination problem is itself $\mathcal{NP}$-complete, increasing exponentially in size as the number of bids. We instead choose to restrict the game in several ways:

1. We search for *ex ante* Bayesian-Nash equilibria, rather than *ex post* equilibria. That is, a player chooses his overall strategy without regard to his own valuations for the various bundles.

2. We assume a linear strategy space, in which an agent changes all his bids by a constant fraction of his true value. A strategy is denoted by a real number on [-1, 1]. Suppose agent i is putting a bidding to buy bundle B_1 and offering to sell bundle B_2, and values the bundles at $v_i(B_1)$and$v_i(B_2)$, respectively. If i plays strategy s, then i will bid $(1 - s)v_i(B_1)$ for B_1, and ask $(1 + s)v_i(B_2)$ for bundle B_2.

All players, before knowing their own payoffs to the various bundles, are symmetric. As a result, the first assumption implies that any Nash equilibria of the restricted game must be symmetric. This greatly simplifies the problem of finding equilibria of the game by allowing us to generate a large number of possible runs of the market and determining the average payoff to a randomly chosen player, rather than trying to aggregate the various best-response strategies across each individual game (which is a difficult problem to solve even for a particular instance of a given market, let alone for all markets). Since strategies constitute a function from values to bids, however, simply having to pick a strategy ahead of the game is not a particularly restrictive assumption.

Assumption 2, however, mandates that all bids are manipulated by the same percentage. This greatly restricts our strategy space. On the one hand this is a good thing, as it allows us to calculate the Nash equilibria more easily. On the other, it means that the space of strategies is less rich and thus may not be as indicative of real behavior as we might theoretically like. Nonetheless, however, it is a reasonable assumption for a first study and allows us to make some general claims– in the future, it would be interesting to try allowing richer functions.

As previously stated, because the game involves choosing a strategy regardless of one's own value, all players are symmetric. Let I be a set of agents, $i \in I$ a particular agent, and let BR_i denote the best response function for player i. BR_i is expressed formally as:

$$\arg\max_{s \in [-1,1]} E(v_i(\lambda_i)|s_{-i}),$$

where s_{-i} denotes the set of strategies played by all players other than player i. Since the game is symmetric, if a strategy is an expected best response for any agent i, then it is a best response for all agents i. Therefore, if for any agent the best response to all other agents playing s is for that agent to play s, then s is a symmetric Bayesian-Nash equilibrium of the game. It follows, then, that finding equilibria of the above-defined game is a matter of finding such strategies s.

Generally in game theory, determining the equilibria of such games is a complex but solvable task. Generally one is searching on finite strategy spaces; one looks up the payoffs to various strategies, and computes the fixed points of the best-response functions. While the game may offer an extremely large number of strategies, the Scarf algorithm (and variants thereof) can solve normal-form games relatively quickly on average, though with poor worst-case complexity. However, this is only feasible if the payoffs themselves are simple to calculate. In the case of combinatorial exchanges, payoffs are computationally very difficult to calculate, and the space is continuous– as a result, such algorithms fail to be of much use.

Just the calculation of any one entry in the payoff matrix is difficult. Calculating the winner of each auction is an $\mathcal{NP}$-complete problem in and of itself. Calculating the Vickrey discounts further compounds this, requiring clearing the auction an additional time for every agent i which participates in the final trade (i.e. for i where $\lambda_i \neq \vec{0}$). Additionally, since we are calculating a Bayesian-Nash equilibrium and thus adding uncertainty into the calculation because we do not know the other agents' values, many trials (drawn samples from the distribution of agent values) are required to determine the expected value of a given strategy, given a set of other agents' strategies- with each trial, as above, being $\mathcal{NP}$-complete.

As a result, the option of discretizing the space of strategies in order to make them sufficiently fine to enumerate all the possible payoffs and then calculate the Nash equilibrium is impractical. Calculating the normal-form matrix of the game with just six possible deviations, for markets with 20 participants and 100 goods, while averaging the results over 100 trials to calculate the *ex ante* payoffs, takes approximately 6 processor hours. This time grows as the square of the number of options given to a participant increases. Allowing a finer array of choices- for example, 50 options, for specificity at the hundredths level from 0 to .50–would take approximately 16-17 days on a dedicated Pentium IV just to calculate the expected payoffs (this does not include actually running the Scarf algorithm to find an equilibrium). Furthermore, expansion to more options would require the number of trials to rise to increase so that the expected payoffs could be calculated more finely. For our searches at this level of accuracy, we generally accepted 500 trials as being sufficient to achieve consistent results. Any time used by the algorithm to find the Nash equilibrium of the game would be in addition to this.

To avoid this problem, we chose to attempt finding the Bayesian-Nash equilibrium of the game through search methods more common in the artificial intelligence community. Algorithms more widely used in the computational game theory community, such as Scarf's and its variants, usually assume some complete representation of the game, either in normal or extensive form, and assume a finite number of possible strategies. While we discretize the space, in our analysis, we do not obliviously calculate all the payoffs and try to find a BNE. Instead, we begin with an extremely coarse discretization, for which calculating the best response function is relatively quick, and we gradually make the discretization finer as we come closer to finding a symmetric BNE. Using these search methods allows a fine grain to be used in determining the equilibria while still ensuring that the calculation is completed in a reasonable amount of time (the above search, to the accuracy of a hundredth over several sets of random data, takes on the order of 5-6 hours for most rules rather than the 16 days mentioned above for the more simplistic discretization). If we find a strategy using traditional search methods for which the best response to the strategy is the strategy itself, then we have found a Nash equilibrium of the game.

Our algorithms are based on the *gradient-descent* (also known as *hill-climbing*) search approach. In general, these algorithms were implemented as follows:

1. Choose a set of strategies S on which to search from the range $[0, 1]$. (Note: for the Large payment rule, a more reasonable search space is $[-.5, .5]$, as the rule provides incentives for an agent to bid more and ask less than his true value for a bundle).

2. Randomly select some $s_0 \in S$.

3. Create some number n of valuation instances, based on the assumed initial distributions of values and bid amounts. In instance of the exchange, randomly choose an agent a_i to follow.

4. Assume that all agents other than a_i follow strategy s_0. Calculate the expected surplus u_i over the n samples for agent a_i for each of some set of strategies S', where $S' \subseteq [0, 1]$ is not necessarily a subset of S (see below).

5. Take the $s^* \in S'$ with the highest expected utility as an estimate of the *ex ante* best response to s_0.

6. Let s_0 take on some new value between s_0 and s^*, and repeat steps 3-5 until convergence– specifically, when S' is of a sufficiently fine detail level and s_0 is a best response to itself.

The above description of this algorithm leaves open the question of how large the set S' should be compared to S, and how it should be chosen. We experiment with 3 methods of choosing S':

- Let $S \subset [0, 1]$ be a discretization of [0, 1] into equal partitions (e.g. 0, .01, .02, ... 1). Let $S' = S$. This means that we always find the exact *ex ante* best response in S to x_0.

- Same as above, but let S' be a smaller and smaller subset of S. As we near an equilibrium we can refine our search to look at options within a smaller and smaller range around s_0, saving computation time at the expense of not being guaranteed to find a global best-response toward which to move.

- Let S = [0, 1] (do *not* discretize the space). We preset a fixed set size n for S' regardless of the range around s_0 being analyzed. We choose S' to test as a function of $|s^* - s_0|$. Having chosen a range, we choose n equidistant values on the range, and evaluate them to approximate the best response.

It is worthwhile noting that the search heuristic we use-specifically, that the Nash equilibrium will be in the direction of the current best-response-is not technically admissible. Thus, there is no guarantee that this algorithm will ever converge, and furthermore, there is no guarantee that if it does converge, it will converge to a Nash equilibrium. When the algorithm finishes, of course, it is guaranteed that we will have found a Nash equilibrium in the final game defined in the last iteration, but this will be necessarily over a small interval. However, in the overall game, it is possible that we have converged to a local maximum, rather than global maximum. To show definitively that we have a Nash equilibrium, having run the search, we assume that all agents play our candidate Nash except for one. We then discretize the space for that one agent, allowing him to choose any option on 0, .01, .02, .03.... If our candidate Nash maximizes this best-response function, then we have shown that we have found an equilibrium accurate to at least the hundredth. Experiments show that in spite of the theoretical possibility of reaching a local rather than a global maximum, that in practice the algorithm finds a global maximum in almost all cases.

The advantage with the third way to choose S' is that we search effectively on a loose granularity at the beginning of the search, when we are most likely farther from the equilibrium, and work towards a finer granularity as the search progresses, and we come closer and closer to finding the symmetric equilibrium strategy. The latter strategy also presents the advantage that it is not limited to searching only on the given set of strategies S, but rather searches on continuous space. Experiments show that this has advantages, since there is no guarantee that even at the level of the hundredth there will necessarily be an exact pure-strategy Nash equilibrium. For example, if .25 has a best response of .24, and .24 has a best response of .25, the first algorithm will cycle forever (unless we check for this case, stopping the algorithm and then shrinking the search space). If the last algorithm comes up with the same problem, it continues searching within the remaining interval for a stable equilibrium.

Our general definition of the algorithm above does not include a fixed definition of how to determine S' under the second and third methods. Also, we did not indicate exactly how to determine the new s_0 at the end of an iteration. After trying various methods, we settled on the following rules:

$$s_0^{\text{new}} = \frac{(s^* + 2s_0)}{3}$$

$$d^{\text{new}} = \begin{cases} 4|s_0^{\text{new}} - s^*| \text{ if } s_0 \neq s^* \\ \frac{3}{4}d \text{ if } s_0 = s_1 \end{cases}$$

3.2 Computational Experiments

For our experiments with the algorithms, we simulated the mechanism defined by each payment rule, with 10 participants in the market (five buyers and five sellers), trading in 20 separate goods, with 100 bids/asks in the exchange (10 per agent). The winner determination problem is formulated as a mixed integer program (MIP), and then we use IBM's OSL version 3, a commercial MIP solver, to calculate the winner. The same method was used to compute the outcome without each agent in the solution (necessary for the calculation of VCG discounts and payments in our scheme).

Figures 3.1 to 3.3 show how the typical runs of the algorithm work. The vertical bars in each graph are the minimum and the maximum of the range on which we search during each iteration of the algorithm; the other lines connect each level's candidate Nash (s_0 as defined in the above algorithm) and the expected best-response for an agent if all others play the candidate Nash (s^*).

The Vickrey rule graph is particularly instructive as an example, since it is known to be dominant-strategy incentive compatible (Figure 3.1). On this graph the additional horizontal line indicates strategy zero. This run shows that we can find a BNE even if we start off with a range that does not contain the BNE– the initial range was to test manipulations from .15 to .65, yet the Nash equilibrium strategy is zero manipulation. The run used just five sample points per iteration, evenly separated on the range. With this information we can see that the best-response curve is not always at zero, because we do not sample that point until after a number of iterations, when the range has gotten sufficiently fine and s_0 is already very close to zero. However, the best-response curve is always as close to zero as we sample. The impact on the range on which we search is also interesting. The range begins by expanding in the direction of the best-response, as the fact that the best response found is a minimum on the range indicates that the range itself was too fine and did not include the true best response. This is correct the first time (when the minimum of the range and the best response was .15), and incorrect the second (when the minimum and the best response was .38). Once we find the best response in the interior of the coarse search space, however, we begin to shrink the space, which makes the search finer. The algorithm then proceeds to search on finer and finer levels, until eventually by the end of the graph the range is just over .001, at which point the algorithm returns.

No Discount (Figure 3.2) is interesting but for the opposite reason from Vickrey– we know that the strategy will encourage manipulation. The same basic pattern applies, however. We start out on the same range as we did for Vickrey, but in this case we expand the range upward first, since the maximum amount of manipulation which we were testing was the best. Unlike with Vickrey, however, this was not because the range did not include the equilibrium, but because the range was course. The following iteration has a larger range and so the calculated best response is on the interior of the range, and thus the narrowing process starts.

The Large rule (Figure 3.3) is interesting because we find a solution not only not in our initial search range, but in fact below zero. Aside from this, however, the steps essentially remain the same as above. The zero line is drawn on the graph to indicate

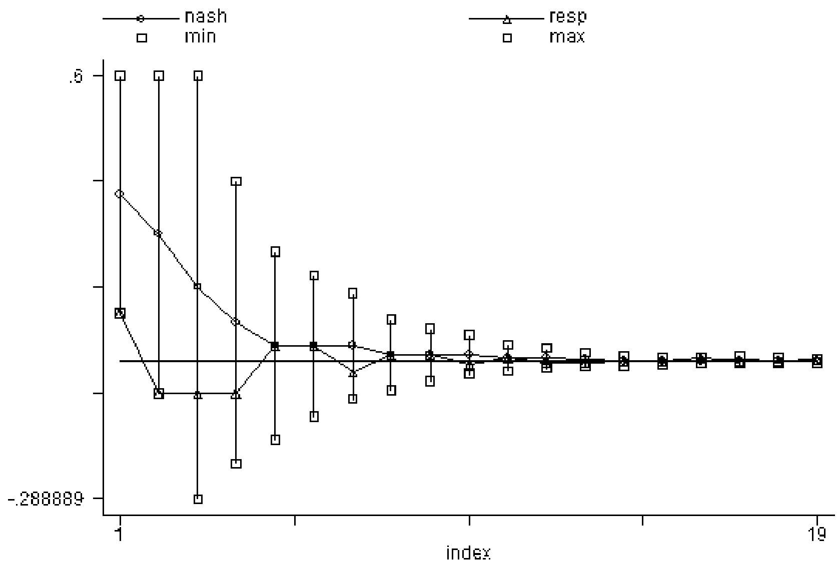

Figure 3.1: Convergence on the Vickrey payment rule

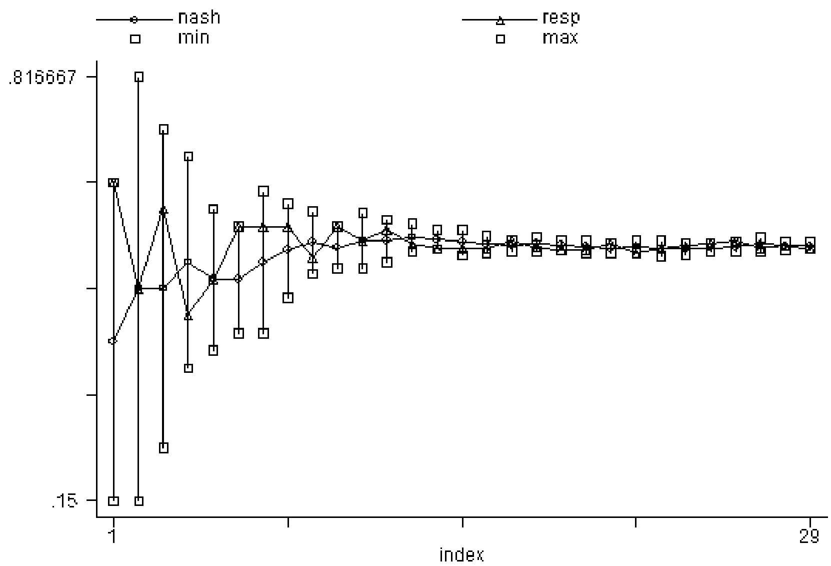

Figure 3.2: Convergence on the No Discount payment rule

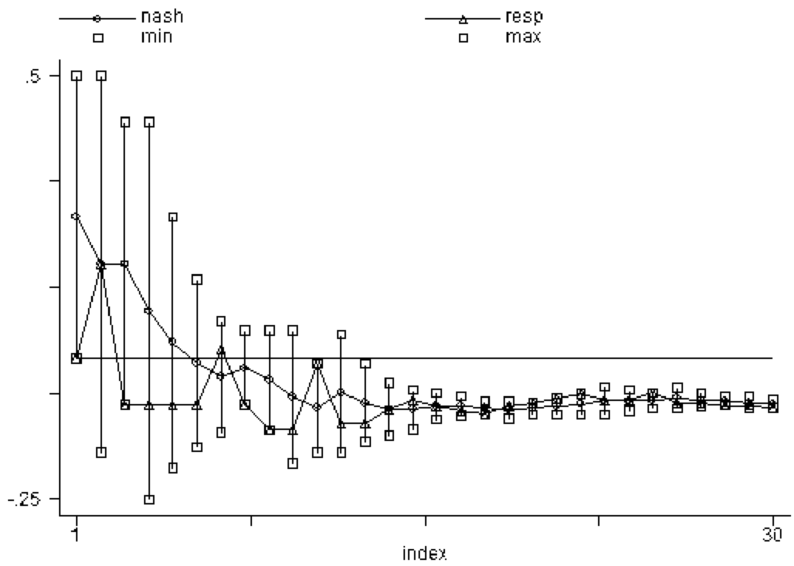

Figure 3.3: Convergence on the Large payment rule

where it crosses; otherwise, this is remarkable mostly in that it is unremarkable, and in that it found an unexpected behavior.

The other rules show similar convergence patterns. After we calculate the BNE for a given payment rule, the final step is to confirm the BNE by checking it against a discretized sample of the space. Specifically, every round hundredth value was checked to ensure that we were finding global, and not local, maxima of the best response function. This is thus as accurately as we can officially claim our algorithm to be. 3.4 graphs the responses that were found during the confirmation step of the Threshold payment rule calculation. The BNE in this case was to play strategy .234.

Timing considerations are also important; if our algorithm were no better than the method of enumerating all the possible strategies and their best responses, then we would have found little. Enumerating all the possible strategies and best responses on a very coarse grid (to the level of the tenth, from 0 to .8) for the Threshold rule took 5535 seconds (92 minutes) of CPU time. This is reasonable. However, to do so on a 50x50 grid (to move to the level of the hundredth, which we found with the search algorithm) would take, given that the overhead time is negligible, would

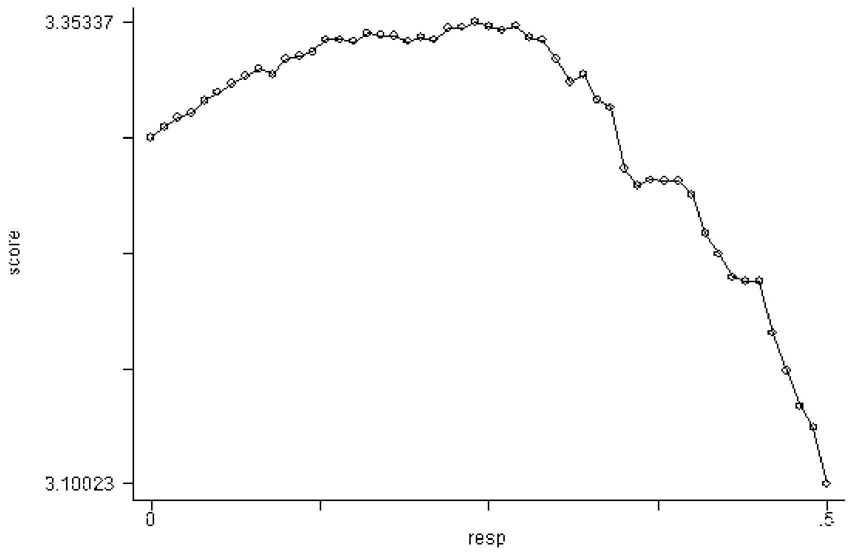

Figure 3.4: Confirmation of a Nash equilibrium for the Threshold rule

take approximately 2.5 *days*(199,000 seconds) of CPU time. This is compared to our algorithm, which returned the answer to the nearest hundredth in just 8,700 seconds (about 2.5 *hours* of CPU time.

Chapter 4

Economic Applications

In PKE, it was intuited that those rules which had maximum gains from trade with minimum manipulation would lead to more market efficiency, since the bids and asks input to the mechanism would be closer to those seen under the VCG framework. We expand on the work in PKE by quantifying the total gains from trade under BNE strategies for the various payment rules and comparing them to the gains from trade under the VCG mechanism (which, being provably efficient, gives the maximum of all gains from trade).

Having calculated the equilibrium behavior of the various agents, it is possible to analyze the efficiency of the setup of the combinatorial exchange. We assume that each player plays their *ex ante* Nash equilibrium strategy and determine the average allocative efficiency of each rule. We define the average allocative efficiency of a payment rule as:

$$E\left(\frac{\sum_{i=1}^{|I|} v_i(\lambda_i)}{\sum_{i=1}^{|I|} v_i(\lambda_{\text{vick},i})}\right)$$

where E denotes the expectation, λ_i is the trade that occurs under the equilibrium strategy for the rule, and $\lambda_{\text{vick},i}$ is the trade that would be made if all agents bid truthfully (as under the Vickrey payment rule).

The payment rules listed in Section 2.2 each define a game with a symmetric Bayesian-Nash equilibrium. We calculate these equilibria using the means described in the previous section. However, simply determining equilibrium behavior under a specific payment rule is not sufficient to determine the usefulness of that payment rule. To compare the expected allocative efficiency of the various rules, we run many markets under each payment rule assuming that players play according to equilibrium strategies. We then can calculate the result of the market, and determine the total gains from trade by totaling the agents' true values to their payments to or from the mechanism. The most efficient rule will have the highest expected gains from trade with, ideally, relatively low variance. That is, it will have both high average market efficiency and tend to be robust to different markets.

For these experiments, first, the adaptive search algorithm given in Section 3.1 was run for each payment rule to calculate a BNE. Having calculated the symmetric BNE for each payment rule, each equilibrium strategy was run on 2,500 randomly generated value distributions. As with the computational experiments, we assumed 5 buyers, 5 sellers, 20 goods, and 100 bids. Table 4 gives the major summary statistics from the experiment. In the table, Equilibrium is the equilibrium strategy played by an agent as defined in Section 2.2. Efficiency is the average allocative efficiency over the trials. We then provide 95% confidence intervals on the efficiency value–note however that this is based on the error induced by randomly selecting values, not on the rounding error from the algorithm. We also give the median and standard deviation of the ratio of the allocative efficiency of a given market under a payment rule to the efficiency of that same market under the VCG mechanism.

4.1 The Large Rule Anomaly

The immediately striking fact from Table 4 is that the Large payment rule–which generates the most efficient allocations of all the Vickrey-based payment schemes studied, allocating efficiently over 99% of all the surplus and achieving 100% of the available surplus nearly 80% of the time–has a Bayes-Nash equilibrium that is *negative*. In fact, what little inefficiencies it does have are a result of trades taking place that are

Rule	Equilibrium	Efficiency	95% conf low	95% conf high	median	std. dev.
Vickrey	0	1	1	1	1	0
Large	-0.098	0.992	0.9914	0.9935	1	.0267
Threshold	0.234	0.953	.9496	0.9561	0.992	0.0831
Fraction	0.294	0.915	0.9098	0.9195	0.958	0.1238
Reverse	0.436	0.811	0.8024	0.8189	0.867	0.2102
Small	0.460	0.784	0.7750	0.7928	0.843	0.2277
Equal	0.465	0.780	0.7712	0.7892	0.838	0.2300
No Discount	0.545	0.697	0.6858	0.7072	0.765	0.2718

Table 4.1: Efficiency data for the Vickrey-based payment rules.

economically inefficient, rather than the normal auction/exchange design inefficiencies caused by too few trades taking place. This implies that buyers in equilibrium overbid for items, and sellers ask for less than their own value for their items. This can be attributed to the incentives in the Large rule. Large gives Vickrey discounts to as many agents as possible, in order from largest Vickrey discount to smallest Vickrey discount. As a result, however, the winners in the exchange benefit unilaterally from bidding as high and asking as low as possible.

Proposition 1 *It can be in the best interest of agents in the mechanism using the Large payment rule to bid more than their true values for a bundle. More specifically, it is in the best interest of an agent to bid more than his value if* $b_i(\lambda) + C_l^* - \Delta_{vick,i} \geq v_i(\lambda)$.

The rule, as noted above, provides full Vickrey discounts to as many agents as possible, in descending order of the size of their discounts. Thus, suppose agent i is selected to trade (there exists bundle λ such that $x_i(\lambda) = 1$), but does not receive a Vickrey discount. Then $\Delta_{\text{vick},i} < C_l^*$. Holding the behavior of the other agents constant, it is then in the interest of i to raise his bid value, as this raises his Vickrey discount for λ. In particular, raising a winning bid (or lowering a winning ask) on a bundle does not affect the choice of the optimum allocation. We know that

$\Delta_{\text{vick},i} = V^* - V^*_{-i}$, where V^* is the total value all agents have for the allocation, and V^*_{-i} is the total reported value all agents would have for an allocation in which agent i did not participate. $V^* = \sum_{\lambda\in\Lambda, j\in I} b_j(\lambda)x_j(\lambda)$ (the sum of all the bids/asks for bundles given by those that were assigned them). Then after manipulation by overbidding or underasking, we have $\Delta'_{\text{vick},i} = V^{*'} - V^{*'}_{-i} = b'_i(\lambda) + \sum_{\lambda\in\Lambda, j\in I|j\neq i} b_j(\lambda)x_j(\lambda) - V^*_{-i} = b'_i(\lambda) - b_i(\lambda) + V^* - V^*_{-i}$. Thus, $\Delta'_{\text{vick},i} - \Delta_{\text{vick},i} = b'_i(\lambda) - b_i(\lambda)$. Thus i increases his apparent Vickrey discount directly with his bid. Because C^*_l increases only when a bid rises above it, agent i needs to increase his discount enough for his bid to rise above C^*_l. This gives $b'_i(\lambda) - b_i(\lambda) \geq C^*_l - \Delta_{\text{vick},i}$, and $b'_i(\lambda) \geq b_i(\lambda) + C^*_l - \Delta_{\text{vick},i}$. Since C^*_l cannot rise any faster than $b_i(\lambda)$ holding the other bids constant, it follows that for any value of b'_i meeting the above condition, including $b'_i | b'_i > v_i$, gives the optimal utility for agent i. Further, if $b_i(\lambda) + C^*_l - \Delta_{\text{vick},i} \geq v_i(\lambda)$, we have $b'_i(\lambda) \geq v_i(\lambda)$, and thus it is in the best interest of the agent to overbid.

The above, however, is an *ex post* calculation. In actual bidding, agents do not know the result of the auction, and thus do not know whether or not they will receive a particular bundle or not. The fact that the Large payment rule can lead to rational behavior of overbidding and underasking in certain circumstances would not matter very much if bidding in such a manner did not make sense on an *ex ante* basis. However, we see from the above experiments that the BNE of the market occurs where the agents deviate in this manner. This is troublesome; in particular, it means we cannot claim *ex post* individual rationality, though we can claim *ex ante* individual rationality.

Proposition 2 *Under the Large payment rule, if all players play the Bayes-Nash equilibrium strategy, some player will lose utility (unless the VCG rule is budget-balanced given the altered bids).*

Assuming that the allocation does not change, we know that not every player can receive a Vickrey discount (assuming the VCG rule is not budget-balanced given the altered bids). But then we know that there exists an agent in the allocation who does not receive a discount. But since that agent played the equilibrium strategy, he bid more than his value and asked less than his value for the bundles he received and sold. Thus, he loses utility.

Corollary 1 *The Large payment rule is not* ex post *individually rational.*

This follows directly from 2, since for a mechanism to be *ex post* individually rational implies that no agent can ever lose utility.

Proposition 3 *The mechanism with payments defined by the Large payment rule is* ex ante *individually rational.*

This follows directly from the fact that equilibrium behavior is the *ex ante* Nash equilibrium of the game. At any *ex ante* Nash equilibrium, we know that the expected utility of the agent must be greater at the Nash equilibrium than anywhere else, given other agents' strategies. We know that if the agent bid truthfully, the mechanism would have non-negative expected value for the agent, because the Large rule is *ex post* individually rational assuming no overbidding or under-asking [PKE01b]. Thus, the expected utility for an agent at any equilibrium point where agents overbid or under-ask must be greater than 0 in order to be a better response than truthfulness, so the mechanism is *ex ante* individually rational.

4.2 Other Rules

The other rules do not give such perverse incentives. In particular, the Threshold rule has a mean efficiency of about 95% of the allocative efficiency of the VCG mechanism, in spite of players altering their bids by more than 20%. Half the time it achieves more than 99% of the efficient allocation. This is quite a good system, particularly when compared with naive systems such as Equal (which divides the surplus among all the agents identically, regardless of the second-price system). Equal saw a mean of 78% of the gains from trade realized, and the median of the efficiencies of the individual market was just 84%.

The overall ordering of the rules makes intuitive sense, and our study confirms the economic findings of PKE using the Bayes-Nash solution concept. The rules are more efficient as they decrease the penalty to those who bid and ask far more than they need to to obtain or sell their bundles, thus cutting the penalty to bidding truthfully or near-truthfully with a high valuation. Threshold and Large each cap the losses to those who bid more or ask less than was required (in the case of Large, it not only caps the losses but rewards those who bid the farthest). Fraction was next; its discount is in proportion to the agent's Vickrey discount, and thus also aids those who overbid or ask too little more than others, but it has no cap on the surplus lost by bidding high and asking low. The other rules either reward bidding close to everyone else, or do not reward any bidder in particular over any other.

All told, the best rule of those studied appears to be Threshold. While it does not lead to outcomes as efficient as the Large rule, it maintains *ex post* budget balance, and one can still expect to achieve 95% of the ideal gains from trade. While Large is more efficient, if it is all right in a given setting to choose an *ex ante* individually rational mechanism for efficiency over an *ex post* individually rational mechanism for the sake of efficiency, the d'AGVA mechanism is available and is perfectly efficient [dG79].

Chapter 5

Conclusion

5.1 Future Work

Future work in this area can go in several directions. The FCC is considering *iterative* combinatorial exchanges– thus a question is whether some of these payment rules can be used or adapted to an iterated mechanism, and whether they have the same levels of efficiency on such exchanges. From both the game theoretic and mechanism design perspective, it would also be interesting to attempt to search on a larger space of strategies than the constant-factor strategy given above. Some ideas of ways to expand the space include allowing separate manipulation functions for bids and asks, or for those with high valuations or low valuations. Particularly interesting may be a study of expanding the space to add a constant factor of the size of the distribution from which values are drawn to the linear function given above, as Chatterjee and Samuelson (as discussed in Section 1.4.1) found that the use of such a function led to the optimum Nash equilibrium in the bargaining game.

5.2 Conclusion

In this thesis, we assess designs for a combinatorial exchange. In the process, we provide an algorithm which, while slow, is useable to determine the Bayes-Nash equilibria of such exchanges. We then used this algorithm to calculate the equilibria of different rules for solving the payment problem in such exchanges, and were then able to analyze the efficiency and property of the exchanges. We thus confirmed, using a stronger solution form, the efficiency results provided in PKE, and can confirm their recommendation on the Threshold rule. We reserve such recommendation, provided in PKE, on the large rule, as we have shown that such a rule encourages behavior that leads to the mechanism not being *ex post* individually rational, which was one of the assumptions that was made when the rule was created (indeed, PKE did not appear to consider the possibility of such perverse incentives). We also feel that our work encourages that of others recently to study continuous-space games using probabilistic and search methods.

Bibliography

[ATY00] Arne Andersson, Mattias Tenhunen, and Fredrik Ygge. Integer programming for auctions with bids for combinations. In *Proc. 4th International Conference on Multi-Agent Systems (ICMAS-00)*, 2000.

[BdVSV01] Sushil Bikchandani, Sven de Vries, James Schummer, and Rakesh V Vohra. Linear programming and Vickrey auctions. Technical report, Anderson Graduate School of Management, U.C.L.A., 2001.

[Cla71] E H Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[CS83] K Chatterjee and W Samuelson. Bargaining under incomplete information. *Operations Research*, 31:835–851, 1983.

[dG79] C d'Aspremont and L Gerard-Varet. Incentives and incomplete information. *J. of Public Economics*, 11:25–45, 1979.

[Fed01] Federal Communications Commission. *2nd FCC Combinatorial Bidding Conference*, http://wireless.fcc.gov/auctions/conferences/combin2001/, October 2001. Wye Woods Conference Center, MD.

[Gro73] Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[HPV89] M D Hirsch, C H Papadimitriou, and S A Vavasis. Exponential lower bounds for finding brouwer fixed points. *Journal of Complexity*, 5:379–416, 1989.

[ijc01] *Proc. 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, 2001.

[KP98] Vijay Krishna and Motty Perry. Efficient mechanism design. Technical report, Pennsylvania State University, 1998. Available at: http://econ.la.psu.edu/˜vkrishna/vcg18.ps.

[KW02] E Kwerel and J Williams. A proposal for a rapid transition to market allocation of spectrum. *Office of Plans and Policy Working Paper*, 38, 2002.

[LJ64] C E Lemke and J T Howson Jr. Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics*, 12:413–423, 1964.

[LLR89] W Leininger, P Linhart, and R Radner. Equilibria of the sealed bid mechanism for bargaining with incomplete information. *Journal of Economic Theory*, 48:63–106, 1989.

[MM96] Richard D McKelvey and Andrew McLennan. *Computation of Equilibria in Finite Games*, volume 1, pages 87–142. Elsevier Science BV, Amsterdam, The Netherlands, 1996.

[MS83] Robert B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.

[NR00] Noam Nisan and Amir Ronen. Computationally feasible VCG mechanisms. In *Proc. 2nd ACM Conf. on Electronic Commerce (EC-00)*, pages 242–252, 2000.

[PKE01a] David C. Parkes, Jayant R. Kalagnanam, and Marta Eso. Achieving budget-balance with vickrey-based payment schemes in combinatorial exchanges. Technical report, IBM Research Report RC 22218, 2001.

[PKE01b] David C Parkes, Jayant R Kalagnanam, and Marta Eso. Achieving budget-balance with Vickrey-based payment schemes in exchanges. In *Proc. 17th International Joint Conference on Artificial Intelligence (IJCAI-01)* [ijc01].

[RH95] Michael H Rothkopf and Ronald M Harstad. Two models of bid-taker cheating in Vickrey auctions. *Journal of Business*, 68(2):257–267, 1995.

[RW03] Daniel Reeves and Michael Wellman. Computing equilibrium strategies in infinite games of incomplete information (unpublished manuscript), 2003.

[Sca73] H Scarf. *The Computation of Economic Equilibria.* Yale University Press, New Haven, CT, 1973.

[SSGL01] Tuomas W Sandholm, Subhash Suri, Andrew Gilpin, and David Levine. CABOB: A fast optimal algorithm for combinatorial auctions. In *Proc. 17th Int. Joint Conf. on Artificial Intelligence* [ijc01].

[Vic61] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.