RC 22218 W0110-065 (10/23/2001)
Computer Science/Mathematics

IBM Research Report

Achieving Budget-Balance with Vickrey-Based Payment Schemes in Combinatorial Exchanges

David C. Parkes

Division of Engineering and Applied Science,
Harvard University,
33 Oxford Street,
Cambridge, MA 02138
parkes@eecs.harvard.edu

Jayant Kalagnanam

IBM Research Division,
T.J. Watson Research Center,
P.O. Box 218,
Yorktown Heights, NY 10598
jayant@us.ibm.com

Marta Eso

IBM Research Division,
T.J. Watson Research Center,
P.O. Box 704,
Yorktown Heights, NY 10598
martaeso@us.ibm.com

This page intentionally left blank.

Achieving Budget-Balance with Vickrey-Based Payment Schemes in Combinatorial Exchanges*

October 23, 2001

Abstract

Generalized Vickrey mechanisms have received wide attention in the combinatorial auction design literature because they are efficient and strategy-proof. However, it is well known that it is impossible for an exchange, with multiple buyers and sellers and voluntary participation, to be efficient and budget-balanced, even relaxing dominant strategy requirements. Except for special cases, a market-maker in an efficient exchange must make more payments than it collects. Taking a constructive approach, we clear exchanges to maximize reported surplus, and explore the efficiency effects of different budget-balanced payment rules. The payment rules are formulated to minimize the distance to Vickrey payments, under different metrics. Different rules lead to different levels of truth-revelation, and therefore efficiency. Experimental and theoretical analysis suggest a simple *Threshold scheme*, which gives surplus to agents with payments further than a certain threshold value from their Vickrey payments, has good properties. The scheme exploits agent uncertainty about bids from other agents to reduce manipulation opportunities.

1 Introduction

The participants in an exchange, or agents, can submit both *bids*, *i.e.* requests to buy items for no more than a bid price, and *asks*, *i.e.* requests to sell items for at least an ask price. Exchanges allow multiple buyers to trade with multiple sellers, with aggregation across bids and asks as necessary to clear the market. An exchange might also allow agents to express complements and substitutes through combinatorial bids and asks, together with logical conditions across bundles of different items. Combinatorial bids allow the agents to express synergies and substitutabilities between items. For example, an agent might want to buy bundle S_1 *and* sell bundle S_2, or perhaps sell S_3 *or* S_4 but not both, with no subsets of bundles acceptable for buys and no supersets acceptable for sells. Following the literature on combinatorial auctions [26, 10] we call this a combinatorial exchange.

Applications of combinatorial exchanges have been suggested to the optimization of excess steel inventory procurement [15], supply chain coordination [34], and to bandwidth exchanges [12]. The FCC has recently expressed interest in exploring two-sided market mechanisms for future wireless spectrum allocations, to enable a reallocation of licenses across both incumbents and new entrants as new wireless spectrum is introduced. One-sided combinatorial auctions are currently used for procurement and logistics problems [18], and we expect applications of combinatorial exchange technology within markets such as the Covisint B2B electronic marketplace.[1] Combinatorial exchanges are essential generalizations of one-sided combinatorial auction mechanisms.

The two core problems in an exchange are *winner-determination* (clearing), i.e. determining what is traded and by which agents; and *pricing*, i.e. determining the net payment to (or from) each agent when the exchange clears. The rules used to clear and price trades within an exchange impact the allocative-efficiency of the exchange, that is the value of trades across all agents. As an example, a mechanism that promotes truth-revelation from agents, with agents bidding their true costs and true values for bundles, elicits enough information to implement the efficient allocation.

Another key consideration in an exchange is one of *timing.* Typically, bid submission, clearing, pricing, and feedback to bidders are iterated in an exchange; with the exchange either clearing *continuously* whenever trade is possible or *periodically.* It is reasonable to expect that periodic clearing will boost the efficiency of a combinatorial exchange, allowing more opportunities for aggregations across multiple bids and asks.

In a non-combinatorial continuous double auction (CDA), such as the NYSE, bids and asks must be for a single type of item, and bids and asks are matched continuously. In the CDA the efficient trade can be priced in equilibrium with non-discriminatory item prices (everyone pays the same for the same item). The pricing problem in a combinatorial exchange is considerably more difficult in the following sense; even if we allow non-linear and

[1] www.covisint.com

discriminatory prices (not seller anonymous *or* buyer anonymous) there are some efficient trades that cannot be priced [5].

Useful economic properties of a mechanism that clears and prices an exchange include:

- **Budget-balance** (BB). The total payments received by the exchange from agents should be at least the total payments made by the exchange to agents.
- **Individual-rationality** (IR). No agent should pay more than its net increase in value for the items it trades.
- **Allocative-efficiency** (EFF). Trade should be executed to maximize the total increase in value over all agents.

In a robust implementation, these properties would be implemented in a *dominant strategy* equilibrium, such that every agent has a dominant bidding strategy whatever the preferences and strategies of the other agents. A weaker, but still useful, solution would implement these properties in a *Bayesian-Nash equilibrium*, such that every agent follows a bidding strategy that maximizes its expected utility given the bidding strategy of every other agent and distributional information about agent preferences. In addition to providing robustness to incorrect assumptions about the rationality and preferences of other agents, a dominant strategy implementation is useful computationally because agents can avoid game-theoretic reasoning about other agents [32].

Unfortunately, the well-known impossibility result of Myerson & Satterthwaite [23] demonstrates that *no* exchange can be efficient, budget-balanced, and individual-rational. This impossibility result applies even with quasi-linear utility functions[2] and Bayesian-Nash implementation. By the *revelation principle* [13, 21] (see Figure 1), this result holds for both direct, in which agents directly report their values for different trades, and indirect mechanisms, in which agents report value information indirectly, for example via responses to market prices. In addition, this result holds with or without *incentive-compatibility*, which states that the equilibrium strategy is for agents to reveal *truthful* information about their preferences. A perfect mechanism for the combinatorial exchange problem (and even the standard exchange problem) is impossible, even putting aside computational considerations.

Recognizing that combinatorial exchanges have great importance as intermediaries in electronic commerce, we take a constructive approach. We implement (BB) and (IR) as hard constraints, and select a clearing rule and payments to maximize efficiency, given rational participation by agents. We are interested in environments in which the market-maker must make a profit to operate the exchange, and in which participation is voluntary, so any reasonable mechanism must be (BB) and (IR).

[2] A quasi-linear utility function assumes that an agent's utility for trade λ is a function $u_i(\lambda, p) = v_i(\lambda) - p$, where $v_i(\lambda)$ is its value for the trade, and p is the price it must pay.

This is a paper about mechanism design, *not* about the computational complexity of clearing a combinatorial exchange. The clearing problem is NP-hard by reduction from the maximum weighted set packing problem, and approximations will be required in truly large instances, in turn affecting the efficiency and truth-revelation properties of the exchange. Special-cases of the winner-determination problem, for example in which agents can receive fractional assignments and with additional restrictions on allowable types of aggregation, can be solved in polynomial time [15]. In addition, observations made about the tractability of certain special-cases of the winner-determination problem in combinatorial auctions [26, 10, 19] carry over to exchanges. In one of the few published computational studies, Sandholm et al. [29] report that CPLEX 7.0 solves hard instances of combinatorial exchanges with multi-unit bids in less than 100 seconds for up to 10 items and 300 bids, but suggest that combinatorial exchanges may be harder to clear in practice than combinatorial auctions.

The mechanisms that we propose apply to a wide variety of combinatorial exchange settings, including exchanges with multiple units of the same items, divisible and indivisible bids, *ex ante* side constraints on feasible trades, different levels of aggregation, and with and without free-disposal. The main restriction is that we take allocative efficiency as our main objective, instead of some other objective such as trade volume.

We limit our attention to *direct mechanisms*, in which agents bids make claims about their values for different trades. It is useful to consider two different types of mechanisms.

(a) *Incentive-compatible mechanisms.* Design a budget-balanced and individual-rational combinatorial exchange in which truthful bidding is in equilibrium.

(b) *Non-incentive-compatible mechanisms.* Design a budget-balanced and individual-rational combinatorial exchange in which truthful bidding is not an equilibrium strategy for agents.

At first glance, mechanisms of kind (b) might seem redundant given the revelation principle, which states that any properties $\mathcal{P}$ that can be achieved with a non-incentive-compatible and perhaps indirect mechanism $\mathcal{M}'$ can be achieved with an incentive-compatible direct-revelation mechanism, $\mathcal{M}$ (see Figure 1). But the revelation principle makes a number of unreasonable computational assumptions. In essence, the revelation principle assumes that agents in the non-incentive-compatible mechanism, $\mathcal{M}'$, are able to compute equilibrium strategies. The new incentive-compatible direct-revelation mechanism, $\mathcal{M}'$, implements the same properties as $\mathcal{M}'$ by simulating the entire system and computing the equilibrium strategies for agents. In fact agents are bounded-rational, with costly and/or limited computation, and agents have only incomplete information about the preferences of other agents. A constructive application of the revelation principle to mechanism design makes a *worst-case* assumption about the strategic abilities agents, that they are able to exploit all opportunities for manipulation that exist in a mechanism, and designs for equilibrium

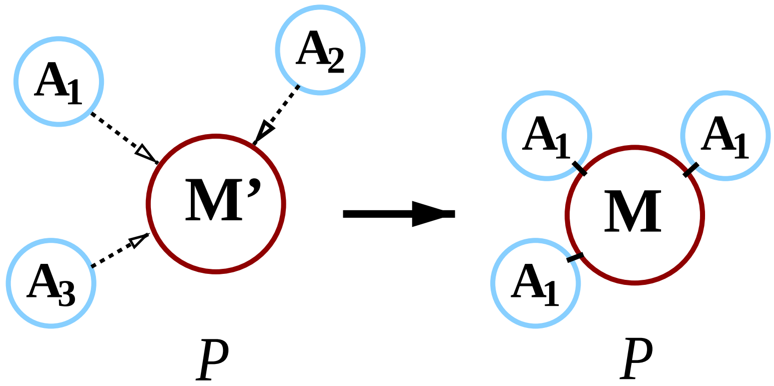

Figure 1: The revelation principle. Any properties $\mathcal{P}$ that can be implemented in Bayesian-Nash equilibrium in an indirect mechanism $\mathcal{M}'$ can be implemented in an incentive-compatible Bayesian-Nash equilibrium in a direct-revelation mechanism $\mathcal{M}$. The dashed arrows between agents $\mathcal{A}_i$ and mechanism $\mathcal{M}'$ indicate indirect revelation while the solid lines between agents and mechanism $\mathcal{M}$ indicate direct revelation.

strategies.[3]

There is a second more pragmatic reason to prefer mechanisms of kind (b) over those of kind (a). The impossibility theorem implies that (BB) and (IR) mechanisms of kind (a), that are incentive-compatible, must deliberately implement inefficient allocations even though the agents reveal enough information to implement an efficient allocation. This approach is adopted in earlier mechanism proposals for single-item exchanges [23, 20, 4], in which the exchange explicitly computes an inefficient outcome. In fact, these authors all propose mechanisms in which truth-revelation is a dominant strategy for agents, know as *strategy-proof* mechanisms.[4]

In comparison, (BB) and (IR) mechanisms of kind (b) *can* clear the exchange to maximize the *reported* values (or *surplus*) of agents. These mechanisms are not incentive-compatible and agents will not necessarily reveal truthful information in equilibrium. Clearing the exchange to maximize reported value does not necessarily implement the efficient allocation and does not violate the impossibility theorem. We believe that it is desirable to clear an exchange to maximize reported surplus, to avoid ex post claims from participants

[3]Note that this objection to the revelation principle is subtly different from the more standard computational objection, see for example Ledyard et al. [17] and Parkes [25]. The more standard observation objects to the use of the revelation principle to focus the design of mechanisms on direct revelation instead of indirect revelation, because the theoretical indirect-to-direct transformation ignores *costs* of complete information revelation (valuation and communication costs) and the computational cost on the mechanism of simulating the system and computing equilibrium outcomes.

[4]This idea of implementing inefficient allocations has a parallel in the *optimal* auction design literature, in which the primary goal is revenue-maximization instead of efficiency. An optimal auction must often be inefficient [22].

that an efficient trade was forfeited.

In this work we propose mechanisms of kind (b), and take the argument to its extreme. We *only* consider mechanisms that clear the exchange to maximize reported value. The mechanism design problem then reduces to designing a *payment scheme* that is budget-balanced and individual-rational, such that the incentives across agents are designed to minimize the benefits of strategic bids, and boost truth-revelation and efficiency. Although truth-revelation is not necessary for allocative-efficiency, if bids are truthful then the exchange implements the efficient allocation. The idea is to design a payment scheme that leverages the bounded-rationality of agents, making the holes for manipulation that must remain in the design difficult for agents to find and exploit.

In particular, we propose a number of budget-balanced variations on the Vickrey-Clarke-Groves (VCG) [33, 7, 14] mechanism for a combinatorial exchange. The VCG mechanism is a *strategy-proof* direct-revelation mechanism, such that it is a dominant strategy for an agent to bid its true values for different trades whatever the bids and preferences of other agents. The VCG mechanism is also (IR) and (EFF), but *not* budget-balanced. Enforcing budget-balance necessarily requires that we lose strategy-proofness, because we continue to clear the exchange to maximize reported surplus.

We interpret Vickrey payments as an assignment of *discounts* to agents after the exchange clears. Budget balance is achieved so long as the market maker distributes no more than the available surplus when the exchange clears.[5] The pricing problem is formulated as an optimization problem, to compute discounts to minimize the distance to Vickrey discounts. We derive payment schemes that correspond to optimal solutions to this problem under a number of different distance functions. The choice of payment scheme changes the incentives within the exchange for misrepresentation of agent values, and in turn influences the efficiency of the exchange.

Theoretical and experimental analysis compares the utility to an agent for misstating its value in bids and asks in each payment scheme across a suite of problem instances. The results, both theoretical and experimental, make quite a compelling argument for a simple *Threshold* payment scheme, that provides discounts to agents with payments more than a threshold distance than their Vickrey payments, and computes the threshold to provide budget-balance.

Intuitively, it is useful to think of payment rules in terms of the residual "degree of manipulation freedom" available to an agent, which is a measure of the best-case gain from manipulation to agent in a particular mechanism. High degree of manipulation freedom indicates that it is quite easy to successfully improve the outcome of a mechanism with non-truthful bidding. The Threshold rule sets the discounts to hide the low-hanging fruit, compensating agents for (ex post) easy opportunites, and leaving more difficult residual opportunities. It is useful to attack low risk manipulation opportunities because agents

[5] The market maker can also take a fraction of the surplus.

have uncertainty about the bids and asks of other agents and cannot fully exploit more difficult opportunities.

The work presented here is a promising first step towards designing "second-best" mechanisms for the combinatorial exchange problem. Work remains to be done to provide a full characterization of the equilibrium strategies under each rule, and to perform simulations with a richer set of agent strategies.

1.1 Example

Let us introduce an example problem, that we will return to later in the paper. Suppose agents 1, 2, 3, 4. Agents 1 and 2 want to sell A and B respectively, with values $v_1(A) = \$10$ and $v_2(B) = \$5$. Agents 3 and 4 want to buy the bundle AB, with values $v_3(AB) = \$51$ and $v_4(AB) = \$40$. The efficient allocation is for agents 1 and 2 to trade with agent 3, for a net increase in value of \$36.

The mechanism design problem is: given bid and ask prices for A, B and AB from the agents, what trades should take place and what payments should be made and received?

Here are a number of individual-rational[6] and budget-balanced payments:

– charge agent 3 its bid price, pay agents 1 and 2 their ask price.

– divide the surplus equally, charging agent 3 a price of \$39, and providing payments of \$22 and \$17 to agents 1 and 2 respectively.

– provide the surplus to the buyer, charging agent 3 a price of \$15 and providing payments equal to bid price to agents 1 and 2.

– divide the surplus across the sellers in proportion to their ask prices; i.e. agent 1 receives $\$10 + (10/15)\$36 = \$34$ and agent 2 receives $\$5 + (5/15)\$36 = \$17$.

Any analysis of the efficiency effect of these and other payment rules requires a consideration of the equilibrium effect of the payment rules on the bids and asks submitted by agents. In Section 4 we propose a set of budget-balanced payment rules, which are analyzed theoretically in Section 5 and experimentally in section 6.

2 Combinatorial Exchanges

The essential element that defines a combinatorial exchange is that the bids and asks of agents are expressive enough to describe values over *bundles* of items. This expressibility allows an agent to represent both complements and substitutes across items.

A simple combinatorial bidding language might allow agents to submit bids and asks explicitly for bundles of items. A *bid*, $B = (S, p_{\text{bid}})$, associates a *bid price* $p_{\text{bid}} \geq 0$ with a bundle of items $S \subseteq \mathcal{G}$, where $\mathcal{G}$ is the set of all items in the exchange. This is the most an

[6] Individual-rational for agents that bid truthfully, or at least for agents that shave bid prices down and ask prices up.

agent will pay for bundle S. An *ask*, $A = (S, p_{\text{ask}})$, associates an *ask price* $p_{\text{ask}} \geq 0$ with a bundle of items S. This is the minimum payment an agent will accept for bundle S. Such a bidding language can also be readily extended to allow agents to bid for *multiple units* of items.

Logical predicates can be introduced to connect these simple bidding elements. One example are "additive-or" connectors, that allow an agent to state that any number of a set of bids and/or asks can be selected simultaneously by the market-maker. Another example are "exclusive-or" connectors, that allow an agent to state that at most one of a set of bids and/or asks can be selected.

Stepping back from the specifics of the bidding language, any language will induce values over different *trades*. Let $\mathcal{I}$ denote the set of agents.

Definition 1 (Trade). *A trade $\lambda_i = (\lambda_i(1), \ldots, \lambda_i(|\mathcal{G}|))$, where $\lambda_i(j) \in \mathbb{Z}$ is an integer, defines a transfer of $\lambda_i(j)$ units of item $j \in \mathcal{G}$ to agent i if $\lambda_i(j) > 0$, and a transfer of $\lambda_i(j)$ units of item j from agent i if $\lambda_i(j) < 0$.*

Let Λ denote the set of all possible trades. Each agent can be engaged in both buying and selling items, so a trade allows the exchange of items in both directions, depending on the sign of $\lambda_i(j)$. We also allow trades of multiple units of the same good, i.e. $\lambda_i(j)$ is not restricted to just {-1, 0, 1} but can be any integer value.[7]

Bids and asks induce a *reported* value, $\hat{v}_i(\lambda_i) \in \mathbb{R}$, for every trade $\lambda_i \in \Lambda$.

Definition 2 (Valuation function). *A valuation function $\hat{v}_i(\lambda_i)$ denotes an agent's net value for the trade.*

Bids indicate positive value for buying a bundle of items, while asks indicate negative value for selling a bundle of items. This explicit representation of the value of a trade does not allow an agent to specify a preference for the *identity* of the buyer or seller with which it executes its trade. However this is without loss of generality, because we can introduce the identity of an agent into the good space with a "dummy item" appended to the bids and asks of that agent, essentially converting attributes into sets of products. Additional side constraints (e.g. a limit on the number of winners, on the volume of trade, etc.) can also be implemented directly within the clearing rules of the exchange.

Returning to our simple bidding language, suppose there are goods $\mathcal{G} = \{A, B, C\}$ and consider bid $(AB, 10)$ and ask $(C, 5)$ from agent 1. These bids induce valuation function, $\hat{v}_1([1,1,0]) = 10, \hat{v}_1([0,0,-1]) = -5, \hat{v}_1([1,1,-1]) = 5$. The values for other trades are constructed to be consistent with value $-\infty$ for selling anything other than item C, a zero value for buying $S \subset \{AB\}$, and no additional value for buying more than bundle AB.

[7]Although we refer to trades over integer values of items, all the mechanisms that are proposed in this paper extend immediately to problems with fractional allocations of items.

2.1 Clearing

As discussed in the introduction, exchanges can clear either periodically, with intervals between periods measured for example in time or in bid volume, or continuously, with any feasible trade cleared immediately. Periodic clearing probably increases allocative-efficiency in combinatorial exchanges, because it provides an opportunity to aggregate across a larger number of bids and asks. In addition, the payment schemes that we present in this paper are most applicable to periodic clearing, because that presents the most choice for surplus distribution. The NYSE is a continuous double auction, in which traders submit bids and asks for immediate execution. The Arizona Stock Exchange (AZX) is a periodic, or *call* market, in which bids and asks are accumulated and cleared at periodic, pre-specified intervals [31]. Call markets are also used to open sessions in continuous markets (e.g. the Bourse de Paris), and used for less active securities and bonds.

Bids may be *open* or *closed*. This refers to the information that is passed on to the bidders by the market maker. In an open bid market all agents may observe all the submitted bids (probably without the identity of the bidder), they may also learn what are the winning bids and how much the winners pay/ are paid. In a closed (or sealed) bid market the agents have no knowledge of the other agents' bids and learn only whether their own bids won or not and their own payment. There are many possibilities between these two extremes; for instance, submitted bids are not revealed to other agents, everyone is notified of his/her trade in the current provisional outcome and non-winners are provided with a minimum bid increment (or ask decrement) to enter the provisional trade if other bids remain unchanged. Open bids, and general market transparency, may encourage better coordination between agents, but also allow more manipulation in thin markets.

The NASDAQ is an example of an exchange with open bids, although trading information is not evenly distributed across participants; access to information about bids and asks waiting to clear (in the "limit book") is distributed among market brokers and held until the trades are cleared. A new plan, the Next Nasdaq, calls for the implementation of a centralized database of bids and asks that is visible to all market participants [31]. In typical procurement examples one would expect closed bids because the information contained in bids may well present useful information to competitors [9].

Following the discussion in the introduction, in this paper we are interested in mechanisms in which the overall trade that is implemented given a set of bids and asks is that which maximizes reported value (or surplus). We define this as the winner determination problem.

Definition 3 (winner determination problem). *The winner determination problem in a combinatorial exchange is to compute the overall trade that maximizes revenue given a set of bids and asks.*

Formally, let $\lambda = (\lambda_1, \ldots, \lambda_{|\mathcal{I}|}) \in \times_{i \in \mathcal{I}} \Lambda$ denote a complete trade across all agents. Given reported valuation functions, $\hat{v}_i(\lambda_i)$ for trades λ_i, a general mathematical programming formulation for the winner determination problem is:

$$\lambda^* = \arg \max_{\lambda \in \times_i \Lambda} \sum_{i \in \mathcal{I}} \hat{v}_i(\lambda_i)$$
$$\text{s.t.} \quad \mathit{feasible}(\lambda)$$

where *feasible*(λ) are a set of constraints to define whether or not the trade is feasible.

A feasible trade must respect supply and demand, for example without allocating the same item to more than one agent. The exact interpretation of *feasible* depends on, amongst other things:

- *Aggregation*: the role of the market-maker in disassembling and reassembling bundles of items.
- *Divisibility*: the ability to allocate fractions of items, and the ability to satisfy a fraction of agents' bids and asks. When an agent wants its bid or nothing, then its bid is called *indivisible*.
- *Ex-ante constraints*: for example, on the volume of trade, assignment constraints, market concentration, etc.
- *Free-disposal*: the ability to match a superset of a bundle of items with an ask, and/or a subset of a bundle of items with a bid.

In combination with divisibility, free-disposal allows a bundle of items to be matched with the maximal fraction of an agent's bid that remains a subset of the bundle of items. The payment rules introduced in this paper are valid in all of these cases.

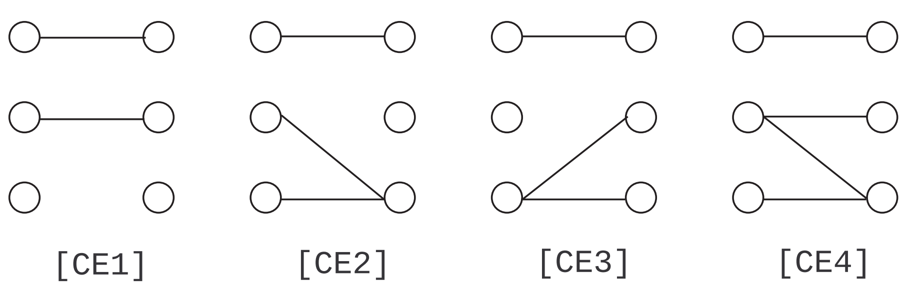

Figure 2: A Hierarchy of clearing models: [CE1] no aggregation or disaggregation; [CE2] buy-side aggregation; [CE3] sell-side aggregation; [CE4] complete aggregation.

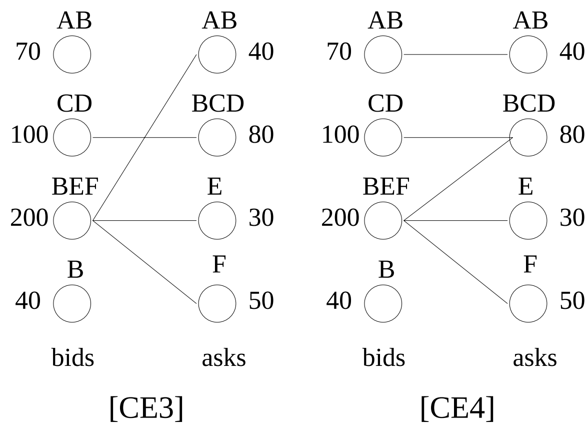

Figure 3: Example 1. Clearing in a combinatorial exchange: [CE3]: sell-side aggregation, $V^* = 100$; [CE4]: complete aggregation, $V^* = 170$.

Figure 2 illustrates a hierarchy of winner determination models, for different levels of auctioneer participation in clearing the exchange. The nodes represent the collection of bids and asks submitted by a single agent. Moving from [CE1] to [CE4] the auctioneer has a more active role, and the maximum possible efficiency of the outcome increases. For example, in [CE1] all trades are bilateral and there is no auctioneer intervention, while in [CE4] there is complete aggregation. The appropriate level of aggregation will depend on the *physical attributes* of the good; e.g. pieces of steel can be cut but not very easily joined (buy-side aggregation [CE2]), conversely computer memory chips can be combined but not split (sell-side aggregation [CE3]). Note that aggregation *does not* imply that the exchange must take physical possession of goods, trades can still be executed directly between agents.

Example

Figure 3 illustrates the trades executed for a simple problem, under aggregation models [CE3] and [CE4]. The exchange receives 4 bids to buy bundles, and 4 bids and 4 asks, and there are 6 items, $\{A, \ldots, F\}$ and no logical constraints across bids and asks. We assume free disposal, such that excess items in an ask can be stripped away at no cost, either by the seller, the buyer, or the exchange.

As expected, the surplus in [CE4] is greater than in [CE3], \$170 compared with \$100. For comparison, the solution for [CE2] is to match ask $(BCD, \$80)$ with bids $(CD, \$100)$ and $(B, \$40)$, for surplus \$90. The optimal bilateral match is ask $(AB, \$40)$ with bid $(AB, \$70)$

and ask $(BCD, \$80)$ with bid $(CD, \$100)$, for a surplus of \$50.

The computational complexity of clearing an exchange depends on the interpretation of feasibility, and in addition on whether bids and asks are single units of items or for multiple-units. The base case, with one item, divisible bids, aggregation possible on both sides of the market, and no constraints is polynomial time solvable by equating supply and demand; the problem remains polynomial time solvable with multiple items. Kalagnanam et al. [15] show that introducing assignment constraints does not make the problem computationally more difficult, the winner-determination problem can be formulated as a network-flow problem. In general, as soon as bids are indivisible the problem becomes NP hard by reduction from maximum weighted set packing. However, if there is no aggregation then the problem can be solved as a maximum weight matching problem in polynomial time irrespective of divisibility. If aggregation is possible on any one side of the exchange and there is at least one indivisible bid and/or ask the problem becomes NP-hard. The complexity analysis for variations that allow piecewise-linear demand and supply curves are discussed in Eso et al. [11] and Sandholm & Suri [28]. In addition, observations made about the tractability of certain special-cases of the winner-determination problem in combinatorial auctions [26, 10, 19] carry over to exchanges.

In practice, despite the NP-hardness of winner-determination in one-sided combinatorial auctions, off-the-shelf general purpose Integer Programming solvers, such as OSL or CPLEX, have been successful in solving quite large winner-determination problem in one-sided combinatorial auctions (with 1000's of items and bids), even for what are believed to be hard distributions [1, 28]. Moreover, recent studies on approximation algorithms (using an approximate linear-programming algorithm together with greedy hill-climbing heuristics) demonstrate solution times comparable with the time to send problem instances over a T1 line [37], and with average-case efficiency within 4% of optimal. In one of the few published computational studies on solving the winner-determination problem in combinatorial exchanges, in this case with full aggregation and indivisible bids, Sandholm et al. [29] conjecture that two-sided problems might be more difficult to solve in practice than one-sided problems.

We cautiously put the computational complexity of winner-determination to one side in this paper, and assume optimal solutions in describing or pricing mechanisms. Although this is reasonable for tractable special cases and small problem instances, approximations will be necessary in large problems. In general one must be careful when introducing approximate solutions within mechanisms, because the presence of an approximation algorithm can break the incentive compatibility of a mechanism [24]. This secondary effect on allocative-efficiency is less of a concern in this paper because we do not try to design fully incentive-compatible mechanisms.

2.2 Pricing

The pricing problem in an exchange is to determine the payments made by agents to the exchange and made by the exchange to agents after the exchange clears. Although in NYSE etc. the price is simply one of the bid or ask prices, in a combinatorial exchange it is useful to allow prices that are not simply the sum over all bids and asks in an agent's trade.

Definition 4 (payment). *The payment $p(i)$ to agent i is the price paid by the agent to the exchange if $p(i) > 0$, and the price paid by the exchange to the agent if $p(i) < 0$.*

As is very common in the literature on mechanism design and auction theory, we assume that agents have a quasi-linear utility function.

Definition 5 (quasi-linear utility). *A utility function $u_i(\lambda_i, p(i))$ is quasi-linear if $u_i(\lambda_i, p(i)) = v_i(\lambda_i) - p(i)$, where $p(i)$ is the agent's payment and $v_i(\lambda_i)$ is its value for trade $\lambda_i \in \Lambda$.*

Implicit in this assumption is that agents are *risk neutral*, with the same expected utility for a certain payment of \$10 and a trade with an expected value of \$10 and zero price.

As discussed in the introduction, the pricing problem is constrained by the characteristics that we require for the overall mechanism that is used to clear and price an exchange. For voluntary participation in an exchange in which agents report truthful values in equilibrium, we require an *individual-rationality* constraint, which states that agents have non-negative utility for participating in the exchange.

We adopt the following constraint for individual-rationality:

$$\hat{v}_i(\lambda_i^*) - p(i) \geq 0, \quad \forall i \in \mathcal{I} \qquad \text{(IR)}$$

where $\hat{v}_i(\lambda_i^*)$ is the reported value to agent i of the trade λ_i^* that is implemented when the exchange clears. In particular, if $\hat{v}_i(\lambda_i) > 0$ then the agent must pay the exchange less than its reported value, and if $\hat{v}_i(\lambda_i) < 0$ then the exchange must pay the agent more than its reported loss in value.

This constraint provides *ex post* individual-rationality for incentive- compatible exchanges in which agents reveal true values via bids and asks, because the agent has non-negative utility over all possible outcomes.[8] In fact, this constraint is too strong in non incentive-compatible mechanisms because agents' revealed values need not be their true value; instead, we might expect agents to shave down bid prices and shave up ask prices. We adopt the constraint because we design for the worst-case, to achieve ex post (IR) in the case in which agents' bids *are* in fact truth-revealing.

[8]Weaker notions of individual-rationality include *ex ante* in which an agent has non negative expected utility given distributional information about its own preferences, and *interim* individual-rationality in which an agent has non negative expected utility given distributional information about the preferences of other agents.

For *budget-balance*, we require:

$$\sum_{i \in \mathcal{I}} p(i) \geq 0 \tag{BB}$$

This is weak budget-balance, as we do not require that the market maker makes a zero surplus, just that the market maker does not make a loss. Again, this is *ex post* budget balance because we require budget-balance in all outcomes, not *ex ante* given distributional information about agent preferences and equilibrium strategies.

As an alternative, we find it useful to express an agent's payment, $p(i)$, as a *discount* from its reported value:

Definition 6 (Discount). *The discount Δ_i to agent i is the reported surplus to agent i at trade λ_i given payment $p(i)$, i.e. $\Delta_i = \hat{v}_i(\lambda_i) - p(i)$.*

Given this definition, equivalent definitions of (IR) and (BB) are:

$$\Delta_i \geq 0, \quad \forall i \in \mathcal{I} \tag{IR'}$$

and

$$\sum_{i \in \mathcal{I}} \hat{v}_i(\lambda_i) \geq \sum_{i \in \mathcal{I}} \Delta_i \tag{BB'}$$

In Section 4 we formulate the pricing problem as one as discount allocation across the agents that trade, subject to constraints (IR') and (BB'), for a number of different objective functions. The participation and budget-balance requirements leave a lot of flexibility over the choice of payment scheme.

3 Mechanism Design for Exchanges

In this section we first describe a straightforward application of the VCG mechanism to a combinatorial exchange. The VCG mechanism is (EFF) and (IR), and also strategy-proof, but as we show in a simple example is often not budget-balanced. In Section 3.2 we make a formal statement of the Myerson-Satterthwaite impossibility theorem for exchanges, and sketch a proof due to Krishna & Perry [16] that establishes impossibility of (EFF), (IR) and (BB) via uniqueness of Groves mechanisms and the revelation principle. We continue to characterize a number of sufficient conditions for (BB) of the VCG mechanism in special cases of the combinatorial exchange problem. In Section 3.4 we review mechanisms that have been proposed to address this impossibility for the standard non-combinatorial exchange.

3.1 The VCG Mechanism

In this section we describe an application of the Vickrey-Clarke-Groves mechanism [33, 7, 14] to a combinatorial exchange.

The VCG mechanism for an exchange computes the revenue-maximizing trade, λ^*, with all bids and asks, and also the revenue-maximizing trade $(\lambda_{-i})^*$, with each agent $i \in \mathcal{I}$ taken out of the exchange in turn. Let V^* denote the revenue from trade λ^*, and $(V_{-i})^*$ denote the revenue from the maximal trade without agent i.

The VCG mechanism implements trade λ^*, and computes the payment to each agent $i \in \mathcal{I}$ as:

$$p_{\text{vick},i} = \hat{v}_i(\lambda_i^*) - (V^* - (V_{-i})^*)$$

Negative payments $p_{\text{vick},i} < 0$ indicate that the agent *receives* money from the exchange after it clears. It is useful to define an agent's Vickrey discount. This is the amount an agent pays less than its reported value for the trade.

Definition 7 (Vickrey discount). *The Vickrey discount to agent i is $\Delta_{\text{vick},i} = V^* - (V_{-i})^*$.*

The Vickrey discount is always non-negative, providing smaller payments to the exchange for agents with reported *positive* net values for the trade, and larger payments from the exchange for agents with *negative* net values for the trade.

Proposition 1 (individual-rational). *The VCG mechanism is individual rational, such that the expected utility to rational agents from participation is also non-negative.*

Proof. It is easy to show that $V^* \geq (V_{-i})^*$ by a simple feasibility argument, so that the Vickrey discount is always non-negative. Truthful bidding is the rational strategy for an agent, so a rational agent will always pay less than its net value for the trade. □

In addition, truth-revelation is a dominant strategy equilibrium.

Proposition 2 (strategy-proof). *The VCG mechanism is strategy-proof.*

Proof. Consider the utility to agent i for reported valuation $\hat{v}_i$.

$$\begin{aligned} u_i(\hat{v}_i) &= v_i(\lambda_i^*) - p_{\text{vick},i} \\ &= v_i(\lambda_i^*) + \sum_{j \neq i} \hat{v}_j(\lambda_j^*) - (V_{-i})^* \end{aligned}$$

Ignoring the final term, which is independent of agent i's bid, agent i should announce a valuation $\hat{v}_i$ to maximize the sum of its own *actual* value for the trade $\lambda^* = (\lambda_1^*, \ldots, \lambda_I^*)$ and the *reported* value to the other agents. Given that trade λ^* is computed to maximize the total reported value of the agents, agent i should announce $\hat{v}_i = v_i$ to align the computation of the market maker in solving the clearing problem with its own interests. □

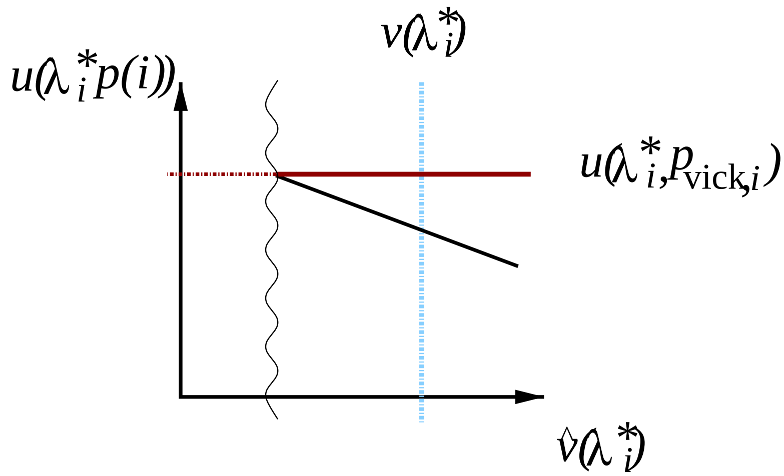

Figure 4: Agent utility against bid price, with all bids except λ_i^* truthful. To the left of the wavy line the maximal trade changes. The horizontal line is the agent's utility in the VCG mechanism, while the inclined line is an agent's utility in a No Discount mechanism.

The following proposition follows as an immediate corollary of strategy-proofness, because the trade λ^* is computed to maximize reported value.

Proposition 3 (efficient). *The VCG mechanism is efficient.*

Figure 4 and 5 provide some intuition about the strategy-proofness of the VCG mechanism. In Figure 4 we plot the agent's utility vs. its reported value on the efficient trade, λ_i^*, and compare the utility in the VCG mechanism with a "pay-what-you-bid", or *No Discount* mechanism. Notice that the agent's utility in the Vickrey mechanism is unchanged while its reported value is high enough to implement the efficient trade as the reported surplus maximizing trade. In comparison, in the No Discount mechanism the agent can maximize its utility by bidding just the minimal value that still implements the efficient trade.

This is generalized in Figure 5, in which we plot the agent's utility for different reported valuation functions. For simplicity we map all valuation functions into a single dimension. In this general case, for reported valuation functions too far from truthful, the utility in the VCG mechanism jumps down to lower plateaus as the implemented trade changes. As before, in the No Discount mechanism the agent can maximize its utility by reporting just the right value.

In the next section we show that no (EFF) and (IR) mechanism can also be (BB) with the Myerson-Satterthwaite impossibility result. The following section presents a number of special cases in which the VCG payments *are* budget balanced.

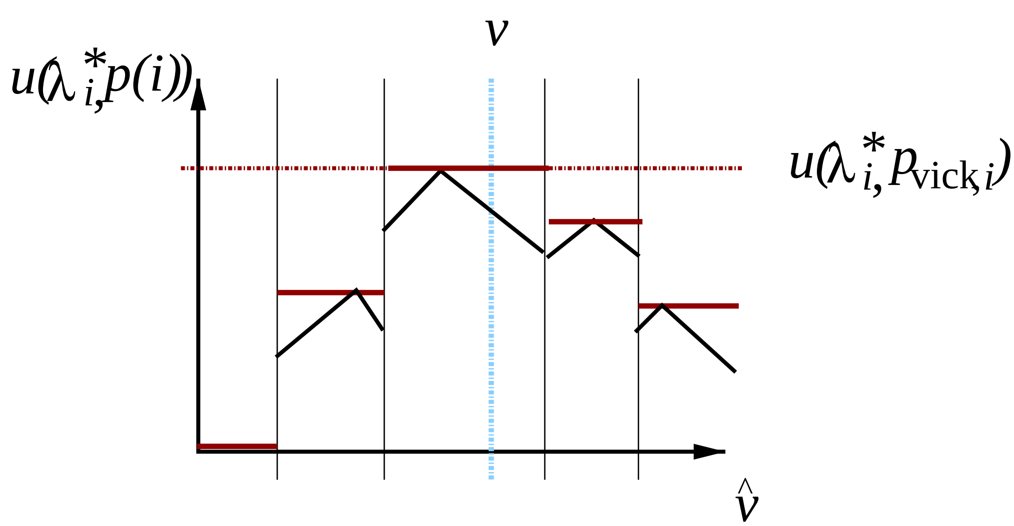

Figure 5: Agent utility against bid price, with all possible reported valuations mapped onto a single axis. The flat plateaus represent utility in the VCG mechanism; the jumps coincide with a change in the maximal trade. The inclined lines represent utility under a No Discount mechanism for different bids.

3.2 Impossibility Result

The Myerson-Satterthwaite impossibility theorem [23] shows that the joint goals of (BB), (IR) and (EFF) are essentially impossible to achieve. The formal statement of the impossibility result is for *ex ante* budget-balance (i.e. even allowing average-case budget-balance) and *interim* individual-rationality (i.e. even allowing average-case non-negative expected utility from participation).

The impossibility result extends via the revelation principle to indirect and non incentive-compatible mechanisms, and includes dominant-strategy implementations and ex-post individual-rationality as special cases.

Following Krishna & Perry [16], we find it convenient to prove the impossibility through a connection with the VCG mechanism. Let us state a special case of an important property of the VCG mechanism:

Theorem 1 (Payment Maximization). *Among all mechanisms that are efficient, interim individual-rational, and Bayesian-Nash incentive-compatible, the VCG mechanism maximizes the expected payments of each agent.*

Proof. See Krishna & Perry [16]. □

Again, this payment maximization property holds for indirect and non incentive-compatible mechanisms, via the revelation principle. It follows from Theorem 1 that there exists an ex-ante (BB), efficient, and interim individual-rational mechanism for an allocation problem *only* in the case that the VCG mechanism for the problem has non-negative expected revenue.

Theorem 2 (Myserson-Satterthwaite impossibility). *It is impossible to achieve allocative-efficiency, ex ante budget-balance and interim individual-rationality in a Bayesian-Nash incentive-compatible mechanism, even with quasi-linear utility functions.*

Proof. [16] Consider a situation of trade with a single seller (agent 1) and a single buyer (agent 2), and one indivisible item owned by the seller. Agent 1 has value $v_1([-1]) \in [0,1]$ for trade $\lambda_1 = [-1]$ and $v_1([0]) = 0$ for no trade. Agent 2 has value $v_2([+1]) \in [0,1]$ for trade $\lambda_2 = [+1]$, and $v_2([0]) = 0$ for no trade. As shorthand, let $v_1([-1]) = v_1$ and $v_2([+1]) = v_2$. To compute the expected revenue to the market-maker with the VCG mechanism consider the following two cases:

(a) $v_1 + v_2 \leq 0$. In this case $\lambda_1^* = \lambda_2^* = [0]$ and $p_{\text{vick},1} + p_{\text{vick},2} = 0$.

(b) $v_1 + v_2 > 0$. In this case $\lambda_1^* = [-1]$ and $\lambda_2^* = [+1]$, and $p_{\text{vick},1} + p_{\text{vick},2} = v_1 - (V^* - (V_{-1})^*) + v_2 - (V^* - (V_{-2})^*) = v_1 - (v_1 + v_2) + 0 + v_2 - (v_1 + v_2) + 0 = -(v_1 + v_2) < 0$.

Therefore, whenever there is a non-zero probability of gains from trade then the expected payment from agents is negative, the VCG mechanism runs at a deficit, and (ex ante) budget-balance fails. The Myerson- Satterthwaite impossibility theorem follows from the saliency of the VCG mechanism, as identified in Theorem 1. □

The impossibility result sets up an interesting second-best design problem. Exchanges, and in particular combinatorial exchanges, have many exciting applications that require mechanisms. But, any feasible design must relax one of (BB), (IR) and (EFF). As we discussed in the introduction, we require (IR) and (BB) and accept some loss in allocative-efficiency.

3.3 Vickrey Budget-Balance: Success & Failure

In this section we characterize conditions that are sufficient for budget-balance with Vickrey payments in an exchange. In fact, we will see that budget-balance failure is quite pervasive with Vickrey payments.

Example

Consider the Vickrey payments in the problem introduced in Section 1.1. In this case, assuming that agents submit these bids in the VCG mechanism, then we have $V^* = 51 - 10 - 5 = 36$, $(V_{-1})^* = (V_{-2})^* = 0$, $(V_{-3})^* = 25$, and $(V_{-4})^* = 36$. The Vickrey payment to agent 1 is $p_{\text{vick},1} = -10 - (36 - 0) = -46$, the Vickrey payment to agent 2 is $p_{\text{vick},2} = -5 - (36 - 0) = -41$, and the Vickrey payment to agent 3's is $p_{\text{vick},3} = 51 - (36 - 25) = 40$. The exchange runs at a loss of \$47 to the market maker.

Budget-balance requires that the sum of the Vickrey discounts across all agents is no greater than the available surplus when the exchange clears.

$$V^* \geq \sum_{i \in \mathcal{I}} (V^* - (V_{-i})^*) \qquad \text{(Vickrey-BB)}$$

The total marginal surplus contribution of each agent must be no greater than the total surplus of the agents taken together. This is a special case of the decreasing marginal returns requirement introduced by Bikchandani & Ostroy [5] to characterize combinatorial allocation problems in which Vickrey payments can be supported in competitive equilibrium.

Immediately, we have the following proposition:

Proposition 4. *If $(V_{-i})^* = 0$ for any agent $i \in \mathcal{I}$ then VCG payments are (BB) if and only if we also have $(V_{-j})^* = V^*$ for every other agent $j \neq i$.*

In other words, if there is one agent that is *critical*, in the sense that there are no revealed gains-from-trade in its absence, then there must be *perfect competition* across the other agents, in the sense that their presence has *no* marginal effect on the possible gains from trade.

It is interesting to strengthen (Vickrey-BB) by enforcing the following per-agent constraint:

$$v_i(\lambda_i^*) \geq V^* - (V_{-i})^*, \quad \forall i \in \mathcal{I} \qquad \text{(Vickrey-BB')}$$

In other words, the marginal contribution of each agent to the revealed surplus is less than its own value for the trade. It is immediate that condition (Vickrey-BB') implies condition (Vickrey-BB).

A necessary condition for this stronger budget-balance requirement is that:

$$v_i(\lambda_i^*) \geq 0, \quad \forall i \in \mathcal{I}$$

because $V^* - (V_{-i})^* \geq 0$. Now, suppose that Vickrey discounts are made to a subset $\mathcal{J} \subseteq \mathcal{I}$ of agents, while other agents receive no discount. Given that $v_i(\lambda_i^*) \geq 0$ for all agents, then [R1] is sufficient for budget-balance:

[R1] Agent j is a *buyer* but not a seller of items, and either free-disposal is possible by the market maker, or always possible by some agent $k \neq j$. This is sufficient because the remaining trades remain feasible at no loss in surplus, and $(V_{-j})^* \geq V^* - v(\lambda_j^*)$.

A decomposition technique is useful to characterize another budget-balanced special case. Suppose that $(\mathcal{C}_1, \ldots, \mathcal{C}_N)$ is a complete partition over agents $\mathcal{I}$. Weaker than (Vickrey-BB'), but still sufficient for budget balance is:

$$\sum_{i \in \mathcal{C}_k} v_i(\lambda_i^*) \geq \sum_{i \in \mathcal{C}_k} (V^* - (V_{-i})^*), \quad \forall k \in \{1, \ldots, N\}$$

Let us define a *safe partition*:

Definition 8 (safe partition). *A partition $(\mathcal{C}_1, \ldots, \mathcal{C}_N)$ is safe if the trades within each cluster $\mathcal{C}_k$ are feasible.*

In a safe partition, removing one or more agents within any one cluster only affects trade feasibility within that cluster, and does not have a ripple effect on the feasibility of agents elsewhere in the partition.

It is budget-balanced to allocate Vickrey discounts to agents in set $\mathcal{J} \subseteq \mathcal{I}$ in the following special case:

[R2] No two agents $l, m \in \mathcal{I}, l \neq m$ are in the same cluster, i.e. only one agent in each cluster receives a Vickrey discount.

To see this, notice that $(V_{\mathcal{I} \setminus \mathcal{C}_k})^* \geq V^* - \sum_{i \in \mathcal{C}_k} v_i(\lambda_i^*)$ because the trades executed by agents outside of the cluster remain feasible. It follows that $\sum_{i \in \mathcal{C}_k} v_i(\lambda_i^*) \geq V^* - (V_{\mathcal{I} \setminus \mathcal{C}_k})^* \geq V^* - (V_{-l})^* = \sum_{i \in \mathcal{C}_k} (V^* - (V_{-i})^*)$, where $l \in \mathcal{C}_k$ is the agent in the cluster that receives its Vickrey discount.

Finally, we can also recursively apply condition [R1] within a cluster in a safe partition, and apply Vickrey discounts to all buyers in a cluster so long as there is no agent in that cluster with negative value for its trade, and so long as there is free-disposal.

Consider the following applications of these two budget-balanced special cases:

The Generalized Vickrey Auction

In the Generalized Vickrey Auction, which is the VCG mechanism for a single-sided combinatorial auction problem, Vickrey discounts are allocated to all bidders in the auction but the seller receives the sum of the payments from the bidders and not her Vickrey payment.

This is budget-balanced by condition [R1] whenever there is free disposal, because the seller is not assumed to have a reservation price for the items. The budget-balance with this one-sided payment rule is quite tight; as soon as the seller receives her Vickrey discount then *every* buyer must pay its bid price.

Bilateral Matching

In a special case of a combinatorial exchange with bilateral matching, it is budget-balanced to provide at most one agent within each match with its Vickrey discount by case [R2], because the pairs of agents in matches define a safe partition.

Neither case is pertinent to our earlier example, of a combinatorial exchange problem in which VCG payments are not budget-balanced. Case [R1] is not useful because the sellers have negative values for the trade. Case [R2] does not add anything, because it is always budget-balanced to provide the Vickrey discount to at most one agent (Proposition 4), and the only safe partition in the example is the trivial partition into clusters $(\{1, 2, 3\}, \{4\})$.

3.4 Budget-Balanced Mechanisms for Single-item Exchanges

In this section we consider earlier proposals for budget-balanced mechanisms for single-item exchanges. We also include the VCG mechanism for completeness.

Figure 6 illustrates the clearing and payment problem in a single-item exchange, or double auction, with periodic clearing. Assume that bids are sorted in descending order, such that $B_1 \geq B_2 \geq \ldots \geq B_n$, while asks are sorted in ascending order, with $A_1 \leq A_2 \leq \ldots \leq A_m$. The efficient trade is to accept the first $l \geq 0$ bids and asks, where l is the maximal index for which $B_l \geq A_l$.

Table 1 surveys double auction mechanisms known in the literature. As expected by the Myerson-Satterthwaite impossibility theorem, no mechanism is (EFF), (BB) and (IR). All mechanisms except the VCG-DA are (BB) and (IR) but not (EFF), and all except the k-DA mechanism are strategy-proof.[9]

Name	traded	p_{buy}	p_{ask}	(EFF)	(BB)	(IR)	equil	(IC)
VCG-DA	l	$\max(A_l, B_{l+1})$	$\min(A_{l+1}, B_l)$	Yes	No	Yes	dom	yes
k-DA [6, 35]	l	$kA_l + (1-k)B_l$	$kA_l + (1-k)B_l$	No	Yes	Yes	Nash	no
TR-DA [3]	$l-1$	B_l	A_l	No	Yes	Yes	dom	yes
McAfee-DA [20]	l or $l-1$	$(A_{l+1} + B_{l+1})/2$ or B_l	$(A_{l+1} + B_{l+1})/2$ or A_l	No	Yes	Yes	dom	yes

Table 1: Double auction mechanisms. The *traded* column indicates the number of trades executed, where l is the efficient number. The *equil* column indicates whether the mechanism implements a dominant strategy or (Bayesian)-Nash equilibrium.

[9] Recently, Yoon [36] has proposed a modified version of the VCG-DA in which participants are charged a fee to enter the auction and balance the budget-loss of the VCG payments. Yoon characterizes conditions on agents' preferences under which the modified VCG-DA is (EFF), (IR) and (BB).

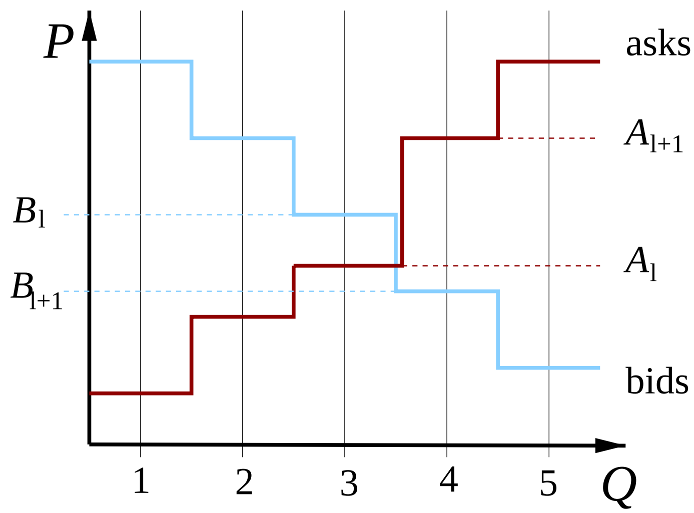

Figure 6: A standard double auction. Bid price vs. Ask price for successive quantities, $\mathcal{Q}$, of the same item. In this example it is efficient to trade $l = 3$ units, and there is a non-discriminatory clearing price $p \in [A_l, B_l]$.

In the k-DA [35, 6, 30], parameter $k \in [0, 1]$ is chosen before the auction begins; the parameter is used to calculate a clearing price somewhere between A_l and B_l. The McAfee DA [20] computes price $p^* = (A_{l+1} + B_{l+1})/2$, and implements this price if $p^* \in [A_l, B_l]$ and trades l units, otherwise $l - 1$ units are traded for price B_l to buyers and A_l to sellers. Babaioff & Nisan [3] have recently proposed the TR-DA rule, which is just the fall-back option of McAfee's DA.[10]

In the introduction we classified (BB) and (IR) mechanisms for exchanges into mechanisms of kind (a), that are (IC) and deliberately clear the exchange to implement an inefficient trade, and kind (b), that are not (IC) and clear the exchange to implement the revealed-surplus maximizing trade. Mechanisms TR-DA and McAfee-DA in Table 1 fall into kind (a), while mechanism k-DA falls into kind (b).

As a special case, we characterize conditions under which the VCG-DA is (weak) budget-balanced.

Theorem 3. *Budget-balance is achieved in a simple exchange for homogeneous items and single-item bids and asks if and only if one (or more) of the following conditions hold: (1)*

[10] Babaioff & Nisan also propose an α-reduction DA (not included in the table), in which a parameter $\alpha \in [0, 1]$ is selected before the auction begins. The TR-DA rule is used with probability α, and the VCG DA rule is used with probability $1 - \alpha$. Parameter α can be chosen to make the expected revenue zero (and achieve *ex ante* BB) with distributional information about agent values, to balance the expected surplus loss in the VCG-DA with expected gain in the TR-DA. The α-reduction DA is (BB) and (IR) but not (EFF), with dominant-strategy incentive-compatibility.

$p^0_{bid} = p^0_{ask}$; (2) $p^0_{bid} = p^{-1}_{bid}$; (3) $p^0_{ask} = p^{-1}_{ask}$.

Proof. BB holds if and only if $\max(p^0_{ask}, p^{-1}_{bid}) \geq \min(p^0_{bid}, p^{-1}_{ask})$, leading to cases: (1) $p^0_{ask} \geq p^{-1}_{bid}$ and $p^0_{bid} \leq p^{-1}_{ask}$; (2) $p^0_{ask} < p^{-1}_{bid}$ and $p^0_{bid} \leq p^{-1}_{ask}$; (3) $p^0_{ask} \geq p^{-1}_{bid}$ and $p^0_{bid} > p^{-1}_{ask}$. □

In other words, either one or more of the supply or demand curves must be "smooth" at the clearing point, with the first excluded bid at approximately the same bid price as the last accepted bid, *or* the winning bid and ask prices must precisely coincide.

4 Vickrey-Based Budget-Balanced Payment Rules

In the introduction we divided the space of (BB) and (IR) mechanisms into mechanisms of kind (a) and kind (b). Mechanisms (a) are incentive-compatible, and must explicitly implement an inefficient trade even though agents reveal enough information to compute the efficient outcome. Mechanisms (b) are non-incentive-compatible, and can clear the exchange to maximize reported surplus without breaking the impossibility theorem. We argued that these two approaches are fundamentally different with bounded-rational and poorly-informed agents, because the revelation principle implies that agents are able to implement non-truthful equilibria in non-incentive-compatible mechanisms.

Our thesis is that well-designed mechanisms of kind (b) may have better efficiency than any mechanism of kind (a) if they are successful in making it difficult for an agent to benefit from following non-truthful bidding strategies. Pragmatically, we further believe that it will often be unacceptable within actual exchanges, such as B2B markets or FCC spectrum wireless markets, to reduce trade and explicitly implement inefficient outcomes to achieve budget-balance.

We take an extreme version of design approach (b), and propose to clear the exchange to maximize reported surplus, which reduces the mechanism design problem to the pricing problem, which can be formulated as a discount allocation problem. We take (BB) and (IR) as hard constraints and propose methods to distribute surplus when an exchange clears to minimize the distance between discounts and Vickrey discounts. The choice of distance function has a distributional effect on the allocation of surplus and changes the incentive-compatibility properties of the exchange.

In this section we do the following:

- Formulate the pricing problem as a mathematical program, to minimize the distance to Vickrey payments with (BB) and (IR) as hard constraints.
- Introduce possible distance functions and construct corresponding budget-balanced payment schemes.
- Derive some theoretical properties that hold for the rules, and present an example application of the rules to the example first presented in the Introduction.

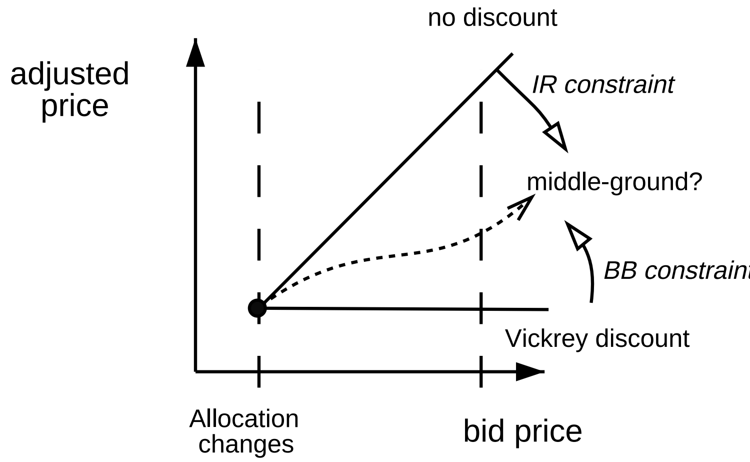

Figure 7: The discount allocation design space. With "no discount", an agent's adjusted price decreases as its bid approaches the smallest price that is still sufficient to implement the same trade. The Vickrey discount provides a flat adjusted price, that is the same for a range of bid prices.

Later, in Section 5 we present a theoretical analysis of each payment scheme in a simple bidding model, and in Section 6 we present an experimental analysis of each payment scheme within a slightly more realistic environment.

4.1 A Surplus-Allocation Model for the Pricing Problem

The pricing problem is to use the available surplus, V^*, computed at trade λ^*, to allocate discounts to agents that have good incentive properties while ensuring (IR) and (BB). Consider the graphical illustration in Figure 7. We must allocate the available discounts across agents that are less than the Vickrey discount in some places, to achieve (BB), and never negative to ensure (IR).

We formulate this problem as a linear program, to allocate surplus to agents to minimize distance to Vickrey discounts. Let V^* denote the available surplus when the exchange clears, before any discounts; let $I^* \subseteq \mathcal{I}$ denote the set of agents that trade. We compute discounts $\mathbf{\Delta} = (\Delta_1, \ldots, \Delta_I)$ to minimize the distance $\mathbf{L}(\mathbf{\Delta}, \mathbf{\Delta}_{\text{vick}})$ to Vickrey discounts, for a suitable distance function $\mathbf{L}$.

$$\min_{\mathbf{\Delta}} \quad \mathbf{L}(\mathbf{\Delta}, \mathbf{\Delta}_{\text{vick}}) \qquad \text{[PP]}$$

$$\text{s.t.} \quad \sum_{i \in I^*} \Delta_i \leq V^* \qquad \text{(BB')}$$

$$\Delta_i \leq \Delta_{\text{vick},i} \quad , \forall\ i \in I^* \qquad \text{(VD)}$$

$$\Delta_i \geq 0 \quad , \forall\ i \in I^* \qquad \text{(IR')}$$

Notice that the discounts are *per-agent*, not per bid or ask, and therefore apply to a wide range of bidding languages (as described in Section 2). Each agent may perform a number of buys and sells, depending on its bids and asks and the bids and asks of other agents.

Constraints (BB') and (IR') over the discount allocation were first introduced in Section 2.2. (BB') gives worst-case (or ex post) budget-balance, such that the exchange never makes a net payment to agents. This constraint can be easily strengthened if it is desirable for the market-maker to take a sliver of the surplus before computing the discounts. The (IR') constraints ensure that truthful bids and asks are (ex post) individual-rational for an agent, such that an agent has non-negative utility for participation whatever the bids and asks received by the exchange.[11]

Constraints (VD) ensure that no agent receives more than its Vickrey discount. The constraints are not redundant for certain distance metrics, such as the $\mathbf{L}_{\mathrm{RE}}(\cdot)$ metric.

In addition to the standard $\mathbf{L}_2(\boldsymbol{\Delta}, \boldsymbol{\Delta}_{\mathrm{vick}}) = \sum_i (\Delta_{\mathrm{vick},i} - \Delta_i)^2$ and $\mathbf{L}_\infty(\boldsymbol{\Delta}, \boldsymbol{\Delta}_{\mathrm{vick}}) = \max_i(\Delta_{\mathrm{vick},i} - \Delta_i)$ distance metrics, we also consider the following functions: (a) $\mathbf{L}_{\mathrm{RE}}(\boldsymbol{\Delta}, \boldsymbol{\Delta}_{\mathrm{vick}}) = \sum_i \frac{\Delta_{\mathrm{vick},i}-\Delta_i}{\Delta_{\mathrm{vick},i}}$, a relative error function; (b) $\mathbf{L}_\Pi(\boldsymbol{\Delta}, \boldsymbol{\Delta}_{\mathrm{vick}}) = \prod_i \frac{\Delta_{\mathrm{vick},i}}{\Delta_i}$, a product error function; (c) $\mathbf{L}_{\mathrm{RE2}}(\boldsymbol{\Delta}, \boldsymbol{\Delta}_{\mathrm{vick}}) = \sum_i \frac{(\Delta_{\mathrm{vick},i}-\Delta_i)^2}{\Delta_{\mathrm{vick},i}}$, a squared relative error function; and (d) $\mathbf{L}_{\mathrm{WE}}(\boldsymbol{\Delta}, \boldsymbol{\Delta}_{\mathrm{vick}}) = \sum_i \Delta_{\mathrm{vick},i}(\Delta_{\mathrm{vick},i} - \Delta_i)$, a weighted error function. The $\mathbf{L}_1$ metric provides no distributional information because any complete allocation of surplus is optimal.[12] These functions do not represent an exhaustive set of distance measures. For example, one could also consider them in combination with pre-selection, in which agents are dropped from consideration using an initial rule, and perhaps stochastically.

We might also substitute an expected surplus $\overline{V}^*$ for V^* and implement average-case (or ex ante) budget-balance. Although we do not explore *ex ante* BB here, it is possible that this weaker constraint will allow stronger incentive-compatibility properties.[13]

4.2 Payment Rules

Rather than solving problem [PP] directly, we can compute an analytic expression for the family of solutions that correspond to each distance function. Each family of solutions is a parameterized *payment rule*. The payment rules are presented in Table 2. We include the

[11] We choose to design for the truth-revealing equilibrium, so this is an appropriate representation of (IR). In fact, this is slightly pessimistic if agents are expected to shave their actual bids in practice.

[12] Divide by zero is avoided in all distance functions by dropping agents with $\Delta_{\mathrm{vick},i} = 0$, and simply setting $\Delta_i = 0$ for these agents.

[13] For example, an extension to the Groves mechanism, the *dAGVA* (or "expected Groves") mechanism, due to Arrow [2] and d'Aspremont and Gerard-Varet [8], demonstrates that it is possible to achieve (EFF) and *ex ante* (BB) with *interim* (IR), such that agents will vountarily participate if they must decide before they know their own preferences, in a Bayesian-Nash incentive-compatible mechanism.

Distance Function	Name	Definition	Parameter Selection
$\mathbf{L}_2$, $\mathbf{L}_\infty$	Threshold	$\max(0, \Delta_{\text{vick},i} - C_t^*)$	$\min C_t$ s.t. (BB')
$\mathbf{L}_{\text{RE}}$	Small	$\Delta_{\text{vick},i}$, if $\Delta_{\text{vick},i} < C_s^*$ 0 otherwise	$\max C_s$ s.t. (BB')
$\mathbf{L}_{\text{RE2}}$	Fractional	$\mu^* \Delta_{\text{vick},i}$	$\mu^* = V^* / \sum_i \Delta_{\text{vick},i}$
$\mathbf{L}_{\text{WE}}$	Large	$\Delta_{\text{vick},i}$, if $\Delta_{\text{vick},i} > C_l^*$ 0 otherwise	$\min C_l$ s.t. (BB')
$\mathbf{L}_\Pi$	Reverse	$\min(\Delta_{\text{vick},i}, C_r^*)$	$\max C_r$ s.t. (BB')
-	No-Discount	0	-
-	Equal	$V^* / \lvert I^* \rvert$	-

Table 2: Distance Functions, Payment Rules, and optimal parameter selection methods. Constraint (BB') states that $\sum_i \Delta_i^* \leq V^*$, and $|I^*|$ (used in the Equal rule) is the number of agents that participate in the trade.

Equal rule, which divides the available surplus equally across all agents $I^* \subseteq \mathcal{I}$ that trade, and the *No-Discount* or "pay-what-you-bid" rule.

Each payment rule is parameterized, for example the Threshold rule, $\Delta_i^*(C_t) = \max(0, \Delta_{\text{vick},i} - C_t)$, depends on the selection of Threshold paramater C_t. The final column in Table 2 summarizes the subproblem that must be solved to determine the optimal parameterization for each rule, e.g. the optimal Threshold parameter, C_t^*, is selected as the smallest C_t for which the discount allocation remains budget-balanced. The optimal parameter for any particular rule is typically not the optimal parameter for another rule. In the Appendix we derive the payment rules, and the methods to compute optimal parameterizations of each rule, using the methods of Lagrangian optimization.

Suppose that the agents are indexed such that their Vickrey discouts are decreasing:

$$\Delta_{\text{vick},I} \leq \Delta_{\text{vick},I-1} \leq \ldots \leq \Delta_{\text{vick},2} \leq \Delta_{\text{vick},1}$$

and add dummy points $\Delta_{\text{vick},0} = \infty$ and $\Delta_{\text{vick},I+1} = 0$, and index the interval $[\Delta_{\text{vick},k+1}, \Delta_{\text{vick},k}]$ by k. Here are the discount allocation policies implemented with each rule, given optimal parameterizations C_l^*, C_s^*, μ^*, C_l^*, and C_r^*:

- **Threshold**. If parameter C_t^* falls into interval k, then agents $i = 1, \ldots, k$ will receive discount $\Delta_i^*(C_t^*) = \Delta_{\text{vick},i} - C_t^*$, and agents $i = k+1, \ldots, I$ will not receive any discounts.

- **Small**. If parameter C_s^* falls into interval k, then agents $i = 1, \ldots, k$ will not receive any discount while agents $i = k+1, \ldots, I$ will receive their Vickrey discounts.

- **Fractional**. Every agent i receives discount $\mu^* \Delta_{\text{vick},i}$, where $\mu^* = V^* / \sum_i \Delta_{\text{vick},i}$;

i.e. agents receive a fraction of the surplus equal to their proportional share of total discount under the VCG mechanism.

- **Large**. If parameter C_l^* falls into interval k then agents $i = k+1, \ldots, I$ will not receive any discounts while agents $i = 1, \ldots, k$ will receive their Vickrey discounts.

- **Reverse**. If parameter C_r^* falls into interval k then agents $i = 1, \ldots, k$ will receive their Vickrey discount while agents $i+1, \ldots, I$ receive a discount in the amount of C.

Let us sketch the construction of the Threshold rule from distance metric $\mathbf{L}_2$. Introducing Lagrange multiplier, $\lambda \geq 0$, we have $\min \sum_i (\Delta_{\text{vick},i} - \Delta_i)^2 + \lambda(\sum_i \Delta_i - V^*)$, s.t. $0 \leq \Delta_i \leq \Delta_{\text{vick},i}$. Now, computing first derivatives with respect to Δ_i and setting to zero, we have $-2(\Delta_{\text{vick},i} - \Delta_i^*) + \lambda = 0$ for all i, where Δ_i^* is the optimal allocation of discount to agent i.[14] Solving, this equalizes the difference between Vickrey discounts and actual discounts across all agents with $\Delta_i^* > 0$, i.e. $\Delta_{\text{vick},1} - \Delta_1^* = \Delta_{\text{vick},2} - \Delta_2^* = \ldots$. Parameter $C_t \geq 0$ denotes this difference, and gives budget balance for $C_t^* = (\sum_{i=1}^{K} \Delta_{\text{vick},i} - V^*)/K$, where index K is such that $\Delta_{\text{vick},K+1} \leq C_t^* \leq \Delta_{\text{vick},K}$. Every agent with Vickrey discount greater than parameter C^* receives a discount $\Delta_{\text{vick},i} - C_t^*$, while the other agents receive no discount.

Example

In Table 3 we compare the payments made with each payment scheme in the simple problem that we presented in the introduction. The payments are all (IR) and (BB) (except for the Vickrey mechanism), but each rule allocates a different discount to each agent.

Rule	Vick	Equal	Frac	Thresh	Reverse	Large		Small	
Agent 1	-46	-22	-25.6	-28	-22.5	-46	or -10	-35	or -10
Agent 2	-41	-17	-20.6	-23	-17.5	-5	-41	-5	-30
Agent 3	40	39	46.2	51	40	51	51	40	40

Table 3: Payments with Different Rules in the Simple Problem.

One thing to notice is that the Large and Threshold rules allocate no discount to the successful buyer, agent 3, and divide all surplus across the sellers. In comparison, the Equal rule actually provides more than the Vickrey discount to agent 3, a problem which is fixed with Reverse. In contrast with Large and Threshold, the Small and Reverse rules allocate the maximal discount to the one buyer and divide the remaining surplus across the two sellers.

[14]First-order conditions are necessary and sufficient for optimality in this problem because the Hessian is positive definite.

It is also interesting that in this simple example neither the Large or Small schemes provide useful guidance about how to distribute the discount across the two sellers because this depends on how the tie is broken.

4.3 Graphical Representations

In Figure 8 we present a graphical illustration of the effect of each payment rule on the adjusted price for different levels of bid prices. Let λ_i^* denote the trade to agent i when it bids truthfully, i.e. with $\hat{v}_i = v_i$, and suppose that the agent's reported value, $v = v_i(\lambda_i^*)$ for the trade is positive. Consider the effect of bidding $b_i \neq v_i(\lambda_i^*)$ on trade λ_i^*, but leaving the reported values for other trades truthful. Let $x = p_{\text{vick},i} = v - \Delta_{\text{vick},i}$, denote a critical value, such that: $V^* - (v - x) = (V_{-i})^*$, and any bid price $b_i < x$ for trade λ_i^* must change the surplus maximizing trade. We assume that $(\lambda_{-i})^*$ is the maximal trade for bid prices less than x.

To keep things simple, assume that trade λ^* continues to solve the winner-determination problem for all bids down to x:

$$\lambda^* = \arg \max_{\lambda \in \times_i \Lambda} \left[\hat{v}_i(\lambda_i) + \sum_{j \neq i} v_j(\lambda_j) \right]$$
$$\text{s.t.} \quad \textit{feasible}(\lambda), \quad \text{and } b \geq x$$

where $\hat{v}_i$ is truthful except for $\hat{v}_i(\lambda_i^*) = b$.

Finally, we will choose to ignore for the moment any effect that a change in bid price can have on the optimal parameterization for a particular payment rule, or the available surplus.

Figure 8 illustrates the adjusted payment to agent i under different rules, as its bid price b changes. For values $b < x$, the payment is zero under all rules because the agent is not in the trade. For $b \geq x$, the Vickrey discount is $\Delta_{\text{vick},i}(b) = V^*(b) - (V_{-i})^*$, where $V^*(b)$ is the surplus from the maximal trade with bid b. Simplifying, $\Delta_{\text{vick},i}(b) = V^* - (v - b) - (V^* - (v - x)) = b - x$. Clearly, the adjusted payment, $p(b) = b - \Delta_{\text{vick},i} = x$, under the Vickrey rule is independent of the bid price. Briefly, here are the derivations of the plots for each rule:

(a) **Fractional**. $p(b) = b - \mu^* \Delta_{\text{vick},i} = b - \mu^*(b - x) = b(1 - \mu^*) + \mu^* x$.

(b) **Equal**. $p(b) = b - V^* / \; |I^*| = b - D$

(c) **Reverse**. $p(b) = b - \min(b - x, C_r^*)$, and $p(b) = b - C_r^*$ if $b > C_r^* + x$, or $p(b) = x$ otherwise.

(d) **Large**. $p(b) = b - (b - x) = x$, if $b > C_l^* + x$, or $p(b) = b$, otherwise.

(e) **Small**. $p(b) = b$, if $b > C_s^* + x$, or $p(b) = x$ otherwise.

(f) **Threshold**. $p(b) = b - \max(0, b - x - C_t^*)$, and $p(b) = x + C$ if $b > x + C_t^*$, or $p(b) = b$ otherwise.

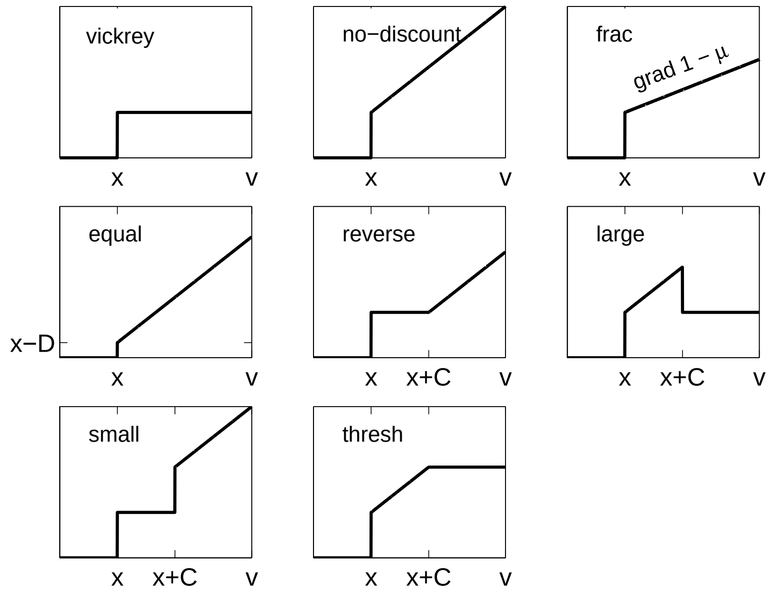

Figure 8: Bid price b_i on trade λ_i^* against adjusted bid price $p(b)$ in each payment scheme. Agent i has value $v > 0$ for λ_i^*, and value x is the agent's Vickrey payment, also the smallest bid price at which trade λ_i^* is executed. The parameters, C, indicate the optimal selections for each rule, and need not be the same across rules; similarly parameter μ is the optimal selection for the Fractional rule. Parameter D in Equal is a share $D = V^* / \, |I^*|$ of the available surplus.

We might conjecture from the strategy-proofness of the VCG mechanism that flat regions in plots of adjusted price vs. bid price lead to good incentive-compatibility properties. Of course, because VCG payments are not budget-balanced we must accept some none flat regions. In Section 7 we draw some conclusions about good positioning of these flat regions, based on the theoretical and experimental analysis of the different Vickrey-based payment rules.

4.4 Theoretical Results

We will characterize some conditions under which an agent's strategy is truth-revelation, given the bids and asks submitted by other agents. Although not a full equilibrium analysis, because the other agents' bids are fixed, we do worry about the effect that an agent's bids can have on the available surplus to allocate as discounts, and in turn on the optimal parameters in each rule. Our analysis will be pessimistic, in that we assume that an agent has perfect information about the bids and asks of other agents.

To keep things simple, we consider the situation only for buyers, that have a positive value for the trade computed to clear the exchange. Let the *maximal trade without agent j* refer to the reported surplus maximizing trade with agent j taken out of the exchange. Let θ denote the amount by which agent i reduces its bid price on trade λ_i^*, i.e. $\hat{v}_i(\lambda_i^*) = v_i(\lambda_i^*) - \theta$; θ can be positive or negative. Consider the following possible assumptions that we might make about bids $\hat{v}_i$ submitted by the agent:

[C1] The maximal trade λ^* with truthful bids, $\hat{v}_i$, from agent i is still implemented with untruthful bids $\hat{v}_i \neq v_i$.

[C2] The maximal trade λ_i^* to agent i is also implemented in the maximal trade without agent j, for all $j \neq i$, for truthful and untruthful bids from agent i.

[C3] In the case that $\theta < 0$, and agent i submits a higher bid, the reported value for trades, $\lambda_i \neq \lambda_i^*$, increases by at most $-\theta$.

[C4] There exists a higher bid $\hat{v}_i(\lambda_i^*)$, for some $\theta < 0$, such that λ_i^* is in the maximal trade without agent j, for all agents $j \neq i$.

It is convenient to denote the new values for surplus, Vickrey discounts, and optimal rule parameters at bids $\hat{v}_i \neq v_i$, as $V^*(\theta)$, $\Delta_{\text{vick},i}(\theta)$, and $C_t^*(\theta)$, etc.

Lemma 1. *Given any bids and asks from other agents, the maximal utility to an agent, for any bid and in any Vickrey-based payment rule, is the utility that it would receive with truthful bidding in the VCG mechanism.*

Proof. This follows immediately from the (VD) constraint, that $\Delta_i \leq \Delta_{\text{vick},i}$ in all payment rules. □

Note that this Lemma is valid irrespective of whether we have truth revelation from the other agents. It is always the case that for any set of bids and asks the agent can do no better than the utility it would receive from truthful bidding in a VCG mechanism.

This allows a number of immediate claims about the optimality of truth-revelation. We do not need to make any assumptions to prove that the agent receives its VCG payment in each of the following special cases when bidding truthfully; the optimality of truth-revelation follows from the Lemma.

Proposition 5. *Truthful bidding is the utlity-maximizing strategy for an agent in Large whenever* $\Delta_{\text{vick},i} \geq C_l^*$.

Proposition 6. *Truthful bidding is the utility-maximizing strategy for an agent in Small whenever* $\Delta_{\text{vick},i} \leq C_s^*$.

Proposition 7. *Truthful bidding is the utility-maximizing strategy for an agent in Reverse whenever* $\Delta_{\text{vick},i} \leq C_r^*$.

The subtle problem with making strong statements about the incentive-compatibility of the Threshold rule for agents that are in the "flat region", and receiving discounts, is that as an agent changes its bids the optimal C_t^* parameter can change to maintain budget-balance. An agent can reduce its adjusted payment, without bidding low enough to achieve $\Delta_{\text{vick},i} < C_t^*$, by lowering the threshold C_t^*. Recall that $C_t^* = (\sum_{i=1}^{K} \Delta_{\text{vick},i} - V^*)/K$, where index K is such that $\Delta_{\text{vick},K+1} \leq C_t^* \leq \Delta_{\text{vick},K}$.

We have the following positive results:

Theorem 4. *An agent that receives a discount in Threshold cannot benefit from submitting a lower bid, with $\theta > 0$, under assumptions [C1] and [C2], while $\theta \leq \Delta_{\text{vick},i} - C_t^*$.*

Proof. We show that the optimal Threshold does not change, and that the agent's Vickrey discount remains to the right of the Threshold. The new Vickrey discount $\Delta_{\text{vick},i}(\theta) = \Delta_{\text{vick},i} - \theta$ by assumption [C1], and $\Delta_{\text{vick},i}(\theta) > C_t^*$ because $\theta \leq \Delta_{\text{vick},i} - C_t^*$. Agent i's discount is θ less than with truth-revelation, exactly compensating for the θ drop in surplus to distribute. In addition, $C_t^*(\theta) = C_t^*$, because the Vickrey discounts are unchanged for all agents $j \neq i$: we have $\Delta_{\text{vick},j}(\theta) = (V^* - \theta) - ((V_{-j})^* - \theta) = \Delta_{\text{vick},j}$ by [C2]. □

Theorem 5. *An agent in Threshold cannot benefit from submitting a higher bid $\hat{v}_i(\lambda_i^*)$, with $\theta < 0$, under assumptions [C1] and [C3].*

Proof. By [C1], $\Delta_{\text{vick},i}(\theta) = \Delta_{\text{vick},i} + \theta$, and there is a neutral effect on the numerator of C_t^* due to the change in surplus and discount to agent i. The discounts to other agents $j \neq i$ can only increase, because $V^*(\theta) = V^* + \theta$ by [C1], and $(V_{-j})^*(\theta) \leq (V_{-j})^* + \theta$ by [C3]. Therefore, C_t^* can increase but not decrease, leaving an agent with $\Delta_{\text{vick},i} > C_t^*$ at best indifferent. In addition, C_t^* cannot increase by more than θ, and bring agent i into the Threshold region, because at best $\Delta_{\text{vick},j}(\theta) - \Delta_{\text{vick},j} \leq \theta$ for every agent j with $\Delta_{\text{vick},j} \geq C_t^*$. □

Turning to Large, we show that it is vulnerable to agents submitting very large bids:

Theorem 6. *An agent in Large that receives no discount with truthful bids can achieve its Vickrey utility by submitting a higher bid, $\hat{v}_i(\lambda_i^*)$, for some finite $\theta < 0$, under assumptions [C1] and [C4].*

Proof. First, $\Delta_{\text{vick},i}(\theta) = \Delta_{\text{vick},i} + \theta$, by assumption [C1]. Then, for some $\theta < M$, for some finite $M > 0$, trade λ_i^* is implemented by agent i in the maximal trade without agent $j \neq i$, for all j by assumption [C4]. Writing $\theta = M + \theta_1$, then $V^*(\theta) = V^* + M + \theta_1$ and $(V_{-j})^*(\theta) \geq (V_{-j})^* + \theta_1$, so that $\Delta_{\text{vick},j}(\theta) \leq V^* + M - (V_{-j})^* = \Delta_{\text{vick},j} + M$. Therefore, there is some finite θ for which $\Delta_{\text{vick},i}(\theta) > \Delta_{\text{vick},j}(\theta)$ for all $j \neq i$, and agent i has the largest Vickrey discount. In this state, if any agent receives its Vickrey discount with Large it will be agent i, and there is always enough surplus to provide the Vickrey discount to any one agent (Prop. 4). □

5 A Simple Analytic Study of Manipulation In Each Payment Rule

In this section we present results for a simple analytic model of agent strategies given payment rules. The analysis assumes an exogeneous amount of surplus is available to allocate as a discount to an agent (independent of its own bids), and compares the optimal amount of manipulation by an agent with each payment rule under ex ante budget-balance. Interestingly, we find that the choice of payment rule can have a significant effect on an agent's optimal level of manipulation, even in the case that the agent receives the same level of expected discount under each scheme.

Although not a full equilibrium analysis, because we do not consider the effect that an agent's bid can have on either the bids of other agents, or on the available surplus, the analysis provides useful insight into the mechanism driving incentives across the schemes. In addition, the results are somewhat validated through consistency with the results of experimental simulation that we describe in the next section.

5.1 The Analytic Model

We focus on a single agent i within an exchange, and consider the effect that misrepresentation of its value can have on its utility. Decision analysis leads to a relationship between agent manipulation and expected surplus in each scheme. To keep things simple, we focus on a single possible trade λ_i^*, with positive value, and assume that agent i bids truthfully for every other possible trade.[15] Let $v = v_i(\lambda_i^*)$ denote the agent's true value, and $b = \hat{v}_i(\lambda_i^*) = v - \theta$ denote the agent's bid for trade λ_i^*, for level of manipulation θ.

The *critical value*, x, captures the idea that there is a smallest bid price that the agent can bid and still be in the maximal trade.

Definition 9 (critical value). *Critical value, x, is the minimal bid, b, for which the trade λ_i^* is selected in the maximal trade.*

By case analysis, when the agent trades λ_i^* in the maximal trade with a *truthful* bid, then $x = p_{\text{vick},i} = v - (V^* - (V_{-i})^*)$, and $x \leq v$. We also allow for the case that the agent needs to bid higher than its value to be successful. Then: $x = v - ((V_{+i})^* - V^*)$, where we use $(V_{+i})^*$ to denote the value from the maximal trade with truthful bids that *must* allocate trade λ_i^* to agent i.

We make the following assumptions:

[A1] the critical value x is uniformly distributed about v, i.e. $x \sim U(v - \delta, v + \delta)$, for some constant $\delta > 0$.

[A2] agent i does not know the actual value of x, but only the distributional information.

[15] Note, this is a slight abuse of notation because trade λ_i^* is not required to be in the maximal trade when agent i bids truthfully (although this is the interesting case).

[A3] the available surplus to allocate as discount to agent i is exogeneous, equal to $\alpha\delta$, where $\alpha > 0$ is a constant.

[A4] *ex ante* (BB) is sufficient, such that the rule is budget-balanced if the expected payment given the distribution of critical value is no greater than the available surplus.

Average case budget-balance is useful in this theoretical analysis because it allows the budget-balanced parameter for each payment rule to be fixed and known to the agent before it submits its bid.

Notice that we do not model the equilibrium effect of an agent's own bidding strategy on the distribution over the critical value x; neither do we model the equilibrium effect of an agent's own bidding strategy on the available surplus to allocate as discounts. This said, we do not need to make any distributional assumptions about an agent's valuations.

In Section 5.2 we determine the optimal bidding strategy of agent i in each payment rule, as a function of the parameter of the rule, such as C_t in Threshold, or μ in Fractional. In Section 5.3 we consider the expected discount to the agent with this optimal bidding strategy, and then compute the optimal manipulation under each payment rule when the parameters are selected to give ex ante budget-balance. The analysis leads to a relationship between the available surplus, $\alpha\delta$, and the manipulation, $\theta_i^*(\alpha)$, selected by the agent in each rule.

Throughout this section we choose to illustrate the analysis with respect to the Threshold rule. We omit the analysis of the other rules, and just give the results of the analysis in tabular form.

5.2 Determining the Optimal Manipulation

Manipulation by an agent has two effects on the expected utility for an agent: (i) the probability of the adjusted bid being accepted decreases, and (ii) the total utility if the bid is accepted can go up because the agent's payment might be reduced. Payment rules change (ii) but not (i), and in turn effect agents' bids and the efficiency of the exchange.

In Figure 9 we plot the *utility* for a particular bid, $b = v - \theta$, as the critical value x varies. Recall that an agent's utility $u_i(\theta)$ for a bid $b = v - \theta$ is

$$u_i(\theta) = \begin{cases} v - p, & \text{if } v - \theta > x \\ 0, & \text{otherwise} \end{cases}$$

where p is the payment by the agent to the exchange when its bid is accepted. Each subplot in Figure 9 is for a different payment scheme, and each line withint a subplot corresponds to the utility to an agent for a different bid as the critical value x changes. We assume in this plot that the agent's value is $v = 1$, and consider critical values $0 \leq x \leq 1$ and bids with $\theta \in \{0, 0.3, 0.5\}$. The Threshold parameter, $C_t = 0.4$, and the Fractional parameter,

$\mu = 0.5$. Although not plotted here, the curves for Equal are similar to the No-Discount case (except shifted higher in utility by a constant amount), and Large is similar to Threshold.

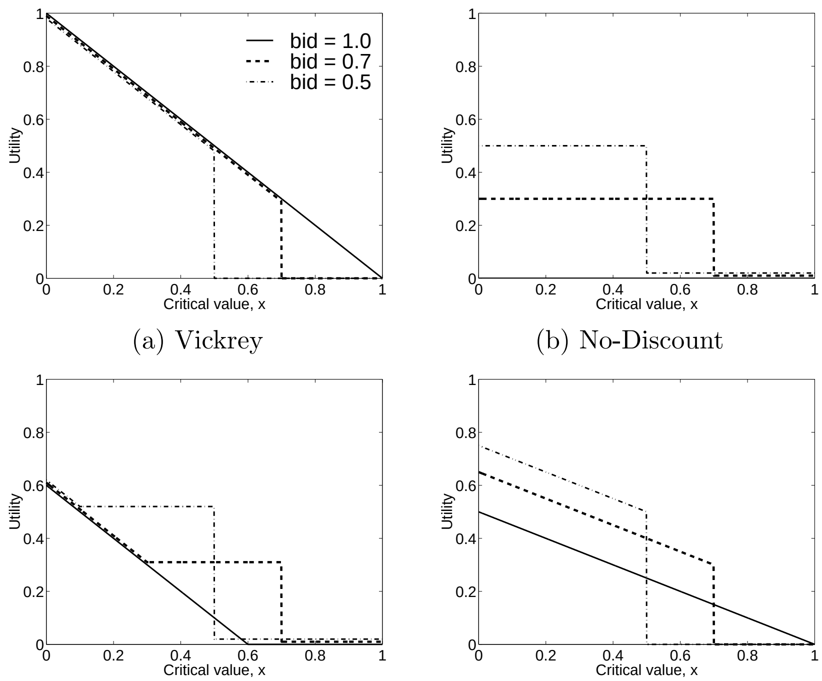

(a) Vickrey (b) No-Discount

(c) Threshold (d) Fractional

Figure 9: Utility of bids $b = v - \theta$ with $\theta \in \{0, 0.3, 0.5\}$, $v = 1$, as the critical value x varies between 0 and 1. $C_t = 0.4$ in Threshold, and $\mu = 0.5$ in Fractional.

Let $\Delta_{\text{vick},i}(b)$ denote the Vickrey discount to agent i given bid b.

$$\Delta_{\text{vick},i}(b) = \begin{cases} b - x, & \text{if } b \geq x \\ 0, & \text{otherwise} \end{cases}$$

In Threshold the discount rule is $\Delta(i) = \max(0, \Delta_{\text{vick},i}(b) - C_t)$, for parameter C_t, where $\Delta_{\text{vick},i}(b)$ is the discount computed using the Vickrey payment rule. The agent's payment is 0 if $b < x$, because it does not trade. On the other hand, if $b \geq x$, we know that $\Delta_{\text{vick},i}(b) = b - x$, and the agent's payment is $b - 0$ if $b \leq x + C_t$, or $b - (b - x - C_t) = x + C_t$ if $b > x + C_t$. Putting this together, the agent's utility, $u_i(\theta, C_t, x)$, given bid $b = v - \theta$ is:

$$u_i(\theta, C_t, x) = \begin{cases} v - (x + C_t), & \text{if } v - \theta \geq x + C_t \\ v - (v - \theta), & \text{if } x + C_t > v - \theta \geq x \\ 0, & \text{otherwise} \end{cases}$$

Looking at Figure 9, in the Vickrey scheme a lower bid reduces the agent's expected utility because it decreases the probability of success without increasing the utility of a successful bid. In comparison, with No Discount, the agent gains utility on all successful bids by the amount of deviation from truthful bidding. In the Threshold scheme a lower bid only reduces the price paid for a limited range of critical values (closer than C_t to the bid price), while in the Fractional scheme a lower bid reduces the price paid on all successful bids (but by less than in the No Discount scheme).

Continuing, we compute the expected utility, $\mathrm{E}u_i(\theta, C)$, in each payment rule as a function of θ and the parameter C in the rule (μ in Fractional). By assumption [A1], the critical value x is uniformly distributed about v. The expected utility for each level of manipulation θ can be computed as the area under a particular curve in a plot like Figure 9. In turn, the expected-utility maximizing bid corresponds to the curve with maximum area.

Returning to Threshold, assume that $C_t \leq \delta$, so that the agent will receive a discount for some choice of $\theta < \delta$. Let $f(x)$ denote the probability density function for critical value x, and consider three cases:

Case 1 ($0 \leq \theta \leq \delta - C_t$) In this case, for almost truthful bids, the agent's Vickrey discount can be both above and below the Threshold parameter C_t, depending on x:

$$\begin{aligned}
&\mathrm{E}u_i(\theta, C_t) = \\
&\int_{x=v-\delta}^{v-\theta-C_t} [v-(x+C_t)] f(x)\mathrm{d}x + \int_{x=v-\theta-C_t}^{v-\theta} [v-(v-\theta)] f(x)\mathrm{d}x + \int_{x=v-\theta}^{v+\delta} 0 f(x)\mathrm{d}x \\
&= \frac{(\delta-\theta-C_t)}{2\delta}(v-C_t) - \frac{1}{4\delta}\left[(v-\theta-C_t)^2 - (v-\delta)^2\right] + \frac{\theta}{2\delta}(\theta + C_t - \theta)
\end{aligned}$$

Case 2 ($\delta - C_t < \theta \leq \delta$) In this case the agent's bid is far enough away from truthful that its Vickrey discount is never above the Threshold parameter C_t.

$$\mathrm{E}u_i(\theta, C_t) = \int_{x=v-\delta}^{v-\theta} [v-(v-\theta)] f(x)\mathrm{d}x = \theta(\delta - \theta)/2\delta$$

Case 3 ($\delta < \theta$) In this case the agent's bid is never greater than the critical value, and $\mathrm{E}u_i(\theta, C_t) = 0$.

Next, we compute the agent's *expected-utility maximizing* bidding strategy, denoted $\theta_i^*(C_t)$. By assumption [A2], the agent itself cannot know the exact critical value x, but only the distributional information about x. Differentiation of $\mathrm{E}u_i(\theta, C_t)$ w.r.t. θ, and then case analysis, we have:

$$\theta_i^*(C_t) = \min\left[C_t,\ \delta/2\right]$$

As special cases, when $C_t = \delta$ (no discount), we have $\theta_i^*(C) = \delta/2$, and with $C_t = 0$ (Vickrey discount), we have $\theta_i^*(C) = 0$.

In Table 4 we present the expressions for the expected utility maximizing strategy in each payment rule, as a function of the parameter δ in the distribution of the critical value, and the parameterization of each rule. It is useful to confirm that all expressions reduce to that for the Vickrey and No-Discount rules at extreme parameter values; e.g. $\mu^* \in \{0, 1\}$ in Fractional, $C_t^* \in \{0, \delta/2\}$ in Threshold, etc.

Rule	Optimal Manipulation, $\theta_i^*(C)$	Expected Discount, $E\Delta_i(\theta^*(C), C)$
No-Discount	$\delta/2$	0
Vickrey	0	$\delta/4$
Fractional	$\max\left[0, \left(\frac{1-\mu}{2-\mu}\right)\delta\right]$	$\min\left[\delta/4, \frac{\delta\mu}{4(2-\mu)^2}\right]$
Threshold	$\min\left[C_t, \delta/2\right]$	$\max\left[0, \frac{(\delta-2C_t)^2}{4\delta}\right]$
Equal	$\frac{\delta-D}{2}$	$\frac{D(\delta+D)}{4\delta}$
Small	$\max\left[0, \min\left(\delta/2, \delta - C_s\right)\right]$	$\min\left[\delta/4, C_s^2/4\delta\right]$
Large	0, if $C_l \leq \delta/\sqrt{2}$	$-C_l^2/4\delta + \delta/4$, if $C_l \leq \delta/\sqrt{2}$
	$\delta/2$, otherwise	0, otherwise
Reverse	$\max\left[0, \frac{\delta-C_r}{2}\right]$	$\min\left[\delta/4, C_r/4\right]$

Table 4: Manipulation and Expected Discount in the Analytic Model.

5.3 Balancing Expected Discount with Surplus

In order to compare the optimal manipulation that is selected by an agent in each payment rule, we compute the expected discount that is allocated to the agent in each rule, and carefully select rule parameters to equalize the expected discount across the rules. Interestingly, we show that the rules have quite different effects on the agent's optimal level of manipulation, even with the same expected discount to an agent.

First, we must compute the *expected discount* allocated to agent i in each payment scheme, as the parameter in the scheme varies, and assuming that the agent follows its optimal bidding strategy (as computed in Table 4).

Returning to the Threshold example, given bid θ, critical value x, and parameter C_t, the discount in the Threshold rule to agent i is 0 when $b < x$, and $\max(0, \Delta_{\text{vick},i}(b) - C_t)$ otherwise. Combining, we can write the discount $\Delta_i(\theta, C_t, x)$ as:

$$\Delta_i(\theta, C_t, x) = \max[0, v - \theta - (x + C)]$$

Continuing, we can compute the *expected* discount, $E\Delta_i(\theta, C_t)$, wrt the distribution over the critical value x. First, in the case that $\theta \leq \delta - C$, we have:

$$\begin{aligned} \mathrm{E}\Delta_i(\theta, C_t) &= \int_{x=v-\delta}^{v-\theta-C_t} [v - \theta - (x + C_t)] f(x) \mathrm{d}x \\ &= (v - \theta - C_t) \frac{(\delta - \theta - C_t)}{2\delta} - \frac{1}{4\delta} \left[(v - \theta - C_t)^2 - (v - \delta)^2 \right] \end{aligned}$$

In the case of $\theta > \delta - C_t$, we have $\mathrm{E}\Delta_i(\theta, C_t) = 0$.

We can now substitute for the agent's optimal bidding strategy, $\theta_i^*(C_t)$, to compute the expected discount, $\mathrm{E}\Delta_i(\theta_i^*(C_t), C_t)$, as a function of rule parameter C_t:

$$\mathrm{E}\Delta_i(\theta_i^*(C_t), C_t) = \max \left[0, \frac{(\delta - 2C_t)^2}{4\delta} \right]$$

In Table 4 we present the expected discount in each payment rule, as a function of the parameter in the rule and the parameter δ of the distribution over the critical value.

Let us now compare how well each payment rule provides incentives to minimize agent manipulation. Figure 10 plots the expected *gain* in utility (in comparison with truthful bidding), $\mathrm{E}u_i(\theta) - \mathrm{E}u_i(0)$, in each payment rule. The parameters in each rule are carefully tuned to ensure that the expected discount to the agent in each rule, when the agent follows its optimal strategy, is the same. In this case we set the parameters to give the expected discount in each rule equal to 0.1δ. Notice that the level of manipulation, θ_i^*, that maximizes the agent's gain in utility is smallest in the Threshold scheme for this value of surplus, with Large not far behind.

Finally, we can understand the effect that the choice of payment rule has on agent's optimal manipulation levels for different levels of expected discount. It is interesting to assume that there is an average surplus of $\alpha\delta$ available to allocate as discount to agent i, and compare the rules for $\alpha \in [0, 1]$. This is assumption [A3].

Returning again to the Threshold rule, we need to compute the parameter $C_t^*(\alpha)$, for surplus $\alpha\delta$, that minimizes agent manipulation but ensures budget-balance; i.e. compute the minimal C_t such that $\mathrm{E}\Delta_i(\theta_i^*(C_t), C_t) \leq \alpha\delta$. By Assumption [A4] we allow *ex ante* budget-balance, such that the discounts balance the surplus on average over the distribution of critical values, x. In the case of Threshold, the solution is:

$$C_t^*(\alpha) = \max \left[0, \frac{\delta}{4} \left(2 - \sqrt{16\alpha} \right) \right]$$

$C_t^*(\alpha)$ is the Threshold parameter that minimizes manipulation and maintains budget-balance. With $\alpha \geq 1/4$, then $C^*(\alpha) = 0$ and the rule implements Vickrey discounts, while for $\alpha = 0$, then $C^*(\alpha) = \delta/2$ and the rule implements no discounts.

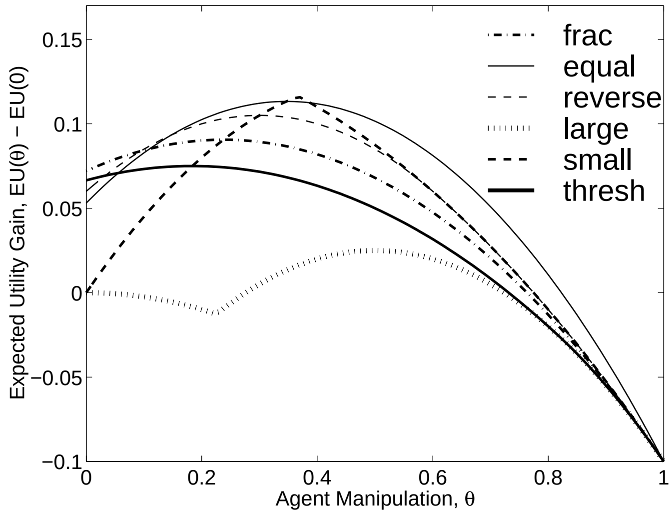

Figure 10: Expected Gain in Utility for different bids $b = v - \theta$ under each payment scheme. The parameters in each rule are carefully set to make each rule allocate the same expected discount, when the agent follows its optimal strategy for that rule. In this case the expected discount is 0.1δ in each rule.

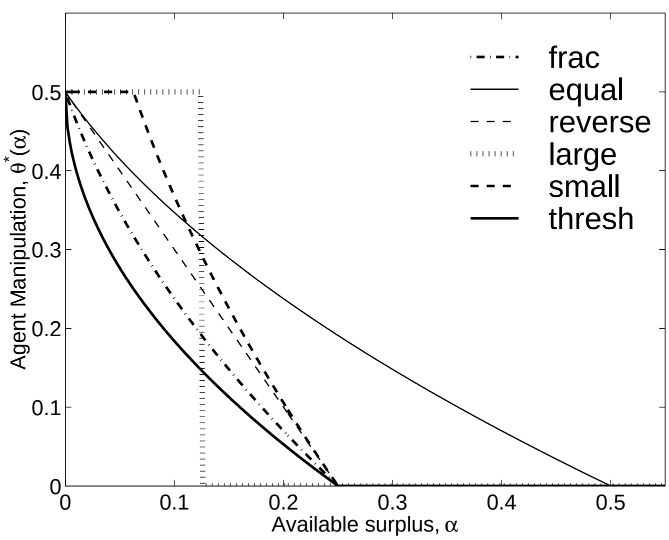

Figure 11: Optimal agent manipulation, $\theta_i^*(\alpha)$, (as a proportion of δ) under each payment scheme as the amount of available surplus increases from 0 to $\delta/2$ per-agent. The parameters are carefully selected at each α to give budget-balance in each rule.

Finally, we have a relationship between the expected discount allocated to agent i and the level of manipulation selected by agent i, when the parameter in Threshold is carefully selected to give budget-balance:

$$\begin{aligned}\theta_i^*(\alpha) &= \min[C_t^*(\alpha), \delta/2] \\ &= C_t^*(\alpha) = \max\left[0, \frac{\delta}{4}\left(2 - \sqrt{16\alpha}\,\right)\right]\end{aligned}$$

Plugging in values, if $\alpha \geq 1/4$ then $\theta_i^*(\alpha) = 0$, while if $\alpha = 0$ then $\theta_i^*(\alpha) = \delta/2$.

Figure 11 plots these relationships between agent manipulation and surplus parameter α under each payment rule. Notice that Vickrey payments can be implemented with surplus $\geq \delta/4$, so all schemes except Equal and No-Discount prevent manipulation completely for $\alpha \geq 1/4$. For smaller amounts of surplus the market maker is forced to deviate from Vickrey, and move left in the Figure, reducing the level of expected discount that can be paid to each agent. Finally, with no available discount, $\alpha = 0$, no scheme can have a beneficial effect on manipulation and the agent will manipulate by $\delta/2$.

Notice that the simple minded Equal scheme appears to have quite bad incentive properties. In fact, the Threshold method dominates all other schemes in this model except Large. Large has an interesting bad-good phase transition at $\alpha = 1/8$, and can prevent manipulation completely for $1/8 \leq \alpha \leq 1/4$ even though agents with small Vickrey discounts might have benefited from manipulation with hindsight.

Essentially, each payment rule imposes different tradeoffs on an agent between the expected loss from bidding below the critical value and losing a beneficial trade, and the positive effect that a lower bid price can have on an agent's payment. Some rules design the tradeoffs to make manipulation less desirable, for example making it the case that an agent can only hope to reduce its payment in the problems in which manipulation is also more likely to drop the agent's bid price below the critical value. In Section 7 we draw some general conclusions about the incentive properties across the payment rules.

6 A Simple Experimental Study of Manipulation in Each Payment Rule

In this section we provide experimental results that compare the manipulation by agents, this time in a more realistic model than the model assumed in the previous section.

The key differences in the experimental analysis are:

- *ex post* budget-balance, with the surplus as computed after the exchange is cleared.
- rule parameters are computed *after* bids and asks are received to implement the optimal discount allocation within each rule.

- we allow agents to adjust their bid and ask prices on multiple trades.

The analysis quantifies the level of manipulation in each scheme for different distributions of problem instances, and also quantifies the *allocative efficiency* in each scheme.

6.1 The Experimental Model

We make the following simplifying assumptions:

[A5] we consider only simple manipulations by $y\%$, where an agent reduces all of its reported values by $y\%$, e.g. shaving bid prices upwards and ask prices downwards.

[A6] we consider only symmetric Nash equilibria in which every agent either manipulates by 0% or $y\%$.

[A7] we assume that agents select the symmetric Nash equilibria which maximizes the expected individual increase in utility from manipulation by $y\%$ in comparison with 0%.[16]

[A8] the valuations of all agents are drawn from the same distributions.

[A9] agents are either buyers or sellers.

Notice that by requiring agents to select a static manipulation policy, $y^*\%$, across all problem instances we are making an implicit assumption that agents do not have anything other than distributional information about the bids and asks from agents in any particular instance of the exchange.

The problem instances are generated by adapting distributions Random, Weighted Random, Decay, and Uniform distributions from Sandholm [27], to generate values on bundles instead of bid prices in a combinatorial auction.

In each instance we generate values on 100 bundles, in a problem with 50 items, and then distribute the bundles across agents to generate valuation functions for the agents. We consider problems with 5, 10, and 20 agents, and buyers and sellers in proportions (buyer/seller) of { 5/5, 7/3, 2/3, 4/1, 10/10, 15/5 }. We assume that a buyer wants to buy *at most one* of the bundles for which it has value, and that a seller wants to sell *at most one* of the bundles for which it has value. The exchange design allows agents to submit exclusive-or bids on bundles, so that an agent can quite easily represent its truthful valuation function. We assume free-disposal, and clean-up any distributions that are generated that do not satisfy this property. The exchange allows full aggregation, such that a bid can be matched with any number of asks, and an ask can be matched with any number of bids.

For each problem configuration we generate 80 problem instances. For each instance, we consider a discrete number of manipulation levels $y\%$. At each level we compute the average

[16] It is not enough to simply measure the absolute utility at different levels of manipulation because even an exchange with Vickrey payments is not safe from collusion. Although it is an optimal strategy for an individual agent in a Vickrey exchange to bid truthfully whatever the bids/asks of other agents, the bids/asks of other agents can have a beneficial effect on its own utility.

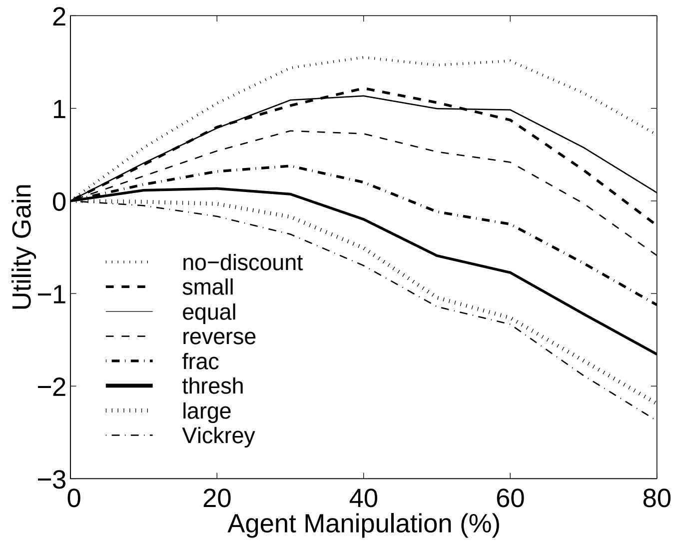

Figure 12: Average Single-Agent Gain in Utility from manipulation by $y\%$ (vs. truthful bidding), in a system in which every other agent manipulates by $y\%$. Problem size: 5 buyers/5 sellers.

single-agent gain in utility from manipulation by $y\%$ in comparison with no manipulation, assuming that every other agent manipulates by $y\%$, with each agent considered in turn. Finally, the single-agent gain in utility is averaged across all agents, and then across all instances for that configuration. The allocative efficiency for each configuration and each level of manipulation is averaged over all instances, and computed for the case that every agent manipulates by $y\%$.

6.2 Experimental Results: Manipulation

Figure 12 plots the average single agent gain in utility for different levels of manipulation, $y\%$, under the assumption that every other agent manipulates by $y\%$, for a problem with 5 buyers and 5 sellers.

By assumption [A7], we compute the symmetric Nash equilibrium for a particular payment rule as the level of manipulation, $y^*\%$, that coincides with the peak of a plot such as that in Figure 12. In this case, for the 5 buyers/5 sellers problem, the equilibrium manipulation level in Large and Threshold is less than under the other rules: around 10% and 20% in Large and Threshold, compared with 30%, 40% and 50% in Fractional, Equal and No-Discount. In addition, the amount of utility gain in Large and Threshold is much less than in the other schemes. In addition, notice that Vickrey, Large and Threshold have lower absolute *values* of utility gain across the range of possible levels of manipulation, and that the maximal individual agent gain from manipulation is lower than in other schemes.

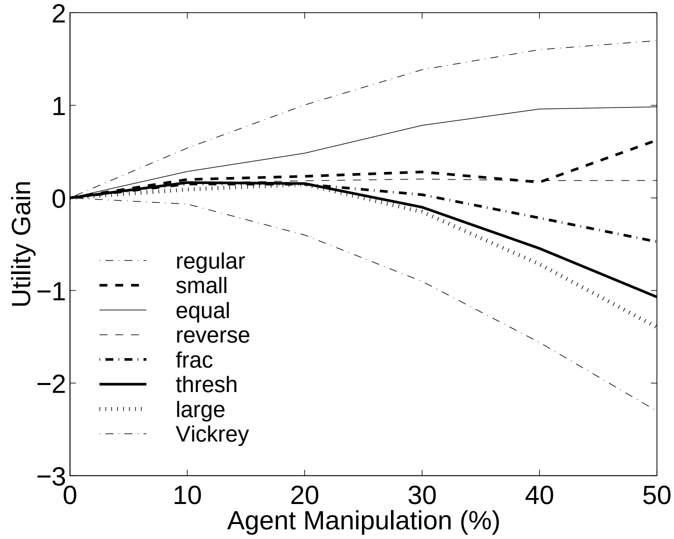

Figure 13: Average Single-Agent Gain in Utility from manipulation by $y\%$ (vs. truthful bidding), in a system in which every other agent manipulates by $y\%$. Problem size: 12 buyers/4 sellers.

By way of comparison, Figure 13 plots the average single-agent utility gain from manipulation in an exchange with more buyers than sellers, in this case 12 buyers and 4 sellers. Although the same basic observations hold, with the equilibrium manipulation level in the Threshold and Large schemes less than in the other schemes, the optimal level of manipulation appears to be higher than in the 5 buyer/ 5 seller problem in Figure 12. In fact this effect might not be too important because the absolute gain from manipulation of Threshold and Large is smaller in this experiment.

In Table 5 we summarize the experimental manipulation results across all problem configurations. We tabulate the average optimal degree of manipulation by agents in each scheme. Over all problems the Large and Threshold schemes perform quite well, with around 20% manipulation; this is in comparison, for example, with approaching 50% manipulation in the Equal, Small, and No Discount schemes. We also compute the average utility gain in each scheme from manipulation at 10%, 20%, and 30%; again Large and Threshold do well under this metric.

6.3 Experimental Results: Efficiency

Figure 14 (a) plots the allocative efficiency of the exchange at different levels of manipulation, averaged across all problem configurations. Notice that the allocative efficiency falls almost linearly (and even point-for-point) with manipulation, from 100% at 0% manipulation to 20% at 80% manipulation. Table 5 presents the average allocative efficiency in the

	No-Discount	Vickrey	Small	Frac
Utility Gain	0.799	-0.195	0.479	0.211
Correlation	0.053	1.0	0.356	0.590
Manipulation, θ^*	48	0	48	32
Efficiency (%)	58	100	58	78
	Threshold	Equal	Large	Reverse
Utility Gain	0.110	0.516	0.029	0.337
Correlation	0.543	0.356	0.176	0.522
Manipulation, θ^*	22	46	18	44
Efficiency (%)	86	62	88	64

Table 5: Experimental results, averaged across all configurations. *Utility gain* is the average single-agent utility gain from manipulation in each scheme, averaged over manip. 10%, 20%, and 30%. *Correlation* with Vickrey discounts is also computed for manip. 10%, 20% and 30%. *Manipulation* is the average equilibrium level of manipulation selected by agents. *Efficiency* is the average allocative efficiency at that level of manipulation.

exchange under each payment rule, at the optimal manipulation equilibrium. Threshold and Large achieve greater than 85% allocative-efficiency, while Equal Small and No Discount manage only around 60% efficiency. The efficiency effect is quite significant at these manipulation levels.

It is also interesting to consider how much of a problem budget-balance is with the Vickrey mechanism in these problems. Figure 14 (b) plots the degree of budget-balance failure with Vickrey payments at different levels of manipulation. The budget deficit with Vickrey discounts falls linearly as the level of manipulation increases, but is still a problem even when agents manipulate by 80%.

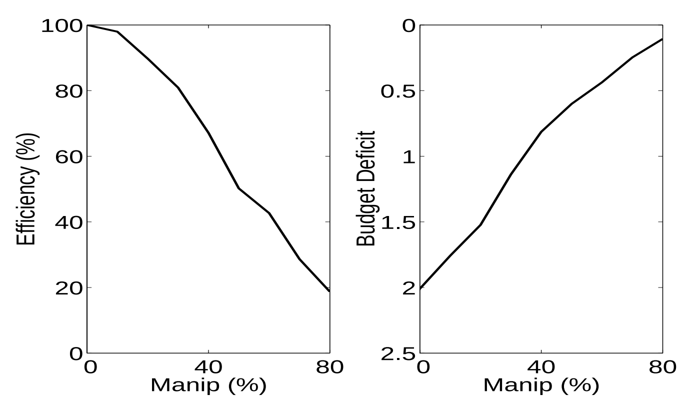

Figure 14: Experimental results, averaged across all problem configurations. (a) Allocative Efficiency; (b) Budget Deficit with Vickrey payments.

We also compare the *correlation* between Vickrey discounts and actual discounts under each scheme in Table 5. While the discounts in Fractional and Threshold are quite well correlated with the Vickrey discounts, notice that the discounts in Large are *not* very well correlated.

7 Discussion

The partial ordering {Large, Threshold } $\succ$ Fractional $\succ$ Reverse $\succ$ {Equal, Small } from the experimental results is remarkably consistent with the results of the simple analytical model that we presented in Section 5. Although the Large scheme generates slightly less manipulation and higher allocative efficiency than Threshold in the experimental tests, the Threshold discounts are quite well correlated with the Vickrey discounts while the Large discounts are quite uncorrelated with the Vickrey discounts. This points to the fact that an agent's discount in Large is very sensitive to its bid. We expect Large to be less robust than Threshold in practice because of this all-or-nothing characteristic.

There are a number of observations to make about the geometric properties of the Large and Threshold rules (see Figure 8), in comparison with the other rules. Recall that the critical value is the minimal price that the agent can bid and still execute the same trade when the exchange clears.

1. Rules with flat sections are better than rules without (e.g. Large, Threshold vs. Equal, Fractional)

2. Rules with flat sections for bids that are a long distance from an agent's critical value (e.g. Large, Threshold) perform better than rules with flat sections for bids that are close to an agent's critical value (e.g. Small, Reverse).

Rules with flat sections for bids that are a long distance from an agent's critical value provide discounts to agents with large Vickrey discounts, but no discounts to agents with small Vickrey discounts. Recognizing that a large Vickrey discount indicates that the agent had a large opportunity to reduce its final payment in a No Discount mechanism, this focus on agents with large Vickrey discounts is useful because it positions incentives to compensate agents with (ex post) easy opportunites for manipulation, while providing no incentives to reduce manipulation to agents with hard opportunities.

To make this more concrete, it is useful to introduce the concept of the *residual degree of manipulation freedom*, RDMF(i), to agent i in a mechanism.

Definition 10 (residual degree of manipulation freedom). *The residual degree of manipulation freedom,* RDMF(i)*, for agent i, given bids and asks by other agents, is the maximal amount that the agent can increase its utility with some bid $\hat{v}_i \neq v_i$, in comparison to its utility from the outcome when it bids truthfully.*

This is a measure of the ex post (i.e. given agents bids and asks, and the allocated discounts) opportunity for manipulation in a payment scheme. Clearly, RDMF$(i) = 0$ in the VCG mechanism, and RDMF$(i) = \Delta_{\text{vick},i}$ with the No Discount rule.

It is easy to show that an agent with perfect information about the bids and asks from other agents can always achieve as much utility as in the VCG mechanism for some bid.

Lemma 2 (maximal utility). *Agent i can achieve its utility in the Vickrey mechanism in any of the Vickrey-based payment schemes, with bid $\hat{v}_i(\lambda_i^*) = v_i(\lambda_i^*) - (V^* - (V_{-i})^*)$, and $\hat{v}_i(\lambda) = 0$ on all $\lambda \neq \lambda_i^*$; where λ_i^* is the trade to agent i in the VCG mechanism.*

This is a tight upper-bound on the maximal utility available to an agent, because by Lemma 1 we know that an agent cannot achieve more than its Vickrey utility.

This allows us to quantify the RDMF(i) to an agent in a particular payment scheme, $\mathcal{M}$. Let $u_{\text{vick},i}$ denote the agent's utility with truth-revelation in the VCG mechanism, and let $u_i(\mathcal{M})$ denote the agent's utility with truth-revelation in $\mathcal{M}$. The RDMF is computed as the maximal gain in utility available to an agent:

$$\begin{aligned}\text{RDMF}(i) &= u_{\text{vick},i} - u_i(\mathcal{M}) \\ &= [v_i(\lambda_i^*) - (v_i(\lambda_i^*) - \Delta_{\text{vick},i})] - [v_i(\lambda_i^*) - (v_i(\lambda_i^*) - \Delta_i)] \\ &= \Delta_{\text{vick},i} - \Delta_i\end{aligned}$$

where $0 \leq \Delta_i \leq \Delta_{\text{vick},i}$ is the discount allocated to agent i in the payment rule $\mathcal{M}$.

Immediately, we have:

Theorem 7. *The Threshold rule minimizes the maximal RDMF amongst all Vickrey-based (IR) and (BB) payment schemes.*

We believe that this is useful because it is rational for an agent to manipulate by larger amounts, and more frequently, as RDMF increases. Therefore, if one considers a convex-increasing plot of manipulation vs. RDMF(i), for every agent i, the optimal strategy to maximize the *total* reduction in manipulation across all agents is to allocate discounts to minimize the maximal RDMF. The optimality of this greedy rule is easy to see because there are decreasing returns in allocating each additional unit of available discount to the same agent.

A Fruit Picking Analogy

It is perhaps useful to think of the payment rules in terms of fruit picking. The mechanism has only limited resources (surplus), but can at least see the fruit (opportunities for manipulation). The goal of the mechanism is to make it as hard as possible for agents to pick fruit. The Large mechanism picks the low-hanging fruit, while the Threshold rule just hides them higher in the tree so that they are more difficult to see from the branches (this takes less effort than actually picking the fruit, which is analogous with the agents not being able to see them at all). The agents cannot see the fruit very clearly, and are reluctant to try to pick fruit that might be too hard to pick.

In terms of efficiency the picture is mixed. In general, reducing agent manipulation can only increase allocative-efficiency, but the easy low fruit that are picked in preference of the difficult high fruit, are also the manipulation opportunities less likely to disrupt rhe efficient trade.

8 Future Work

The most immediate direction for future work is to complete an equilibrium analysis of the most interesting payment rules (e.g. Large, Threshold, Fractional), where the surplus and bids from other agents are endogenous within the model, and computed in Bayesian-Nash equilibrium.

Interesting extensions include: consider the effect of strategic manipulation through timing of bids and asks; complete a more complete experimental analysis with a richer strategy space; complete a theoretical analysis to establish rules with worst-case optimal performance; examine the effect of solving the WD and pricing problems approximately on the manipulations; and consider randomized payment rules, and payment rules with multiple flat regions.

9 Conclusions

We proposed a family of budget-balanced and individual-rational mechanisms for combinatorial exchanges, that clear the exchange to maximize reported surplus and then allocate discounts to agents to minimize a suitable distance metric to Vickrey payments. The mechanisms are explicitly designed *not* to be incentive-compatible, both to be able to clear the exchange to maximize reported surplus, and also to leverage the bounded-rationality and limited information available to agents, that makes them unable to fully exploit opportunites for manipulation.

The analytical and experimental results both suggest that a simple Threshold rule has useful incentive properties, and provides higher allocative efficiency than other rules. The effect of the Threshold rule is to remove easy opportunities for manipulation, without attempting to provide incentives for truth-revelation to agents with hard opportunites for manipulation.

Finally, we note that the schemes outlined here can also allow a market maker to make a small profit by taking a sliver of budget-balance, or used in combination with a participation charge to move payments closer to Vickrey payments.

10 Acknowledgments

We thank William Walsh, anonymous reviewers, and seminar participants for their helpful comments on an earlier draft of this paper. This research was funded in part by National Science Foundation Grant SBR 97-08965. The first author also gratefully acknowledges financial support from an IBM Research Fellowship.

11 Appendix: Computing the Distance Functions

In what follows we will give details about how to derive the discount (payment) function for the different distance functions of Section 4. The payment functions obtained are summarized in Table 2.

To solve the problem captured by the mathematical model we will apply the Lagrangian relaxation technique. We relax the budget balance constraint by introducing an associated Lagrange multiplier λ and moving the constraint into the objective:

$$z(\lambda) = \min_{\mathbf{\Delta}} \mathbf{L}(\mathbf{\Delta}, \mathbf{\Delta}_{\text{vick}}) + \lambda \left(\sum_i \Delta_i - V^* \right)$$

$$\text{s.t.,} \quad \mathbf{0} \leq \mathbf{\Delta} \leq \mathbf{\Delta}_{\text{vick}}.$$

Note that for a given multiplier both the optimal solution and its value will be functions of the multiplier, we denote this solution by $\mathbf{\Delta}^*(\lambda)$ and its value by $z(\lambda)$. Sometimes it will be more convenient to use a function of λ to parameterize the solution, these are the C's and μ in the various payment rules.

Obviously $z(\lambda)$ is a lower bound on the optimal value of the original mathematical model [PP]. Lagrangian optimization involves maximizing $z(\lambda)$ in λ, that is, obtaining the best possible lower bound $z(\lambda^*)$ for the optimal value of [PP]. Weak duality implies that if the solution $\mathbf{\Delta}^*(\lambda^*)$ satisfies the budget balance constraint then it will also be an optimal solution for the original problem [PP]. In this case $z(\lambda^*)$ is the optimal objective value for [PP].

Thus for a given distance function we will need to do the following:

1. solve the Lagrangian subproblem for a given value of the multiplier to obtain a solution parameterized by λ,

2. find the multiplier value that maximizes $z(\lambda)$, and

3. show that the solution to the Lagrangian subproblem for the optimal multiplier also satisfies the relaxed constraint and thus it is an optimal solution for [PP].

Note that for all but one of the distance functions we considered the optimal solution to [PP] is also unique; that is, there is exactly one value for the Lagrange multiplier where the optimal solution to the Lagrangian subproblem will be budget balanced. This means that for values of the parameter on one side of the optimal multiplier the optimal solution to the Lagrangian subproblem will be budget balanced and for values on the other side of the optimal multiplier the solution will never be budget balanced.

Assumptions about the distribution of the Vickrey discounts $\boldsymbol{\Delta}_{\text{vick}}$ and the proportion of their sum to the available surplus V^* are sometimes needed in the derivations. Therefore, we will assume that all Vickrey discounts are positive and that the agents are indexed so that their Vickrey discounts are decreasing (let I denote the total number of agents who might receive discounts):

$$\Delta_{\text{vick},I} \leq \Delta_{\text{vick},I-1} \leq \ldots \leq \Delta_{\text{vick},2} \leq \Delta_{\text{vick},1}.$$

We add dummy points $\Delta_{\text{vick},0} = \infty$ and $\Delta_{\text{vick},I+1} = 0$ and index the interval $[\Delta_{\text{vick},k+1}, \Delta_{\text{vick},k}]$ by k. (To properly partition $[0, \infty)$ these intervals have to be open on one end. In some cases it does matter which end of the interval is open, we will indicate this at the description of the individual distance functions.) We also assume that the sum of Vickrey discounts across all agents exceeds the available surplus ($\sum_1^I \Delta_{\text{vick},i} > V^*$) since otherwise the exchange could distribute Vickrey discounts to the agents and achieve full incentive compatibility; and that the available surplus is positive ($V^* > 0$).

11.1 L_2 and L_∞ norms

We will discuss the L_2 norm in detail, the L_∞ norm can be handled similarly (notice that the L_∞ norm can be thought of as the limit of the L_k norm as k goes to infinity). For the L_2 norm the Lagrangian subproblem is

$$z(\lambda) = \min_{\Delta} \sum_{i=1}^{I} (\Delta_{\text{vick},i} - \Delta_i)^2 + \lambda(\sum_{i=1}^{I} \Delta_i - V^*)$$

$$\text{s.t.} \quad 0 \leq \Delta_i \leq \Delta_{\text{vick},i} \quad i = 1, \ldots, I.$$

11.1.1 Solving the Lagrangian subproblem for a given multiplier

Note that this problem decomposes into smaller problems for each i:

$$\min_{\Delta_i} (\Delta_{\text{vick},i} - \Delta_i)^2 + \lambda \Delta_i$$

$$\text{s.t.} \quad 0 \leq \Delta_i \leq \Delta_{\text{vick},i}.$$

From the first order condition

$$-2(\Delta_{\text{vick},i} - \Delta_i) + \lambda = 0$$

and the bounds on the variable we obtain

$$\Delta_i^*(\lambda) = \max(0, \Delta_{\text{vick},i} - \lambda/2)$$

The second derivative is positive thus the objective function is convex.

Let us introduce C_t for $\lambda/2$, thus

$$\Delta_i^*(C_t) = \max(0, \Delta_{\text{vick},i} - C_t), \quad i = 1, \dots, I \tag{Threshold}$$

Note that if C_t falls into interval k agents $i = 1, \dots, k$ will receive discounts $\Delta_i^*(C_t) = \Delta_{\text{vick},i} - C_t$ while agents $i = k+1, \dots, I$ will not receive any discounts.

11.1.2 Finding the best multiplier value

Substituting $\Delta_i^*(C_t)$ into $z(C_t)$, the function in interval k is

$$z^k(C_t) = \sum_{i=k+1}^{I} (\Delta_{\text{vick},i})^2 + 2C_t \left(\sum_{i=1}^{k} \Delta_{\text{vick},i} - V^* \right) - kC_t^2.$$

Observe that the $z(C_t)$ function, which is pieced together from the $z^k(C_t)$-s, is continuous everywhere (it is continuous within each interval and $z^k(\Delta_{\text{vick},k}) = z^{k-1}(\Delta_{\text{vick},k})$). The first derivative of the function in interval k is

$$(z^k)'(C_t) = -2kC_t + 2 \left(\sum_{i=1}^{k} \Delta_{\text{vick},i} - V^* \right),$$

thus $z'(C_t)$, which is pieced together from the $(z^k)'(C_t)$-s, is a continuous everywhere. Also, $z'(C_t)$ is a monotone decreasing function. For a very small value of the parameter

$$z'(\epsilon) = (z^I)'(\epsilon) = -2I\epsilon + 2(\sum_{i=1}^{I} \Delta_{\text{vick},i} - V^*) > 0$$

since the sum of the Vickrey discounts exceeds the available surplus. For any parameter value in interval 0

$$z'(C_t) = (z^0)'(C_t) = -2V^* < 0.$$

Therefore, because of the continuity of the first derivative, there exists an interval K and a unique point C_t^* within this interval so that $z'(C_t^*) = (z^K)'(C_t^*) = 0$; that is,

$$\Delta_{\text{vick},K+1} \le C_t^* = \frac{\sum_{i=1}^{K} \Delta_{\text{vick},i} - V^*}{K} \le \Delta_{\text{vick},K}.$$

11.1.3 Finding the overall best solution

We show that budget balance holds for $\mathbf{\Delta}^*(C_t^*)$, the solution to the Lagrangian subproblem with the best parameter C_t^* computed above. The sum of discounts for parameter C_t^* (which falls into interval K) is

$$\sum_{i=1}^{K}(\Delta_{\text{vick},i} - C_t^*) + \sum_{i=K+1}^{I} 0 = \sum_{i=1}^{K} \Delta_{\text{vick},i} - KC_t^* = V^*,$$

the last equality implied by the definition of C_t^*. This implies that $\mathbf{\Delta}^*(C_t^*)$ is an optimal solution to [PP] of objective value

$$z^K(C_t^*) = \sum_{i=K+1}^{I} (\Delta_{\text{vick},i})^2 + \frac{1}{K}\left(\sum_{i=1}^{K} \Delta_{\text{vick},i} - V^*\right)^2.$$

11.2 The L_{RE} distance functions

For the L_{RE} distance function the Lagrangian subproblem is

$$z(\lambda) = \min_{\Delta} \sum_{i=1}^{I} \frac{(\Delta_{\text{vick},i} - \Delta_i)}{\Delta_{\text{vick},i}} + \lambda(\sum_{i=1}^{I} \Delta_i - V^*)$$

$$\text{s.t.} \quad 0 \leq \Delta_i \leq \Delta_{\text{vick},i} \quad i = 1, \ldots, I.$$

11.2.1 Solving the Lagrangian subproblem for a given multiplier

Note that this problem decomposes into smaller problems for each i:

$$\min_{\Delta_i} 1 + \Delta_i \left(\lambda - \frac{1}{\Delta_{\text{vick},i}}\right)$$

$$\text{s.t.} \quad 0 \leq \Delta_i \leq \Delta_{\text{vick},i}.$$

The minimum is attained at zero if the coefficient of the variable is positive, and at the upper bound of the variable if the coefficient is negative. When the coefficient is zero the variable could take any value, we choose to set it to zero in this case. Let us introduce C_s for $1/\lambda$, thus

$$\Delta_i^*(C_s) = \Delta_{\text{vick},i} \ \text{ if } \ C_s > \Delta_{\text{vick},i}, \ 0 \ \text{ otherwise} \tag{Small}$$

Note that if C_s falls into the k^{th} interval $(\Delta_{\text{vick},k+1}, \Delta_{\text{vick},k}]$ then agents $i = 1, \ldots, k$ will not receive any discounts while agents $i = k+1, \ldots, I$ will receive their Vickrey discounts.

11.2.2 Finding the best multiplier value

Substituting $\Delta_i^*(C_s)$ into $z(C_s)$ the function in interval k is

$$z^k(C_s) = k + \frac{1}{C_s}\left(\sum_{i=k+1}^{I} \Delta_{\text{vick},i} - V^*\right).$$

It is easy to see that $z(C_s)$ is continuous. The first derivative of the function in interval k is

$$(z^k)'(C_s) = -\frac{1}{C^2}\left(\sum_{i=k+1}^{I} \Delta_{\text{vick},i} - V^*\right),$$

a monotone decreasing function (but not continuous in the breakpoints). Note that $\sum_{i=k+1}^{I} \Delta_{\text{vick},i} - V^*$ is a constant within each interval and it "jumps" by $\Delta_{\text{vick},k}$ from interval k to $k-1$. For $k = I$ the constant is negative since $V^* > 0$, and for $k = 0$ it is positive since the total of the Vickrey discounts exceeds the available surplus. Thus there exists a unique index K so that the constant is positive for intervals $K, \dots, I$ and negative for intervals $0, \dots, K-1$. This implies that the first derivative is positive, and therefore $z(C_s)$ is increasing, for $C_s \leq \Delta_{\text{vick},K}$, and that the first derivative is negative, and therefore $z(C_s)$ is decreasing, for $C_s > \Delta_{\text{vick},K}$.

That is, $C_s^* = \Delta_{\text{vick},K}$ for the index K where $\sum_{i=K+1}^{I} \Delta_{\text{vick},i} \leq V^*$ but $\sum_{i=K}^{I} \Delta_{\text{vick},i} > V^*$.

11.2.3 Finding the overall best solution

It is obvious that budget balance holds for $\mathbf{\Delta}^*(C_s^*)$, the solution to the Lagrangian subproblem with the best parameter C_s^* computed above. The sum of discounts for parameter value $C_s^* = \Delta_{\text{vick},K}$ is

$$\sum_{i=1}^{K} 0 + \sum_{i=K+1}^{I} \Delta_{\text{vick},i} \leq V^*$$

by the choice of the index K. This implies that $\mathbf{\Delta}^*(C_s^*)$ is an optimal solution to [PP] of objective value

$$z^K(\Delta_{\text{vick},K}) = K + \frac{1}{\Delta_{\text{vick},K}}\left(\sum_{i=K+1}^{I} \Delta_{\text{vick},i} - V^*\right).$$

11.3 The L_{WE} distance function

For the L_{WE} distance function the Lagrangian subproblem is

$$z(\lambda) = \min_{\Delta} \sum_{i=1}^{I} \Delta_{\text{vick},i}\,(\Delta_{\text{vick},i} - \Delta_i) + \lambda(\sum_{i=1}^{I} \Delta_i - V^*)$$

$$\text{s.t. } 0 \le \Delta_i \le \Delta_{\text{vick},i} \quad i = 1, \dots, I.$$

11.3.1 Solving the Lagrangian subproblem for a given multiplier

Note that this problem decomposes into smaller problems for each i:

$$\min_{\Delta_i} (\Delta_{\text{vick},i})^2 + \Delta_i(\lambda - \Delta_{\text{vick},i})$$

$$\text{s.t. } 0 \le \Delta_i \le \Delta_{\text{vick},i}.$$

The minimum is attained at zero if the coefficient of the variable is positive, and at the upper bound of the variable if the coefficient is negative. When the coefficient is zero the variable could take any value up to its upper bound, we choose to set it to zero in this case. Let us introduce C_l for λ, thus

$$\Delta_i^*(C_l) = \Delta_{\text{vick},i} \text{ if } C_l < \Delta_{\text{vick},i}, \ 0 \text{ otherwise} \tag{Large}$$

Note that if C_l falls into the k^{th} interval $[\Delta_{\text{vick},k+1}, \Delta_{\text{vick},k})$ then agents $i = k+1, \dots, I$ will not receive any discounts while agents $i = 1, \dots, k$ will receive their Vickrey discounts.

11.3.2 Finding the best multiplier value

Substituting $\Delta_i^*(C_l)$ into $z(C_l)$ the function in interval k is

$$z^k(C_l) = \sum_{i=k+1}^{I} (\Delta_{\text{vick},i})^2 + C_l(\sum_{i=1}^{k} \Delta_{\text{vick},i} - V^*).$$

It is easy to see that $z(C_l)$ is continuous. Note that for $k = 0$ the coefficient of C_l is negative since $V^* > 0$, and that for $k = I$ the coefficient is positive since the total of the Vickrey discounts exceeds the available surplus. Therefore, there exists a unique index K so that the constant is non-positive for $k \le K$ and is strictly positive for $k \ge K + 1$. Thus the function $z(C_l)$ is increasing for $C_l < \Delta_{\text{vick},K+1}$ and decreasing for $C_l \ge \Delta_{\text{vick},K+1}$.

That is, the optimum of $z(C_l)$ is attained at $C_l^* = \Delta_{\text{vick},K+1}$. for the index K where $\sum_{i=1}^{K} \Delta_{\text{vick},i} \le V^*$ but $\sum_{i=1}^{K+1} \Delta_{\text{vick},i} > V^*$.

11.3.3 Finding the overall best solution

It is obvious that budget balance holds for $\boldsymbol{\Delta}^*(C_l^*)$, the solution to the Lagrangian subproblem with the best parameter C_l^* computed above. The sum of discounts for parameter value $C_l^* = \Delta_{\text{vick},K+1}$ is

$$\sum_{i=1}^{K} \Delta_{\text{vick},i} + \sum_{i=K+1}^{I} 0 \le V^*$$

by the choice of the index K. This implies that $\mathbf{\Delta}^*(C_l^*)$ is an optimal solution to [PP] of objective value

$$z^K(\Delta_{\text{vick},K+1}) = \sum_{i=K+1}^{I} (\Delta_{\text{vick},i})^2 + \Delta_{\text{vick},K+1}(\sum_{i=1}^{K} \Delta_{\text{vick},i} - V^*).$$

11.4 The L_{RE2} distance function

For the L_{RE2} distance function the Lagrangian subproblem is

$$z(\lambda) = \min_{\Delta} \sum_{i=1}^{I} \frac{(\Delta_{\text{vick},i} - \Delta_i)^2}{\Delta_{\text{vick},i}} + \lambda(\sum_{i=1}^{I} \Delta_i - V^*)$$

$$\text{s.t.} \;\; 0 \leq \Delta_i \leq \Delta_{\text{vick},i} \;\; i = 1, \ldots, I.$$

11.4.1 Solving the Lagrangian subproblem for a given multiplier

Note that this problem decomposes into smaller problems for each i:

$$\min_{\Delta_i} \frac{(\Delta_{\text{vick},i} - \Delta_i)^2}{\Delta_{\text{vick},i}} + \lambda\Delta_i$$

$$\text{s.t.} \;\; 0 \leq \Delta_i \leq \Delta_{\text{vick},i}.$$

From the first order condition

$$-\frac{2}{\Delta_{\text{vick},i}}(\Delta_{\text{vick},i} - \Delta_i) + \lambda = 0$$

we obtain

$$\Delta_i = (1 - \frac{\lambda}{2})\Delta_{\text{vick},i}.$$

Let us introduce μ for $1 - \lambda/2$. Thus

$$\Delta_i^*(\mu) = \mu\Delta_{\text{vick},i} \;\; \text{if } 0 \leq \mu \leq 1, \;\; 0 \;\; \text{if } \mu < 0 \;\; \text{and } \Delta_{\text{vick},i} \;\; \text{if } \mu > 1 \qquad \text{(Fractional)}$$

That is, if μ is between zero and one, all agents receive discounts proportional to their Vickrey discounts.

11.4.2 Finding the best multiplier value

Substituting $\Delta_i^*(\mu)$ into $z(\mu)$ the function for $\mu < 0$ is

$$z^{<0}(\mu) = \sum_{i=1}^{I} \Delta_{\text{vick},i} - 2(1-\mu)V^*,$$

an increasing function since the coefficient of μ is positive. On the other hand, for $\mu > 1$,

$$z^{>1}(\mu) = 2(1-\mu)(\sum_{i=1}^{I} \Delta_{\text{vick},i} - V^*),$$

a decreasing function since the coefficient of μ is negative. The function for $0 \leq \mu \leq 1$ can be written as

$$z^{[0,1]}(\mu) = (1-\mu)^2 \sum_{i=1}^{I} \Delta_{\text{vick},i} + 2\mu(1-\mu)(\sum_{i=1}^{I} \Delta_{\text{vick},i} - V^*) =$$

$$= (1-\mu^2) \sum_{i=1}^{I} \Delta_{\text{vick},i} - 2(1-\mu)V^*.$$

It is easy to check that $z(\mu)$ is continuous. First order conditions for $z^{[0,1]}$ imply

$$-2\mu \sum_{i=1}^{I} \Delta_{\text{vick},i} + 2V^* = 0,$$

and since the second derivative is negative,

$$\mu^* = \frac{V^*}{\sum_{i=1}^{I} \Delta_{\text{vick},i}}.$$

μ^* indeed falls into the interval $[0,1]$ since the sum of Vickrey discounts exceeds the available surplus and both the numerator and denominator are positive.

11.4.3 Finding the overall best solution

Budget balance holds trivially since for the optimal parameter value μ^* the total discount awarded is

$$\sum_{i=1}^{I} \mu^* \Delta_{\text{vick},i} = \frac{V^*}{\sum_{i=1}^{I} \Delta_{\text{vick},i}} \sum_{i=1}^{I} \Delta_{\text{vick},i} = V^*.$$

11.5 The L_Π distance function

Before introducing the Lagrange multiplier we take the logarithm of the L_Π distance function. The optimum after the transformation will be attained at the same value as before the transformation since the logarithm function is continuous and monotone. Note that since the Vickrey discounts are positive the logarithm function can be applied.

Thus after the transformation the Lagrangian subproblem becomes

$$z(\lambda) = \min_{\Delta} \sum_{i=1}^{I} \log(\Delta_{\text{vick},i}) - \sum_{i=1}^{I} \log(\Delta_i) + \lambda(\sum_{i=1}^{I} \Delta_i - V^*)$$
$$\text{s.t.} \quad 0 \leq \Delta_i \leq \Delta_{\text{vick},i} \quad i = 1, \dots, I.$$

11.5.1 Solving the Lagrangian subproblem for a given multiplier

After taking the logarithm the problem decomposes into smaller problems for each i:

$$\min_{\Delta_i} \log(\Delta_{\text{vick},i}) - \log(\Delta_i) + \lambda \Delta_i$$
$$\text{s.t.} \quad 0 \leq \Delta_i \leq \Delta_{\text{vick},i}.$$

Equating the first derivative with zero implies $\Delta_i = 1/\lambda$. Let us introduce C_r for $1/\lambda$. Thus from the first order condition and the bounds on the variables it follows that

$$\Delta_i^*(C_r) = \min(C_r, \Delta_{\text{vick},i}) \tag{Reverse}$$

Note that if C_r falls into interval k then agents $i = 1, \dots, k$ get a discount in the amount of C_r, while agents $i = k+1, \dots, I$ will receive their Vickrey discounts.

11.5.2 Finding the best multiplier value

Substituting $\Delta_i^*(C_r)$ into $z(C_r)$ the function in interval k is

$$z^k(C_r) = \sum_{i=1}^{k} \log(\Delta_{\text{vick},i}) + k - k \log(C_r) + \frac{1}{C_r} \left(\sum_{i=k+1}^{I} \Delta_{\text{vick},i} - V^* \right).$$

It is easy to check that $z^k(C_r)$ is continuous. The first derivative in interval k is

$$(z^k)'(C_r) = -\frac{k}{C_r} - \frac{1}{C_r^2} \left(\sum_{i=k+1}^{I} \Delta_{\text{vick},i} - V^* \right).$$

$z'(C_r)$ is also a continuous function but not necessarily monotone. For a very small value of the parameter

$$z'(\epsilon) = (z^I)'(\epsilon) = \frac{1}{\epsilon^2}(V^* - \epsilon I) > 0$$

for a small enough ϵ. On the other hand for a large parameter value (that falls into interval zero)

$$z'(C_r) = -\frac{1}{C_r^2}\left(\sum_{i=1}^{I} \Delta_{\text{vick},i} - V^*\right) < 0$$

since the sum of the Vickrey discounts exceeds the available surplus. Therefore, because of the continuity of the first derivative, there exists at least one parameter value where it is zero. Choose C_r^* to be the smallest parameter value (and K to be the corresponding interval) for which the first derivative is zero:

$$C_r^* = -\frac{1}{K}\left(\sum_{i=K+1}^{I} \Delta_{\text{vick},i} - V^*\right).$$

11.5.3 Finding the overall best solution

For the parameter value C_r^* (that falls into the interval K) total of the awarded discounts is

$$KC_r^* + \sum_{i=K+1}^{I} \Delta_{\text{vick},i} = V^*$$

by definition of C_r^*.

References

[1] 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.

[2] Kenneth J Arrow. The property rights doctrine and demand revelation under incomplete information. In M Boskin, editor, *Economics and Human Welfare*. Academic Press, New York, 1979.

[3] Moshe Babaioff and Noam Nisan. Concurrent auctions across the supply chain. In *Proc. 3rd ACM Conference on Electronic Commerce*, pages 1–10. ACM Press, 2001.

[4] Salvador Barbera and Matthew O Jackson. Strategy-proof exchange. *Econometrica*, 63(1):51–87, 1995.

[5] Sushil Bikchandani and Joseph M Ostroy. The package assignment model. Technical report, Anderson School of Management and Department of Economics, UCLA, 1999.

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

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

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

[9] Andrew J Davenport and Jayant R Kalagnanam. Price negotiations for procurement of direct inputs. Technical report, IBM Research Report RC 22078, 2001. To appear, Proc. of IMA Workshop: Mathematics of the Internet: E-Auction and Markets, Minneapolis, MN, December 2000.

[10] Sven de Vries and Rakesh Vohra. Combinatorial auctions: A brief survey. Technical report, MEDS Department, Kellogg Graduate School of Management, Northwestern University, 2000.

[11] Marta Eso, Soumyadip Ghosh, Jayant R Kalagnanam, and Laszlo Ladanyi. Bid evaluation in procurement auctions with piece-wise linear supply curves. Technical report, IBM TJ Watson Research Center, 2001. in preparation.

[12] Marta Eso, Jayant R Kalagnanam, Laszlo Ladanyi, and Y Grace Li. Winner determination in bandwidth exchanges. Technical report, IBM TJ Wastson Research Center, 2001. In preparation, presentation available at www.research.ibm.com/auctions/index.htm.

[13] Alan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41:587–602, 1973.

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

[15] Jayant R Kalagnanam, Andrew J Davenport, and H S Lee. Computational aspects of clearing continous double auctions with assignment constraints and indivisible demand. *Electronic Commerce Journal*, 1(3), 2001.

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

[17] John O Ledyard. The design of coordination mechanisms and organizational computing. *Journal of Organizational Computing*, 1:121–134, 1993.

[18] John O Ledyard, Mark Olson, David Porter, Joseph A Swanson, and David P Torma. The first use of a combined value auction for transportation services. Technical Report Social Science Working Paper 1093, California Institute of Technology, 2000.

[19] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. In *Proc. 3rd ACM Conference on Electronic Commerce*, pages 18–28. ACM Press, 2001.

[20] R Preston McAfee. A dominant strategy double auction. *J. of Economic Theory*, 56:434–450, 1992.

[21] Robert B Myerson. Incentive compatibility and the bargaining problem. *Econometrica*, 47:61–73, 1979.

[22] Robert B Myerson. Optimal auction design. *Mathematics of Operation Research*, 6:58–73, 1981.

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

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

[25] David C Parkes. An iterative generalized Vickrey auction: Strategy-proofness without complete revelation. In *Proc. AAAI Spring Symposium on Game Theoretic and Decision Theoretic Agents*. AAAI Press, March 2001.

[26] Michael H Rothkopf, Aleksandar Pekeč, and Ronald M Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.

[27] Tuomas W Sandholm. An algorithm for optimal winner determination in combinatorial auctions. In *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, pages 542–547, 1999.

[28] 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*, 2001.

[29] Tuomas W Sandholm, Subhash Suri, Andrew Gilpin, and David Levine. Winner determination in combinatorial auction generalizations. In *Proc. International Conference on Autonomous Agents, Workshop on Agent-based Approaches to B2B*, pages 35–41, 2001.

[30] Mark A Satterthwaite and Steven R Williams. Bilateral trade with the sealed bid k-double auction: Existence and efficiency. *Journal of Economic Theory*, 48:107–133, 1989.

[31] Sayee Srinivasan, Jan Sattaert, and Andrew B Whinston. Portfolio trading and electonic networks. Technical report, Center for Research in Electronic Commerce, University of Texas at Austin, 1998.

[32] Hal R Varian. Economic mechanism design for computerized agents. In *Proc. USENIX Workshop on Electronic Commerce*, 1995. Minor update, 2000.

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

[34] William E Walsh, Michael P Wellman, and Fredrik Ygge. Combinatorial auctions for supply chain formation. In *Proc. ACM Conference on Electronic Commerce*, pages 260–269, 2000.

[35] R Wilson. Incentive efficiency of double auctions. *Econometrica*, 53:1101–1115, 1985.

[36] Kiho Yoon. The modified vickrey double auction. *Journal of Economic Theory*, 2001.

[37] Edo Zurel and Noam Nisan. An efficient approximate allocation algorithm for combinatorial auctions. In *Proc. 3rd ACM Conference on Electronic Commerce*, pages 125–136. ACM Press, 2001.