



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

Citation

Parkes, David C., Jayant R Kalagnanam, and Marta Eso. 2001. Achieving budget-balance with Vickrey-based payment schemes in exchanges. In IJCAI-01: Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence: August 4-10, 2001, Seattle, Washington, ed. B. Nebel, International Joint Conferences on Artificial Intelligence, and American Association for Artificial Intelligence, 1161-1168. San Francisco, C.A.: Morgan Kaufman.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:4101693>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story.](#)

[Accessibility](#)

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

David C. Parkes

Computer and Information Science Department
University of Pennsylvania
200 South 33rd Street, Philadelphia, PA 19104
dparkes@unagi.cis.upenn.edu

Jayant Kalagnanam and Marta Eso

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

Abstract

Generalized Vickrey mechanisms have received wide attention in the literature because they are efficient and strategy-proof, i.e. truthful bidding is optimal whatever the bids of other agents. However it is well-known that it is impossible for an exchange, with multiple buyers and sellers, to be efficient and budget-balanced, even putting strategy-proofness to one side. A market-maker in an efficient exchange must make more payments than it collects. We enforce budget-balance as a hard constraint, and explore payment rules to distribute surplus after an exchange clears to minimize distance to Vickrey payments. Different rules lead to different levels of truth-revelation and 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. The scheme appears able to exploit agent uncertainty about bids from other agents to reduce manipulation and boost allocative efficiency in comparison with other simple rules.

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 logical conditions across bundles of different items; for example, an agent might want to buy “*A* and *B*”, or sell “*A* and *B*, or *C*”. Following the literature on combinatorial auctions (Rothkopf *et al.* 1998; de Vries & Vohra 2000) we call this a *combinatorial exchange*. Applications of combinatorial exchanges have been suggested to excess steel inventory procurement (Kalagnanam *et al.* 2000) and to supply chain coordination (Walsh *et al.* 2000).

The market maker in an exchange collects bids and asks and clears the exchange by computing: (i) a set of trades, and (ii) the payments made and received by agents. In designing a mechanism to compute trades and payments we must consider the bidding strategies of self-interested agents, i.e. rational agents that follow expected-utility maximizing strategies. We take as our primary goal that of *allocative-efficiency*: to compute a set of trades that maximize value. In addition, we require:

Copyright © 2002, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

–*individual-rationality* (IR), or voluntary participation, such that all agents have positive expected utility to participate.

–*budget-balance* (BB), such that the exchange does not run at a loss.

Another useful property is *incentive-compatibility* (IC), which states that truthful bidding (submitting bid and ask prices equal to an agent’s value) forms a Bayesian-Nash equilibrium. In other words, every agent can maximize its expected utility by bidding its true values, given that every other agent also bids truthfully. A stronger condition is *strategy-proofness*, such that truthful bidding is optimal whatever the bids of other agents. Strategy-proofness is useful computationally because agents can avoid game-theoretic reasoning about other agents.

Unfortunately, the well-known result of Myerson & Satterthwaite (1983) demonstrates that *no* exchange can be efficient, budget-balanced (even in the average-case), and individual-rational. This impossibility result holds with or without incentive-compatibility¹, and even in Bayesian-Nash equilibrium. Instead, one must:

- (a) impose BB and IR, and design a fairly efficient but incentive-compatible (or perhaps strategy-proof) scheme.
- (b) impose BB and IR, and design a fairly efficient and fairly incentive-compatible scheme.

We follow (b), and design a mechanism for combinatorial exchanges (with multi-unit and regular exchanges as special cases) that promotes reasonable truth-revelation and reasonable allocative-efficiency. The mechanism computes the value-maximizing allocation given agent bids, and computes payments to reduce the utility for non-truthful bidding.

Earlier authors (Myerson & Satterthwaite 1983; McAfee 1992; Barbera & Jackson 1995) have followed approach (a), deliberately computing allocations that are inefficient for truthful bids from agents to achieve *incentive-compatibility* or strategy-proofness. We do not believe their schemes extend easily to combinatorial problems. Furthermore, we believe that our scheme is particularly useful with bounded-rational agents with incomplete information about other agents, because such agents are unable to fully exploit the “holes” for manipulation that remain in our designs.

A Vickrey-Based Payment Scheme

Our particular approach takes the Vickrey payment scheme, and adapts it to make it budget-balanced. Without the prob-

¹As it must, by the revelation principle.

lem of BB, Vickrey payments support an efficient, IR, and strategy-proof exchange.

We interpret Vickrey payments as an assignment of discounts to agents after the exchange clears. BB is achieved so long as the market maker distributes no more than the available surplus when the exchange clears. The pricing problem is formulated as an optimization problem, to compute discounts to minimize the distance to Vickrey discounts. We derive the payment schemes that correspond to optimal solutions to a number of different distance functions.

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 which provides discounts to agents with payments more than a threshold distance than their Vickrey payments.

The Threshold rule increases the amount by which an agent with a large degree of manipulation freedom must adjust its bid to have a useful effect on the price it finally pays, while leaving unaffected the manipulation properties for agents with a small degree of manipulation freedom. The effect is to reduce manipulation and boost allocative-efficiency in comparison with other schemes.

Let us introduce an example problem, that we will return to later in the paper.

Example. 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?

Vickrey Based Surplus Distribution

The market maker in an exchange has two problems to solve: *winner determination*, to determine the trades executed, and *pricing*, to determine agent payments. A common goal in winner-determination is to compute trades that maximize *surplus*, the difference between bid prices and ask prices.² These trades implement the efficient allocation with truthful bids and asks.

The pricing problem is to determine agent payments when the exchange clears. In this section we describe an application of the Vickrey-Clarke-Groves pricing mechanism (Vickrey 1961; Clarke 1971; Groves 1973) to an exchange, which often fails BB. The presentation is for a combinatorial exchange, in which agents can bid and ask for bundles of items and express logical constraints, e.g. “exclusive-or” and “additive-or” constraints, across bids and asks.³

²The payment schemes presented in this paper are also applicable with any (*ex ante* fixed) constraint on feasible trades; e.g. any level of aggregation in matching trades, or side constraints, e.g. on the volume of trade or degree of dominance by a single agent.

³Vickrey payments in exchanges for homogeneous items, with and without multi-unit bids can be derived as special cases (Wur-

Computing payments in a Vickrey-based exchange also requires solving a number of winner-determination problems, once without each agent that trades. Winner-determination is NP-hard for general combinatorial exchange problems and intractable as problems become large. However, our current focus is on the incentive properties of novel Vickrey-based payment schemes. Tractable winner-determination is not our present concern. This noted, the payment schemes proposed are immediately applicable to tractable special cases of combinatorial exchanges (see Kalagnanam et al.) Future work should explore the effect of layering our schemes on top of approximate winner-determination algorithms.

We first define the Vickrey payments in an exchange, and then argue that the failure of BB is quite pervasive with Vickrey payments in exchanges.

Vickrey Payments

Let \mathcal{L} denote the set of agents and $\mathcal{G} = \{A, B, C, \dots\}$ denote the set of items. We need notation for a *trade*; let $T_l \in \{-1, 0, 1\}^{|\mathcal{G}|}$ denote an *indicator vector* for a trade, such that agent l buys items $\{x \mid T_l(x) = 1, x \in \mathcal{G}\}$ and sells items $\{x \mid T_l(x) = -1, x \in \mathcal{G}\}$. Let $\mathbf{T} = (T_1, \dots, T_{|\mathcal{L}|})$ denote a complete trade between all agents.

Bids and asks define a *reported value*, $\hat{v}_l(T_l)$ for a trade T_l , comprising buys and sells. Bids indicate positive value for buying a bundle of items, while asks indicate negative value for selling a bundle of items. For example, if agent 1 submits bid $(AB, 10)$ and ask $(C, 5)$, then $\hat{v}_1([1, 1, 0]) = 10$, $\hat{v}_1([0, 0, -1]) = -5$, $\hat{v}_1([1, 1, -1]) = 5$. The values for other trades are constructed to be consistent with value $-\infty$ for selling anything other than item C , zero value for buying $S \subset \{AB\}$, and no additional value for buying more than bundle AB .

Let \mathbf{T}^* denote the value-maximizing trade, given reported values, $\hat{v}_l(T_l)$, from each agent, with total surplus $V^* = \sum_l \hat{v}_l(T_l^*)$. Trades must be *feasible*, so that supply and demand is balanced, given a model of aggregation. Also, let $(V_{-l})^*$ denote surplus from the value-maximizing trade *without* bids (or asks) from agent l .

By definition, the Vickrey payment to agent l is computed as:

$$p_{\text{vick},l} = (V_{-l})^* - V_{-l}^*$$

where V_{-l}^* is the value of trade \mathbf{T}^* to all agents except agent l , i.e. $V_{-l}^* = V^* - \hat{v}_l(T_l^*)$. Negative payments $p_{\text{vick},l} < 0$ indicate that the agent *receives* money from the exchange after it clears.

We can express an agent’s Vickrey payment as a discount, $\Delta_{\text{vick},l}$, from the payment, $\hat{v}_l(T_l^*)$, the agent would make given its bid and ask prices; i.e. $p_{\text{vick},l} = \hat{v}_l(T_l^*) - \Delta_{\text{vick},l}$, where the *Vickrey discount* is computed as:

$$\Delta_{\text{vick},l} = V^* - (V_{-l})^*$$

The Vickrey discount is always non-negative, representing smaller payments by buyers and higher payments to sellers.

man et al. 1998).

Economic Properties. Vickrey payments are IR, because $V^* \geq (V_{-l})^*$ by a simple feasibility argument, and also strategy-proof. The proof of strategy-proofness is omitted due to lack of space, but closely follows standard Vickrey proofs, for example see Varian & MacKie-Mason (1995). However, BB will often fail in an exchange, as we show in the next section.

Vickrey Budget-Balance: Success & Failure

Now that we have defined Vickrey payments in a combinatorial exchange, let us outline some cases in which BB is achieved and some cases in which BB fails. We will see that budget-balance failure is quite pervasive with Vickrey payments in exchanges.

Standard Exchange. First, consider a standard exchange with bids and asks for single units of a homogeneous item. In this case the exchange is cleared by sorting bids in order of decreasing price and asks in order of increasing price. Bids are matched with asks while the bid price is greater than the ask price. It is well known that Vickrey payments are not BB in this environment.

Let p_{bid}^0 denote the smallest successful bid and p_{bid}^{-1} denote the largest unsuccessful bid. Similarly, let p_{ask}^0 denote the largest successful ask and p_{ask}^{-1} denote the smallest unsuccessful ask. In the Vickrey scheme, every winning seller receives payment $p_{vick, sell} = \min(p_{bid}^0, p_{ask}^{-1})$, whatever its own ask price, and every winning buyer pays $p_{vick, buy} = \max(p_{ask}^0, p_{bid}^{-1})$, whatever its own bid price. The following condition is required for BB:

Claim 1. *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) $p_{bid}^0 = p_{ask}^0$; (2) $p_{bid}^0 = p_{bid}^{-1}$; (3) $p_{ask}^0 = p_{ask}^{-1}$.*

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

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. Thus, we cannot expect BB with Vickrey payments even in a standard (non-combinatorial) exchange except in special cases.

Combinatorial Exchange As an example of BB failure, consider that agents submit truthful bids in the earlier example; i.e. asks $(A, \$10)$, $(B, \$5)$ and bids $(AB, \$51)$, $(AB, \$40)$. $V^* = 51 - 10 - 5 = 36$, $(V_{-1})^* = (V_{-2})^* = 0$, $(V_{-3})^* = 25$, and $(V_{-4})^* = 36$. Agent 1’s Vickrey payment is $-10 - (36 - 0) = -46$, agent 2’s is $-5 - (36 - 0) = -41$, agent 3’s is $51 - (36 - 25) = 40$. The exchange runs at a loss of \$47 to the market maker.

One-Sided Vickrey-Payments First, a positive special-case. Claim 2 gives a sufficient condition for BB in the special-case that Vickrey discounts are only allocated to

agents on one-side of an exchange, i.e. to all buyers or to all sellers (but not to buyers and sellers).

We define *aggregation* on the sell-side as when bids from multiple buyers can be combined to match an ask from a single seller, and aggregation on the buy-side as when asks from multiple sellers can be combined to match a bid from a single buyer.

Claim 2. *Budget-balance holds if Vickrey payments are implemented on one-side of an exchange, and when that side has no aggregation.*

Proof sketch. Simple, just show that this BB holds for each “cluster” of trading agents, and therefore for the entire exchange.

Bilateral matching is a special-case, with no aggregation on either side; i.e. Vickrey payments are budget-balanced if implemented for at most one agent in each trade, for example with trades cleared at either the ask price (buy-side strategy-proofness) or the bid price (sell-side strategy-proofness). Similarly, the single-item Vickrey auction is a special case (and strategy-proof to buyers but not the seller).

The Generalized Vickrey Auction (GVA) is the VCG mechanism for a combinatorial auction, in which there is a single seller and sell-side aggregation. The GVA is BB because the buyers, but not the seller, receive Vickrey payments. The auctioneer simply collects the total payment made by the buyers and passes it on to the seller. As such the GVA is strategy-proof for buyers but not for the seller. Another problem is that the seller can sometimes receive less than her ask price. Consider a seller with an ask price of $(AB, \$10)$ and bids of $(A, \$8)$ and $(B, \$8)$ from different buyers. Each buyer receives Vickrey discount \$6 and pays \$2, but the seller needs at least \$10.

One-to-N models We can state a general negative result for Vickrey payments to all agents (buyers and sellers) in a combinatorial auction.

Claim 3. *Budget-balance fails with Vickrey payments to all agents in a combinatorial auction except in the case that no buyer requires a Vickrey discount.*

Proof sketch. Simple, just show that the seller extracts *all* of the surplus as its Vickrey discount.

Intuitively, BB fails in this case unless the marginal value contributed by each buyer is zero, i.e. unless the surplus in the exchange is the same with any one of the buyers removed.

Budget-Balanced Payment Rules

In this section 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 a later section we demonstrate useful truth-revelation properties for the Vickrey-based schemes.

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.
- Present a theoretical analysis of each payment scheme in a simple bidding model.

Mathematical Programming Model

We formulate the pricing problem as a linear program, to assign surplus to agents to minimize distance to Vickrey discounts. Let V^* denote the available surplus when the exchange clears, before any discounts, and $L^* \subseteq \mathcal{L}$ denote the set of agents that trade. Each agent may perform a number of buys and sells, depending on its bids and asks of other agents. We compute discounts $\Delta = (\Delta_1, \dots, \Delta_L)$ to minimize the distance $\mathbf{L}(\Delta, \Delta_{\text{vick}})$ to Vickrey discounts, for a suitable distance function \mathbf{L} .

$$\begin{aligned} \min_{\Delta} \quad & \mathbf{L}(\Delta, \Delta_{\text{vick}}) && [\text{PP}] \\ \text{s.t.} \quad & \sum_{l \in L^*} \Delta_l \leq V^* && (\text{BB}) \\ & \Delta_l \leq \Delta_{\text{vick},l} \quad , \forall l \in L^* && (\text{VD}) \\ & \Delta_l \geq 0 \quad , \forall l \in L^* && (\text{IR}) \end{aligned}$$

Constraint (BB) gives worst-case (or ex post) budget-balance, such that the exchange never makes a net payment to agents. We might also substitute an expected surplus \bar{V}^* for V^* and implement average-case (or ex ante) budget-balance. Constraints (IR) ensure that truthful bids and asks are individual-rational for an agent, with a worst-case (or ex post) non-negative expected utility. Constraints (VD) ensure that no agent receives more than its Vickrey discount.⁴

In addition to the standard \mathbf{L}_2 and \mathbf{L}_{∞} distance metrics, we also consider the following functions: (a) $\mathbf{L}_{\text{RE}}(\Delta, \Delta_{\text{vick}}) = \sum_l \frac{\Delta_{\text{vick},l} - \Delta_l}{\Delta_{\text{vick},l}}$, a relative error function; (b) $\mathbf{L}_{\Pi}(\Delta, \Delta_{\text{vick}}) = \prod_l \frac{\Delta_{\text{vick},l}}{\Delta_l}$, a product error function; (c) $\mathbf{L}_{\text{RE2}}(\Delta, \Delta_{\text{vick}}) = \sum_l \frac{(\Delta_{\text{vick},l} - \Delta_l)^2}{\Delta_{\text{vick},l}}$, a squared relative error function; (d) $\mathbf{L}_{\text{WE}}(\Delta, \Delta_{\text{vick}}) = \sum_l \Delta_{\text{vick},l} / (\Delta_{\text{vick},l} - \Delta_l)$, a weighted error function. The \mathbf{L}_1 metric provides no distributional information.

We drop agents with $\Delta_{\text{vick},l} = 0$ from all models, and simply set $\Delta_l = 0$ for these agents.

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*. For example, the *Threshold* rule, $\Delta_l = \max(0, \Delta_{\text{vick},l} - C)$ for some parameter $C \geq 0$, solves [PP] for the \mathbf{L}_2 distance metric. For large C , Threshold allocates small, or no, discounts, while for $C = 0$ Threshold allocates Vickrey discounts.

To understand the construction of Threshold from \mathbf{L}_2 consider the simplest case, when constraints (VD) and (IR) are

⁴The (VD) constraints are not redundant for certain distance metrics, such as the $\mathbf{L}_{\text{RE}}(\cdot)$ metric.

Distance Function	Payment Scheme	Discount Definition
$\mathbf{L}_2, \mathbf{L}_{\infty}$	Threshold	$\max(0, \Delta_{\text{vick},l} - C)$
\mathbf{L}_{RE}	Small	$\Delta_{\text{vick},l}$ if $\Delta_{\text{vick},l} \leq C$
\mathbf{L}_{RE2}	Fractional	$\mu \Delta_{\text{vick},l}$
\mathbf{L}_{WE}	Large	$\Delta_{\text{vick},l}$ if $\Delta_{\text{vick},l} \geq C$
\mathbf{L}_{Π}	Reverse	$\min(\Delta_{\text{vick},l}, C)$
-	No-Discount	0
-	Equal	D

Table 1: Distance Functions and Payment Schemes.

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
Agent 3	40	39	46.2	51	40	51	40

Table 2: Payments with Different Rules in the Simple Problem.

not binding, and perform Lagrangian optimization. Dropping the outer square root from the \mathbf{L}_2 metric and introducing Lagrange multiplier $\lambda \geq 0$, we have $\min \sum_l (\Delta_{\text{vick},l} - \Delta_l)^2 + \lambda(\sum_l \Delta_l - V^*)$. Now, computing first derivatives w.r.t. Δ_l and setting to zero, we have $-2(\Delta_{\text{vick},l} - \Delta_l) + \lambda = 0$ for all l .⁵ Solving, this equates the distance to Vickrey discounts across all agents, $\Delta_{\text{vick},1} - \Delta_1 = \Delta_{\text{vick},2} - \Delta_2 = \dots$, and with budget-balance we find $\Delta_l = \Delta_{\text{vick},l} - (\sum_{l'} \Delta_{\text{vick},l'} - V^*) / |L^*|$. This is the *Threshold* rule with parameter $C = (\sum_l \Delta_{\text{vick},l} - V^*) / |L^*| = \lambda/2$.

Table 1 tabulates the payment rules for each distance function, and also includes the *Equal* rule which is not Vickrey-based but divides surplus equally across all agents, and the *No-Discount* rule (see also Figure 1). Each payment rule is parameterized with $C \geq 0$, except for *Fractional*, which has parameter $0 \leq \mu \leq 1$. The parameters that give BB in each scheme can be easily computed from Vickrey discounts and available surplus.

Example. In Table 2 we compare the payments made with each payment scheme in our simple problem. Notice that neither the Large or Small schemes provide useful guidance about how to distribute the discount across the two sellers, this depends on how the tie is broken.

Theoretical Analysis

In this section we develop simple analytic results for the amount of manipulation an agent will perform with each payment scheme. The model permits tractable analysis, and proves interesting both for the insight it provides and for the close correspondence that we find with later experimental results for combinatorial exchange problems.

We choose to analyze an exchange in which bids and asks are for *single items*. Later, in our experimental analysis we compare the payment schemes in combinatorial problems.

For buyers (the analysis is symmetric for sellers):

- (1) Every agent l has value v_l for a single item (drawn from some distribution $F_l(v)$) and chooses to manipulate by

⁵First-order conditions are necessary and sufficient for optimality in this problem because the Hessian is positive definite.




Figure 1: Bid price b_l (x-axis) against adjusted bid price $b_l - \Delta_l$ (y-axis) in each payment scheme. Agent value v , highest outside bid x , Payment scheme parameters C, μ, D .

$\theta_l \geq 0$, and bid $b_l = v_l - \theta_l$.

- (2) The maximum bid x_l from another agent for the item, or ask price (whichever is higher), is uniformly distributed about v_l , i.e. $x_l \sim U(v_l - \delta, v_l + \delta)$ for some constant, $\delta > 0$.
- (3) The average surplus available to the market maker when the exchange clears is $\alpha\delta$ per-agent, for some constant $\alpha > 0$ that defines the amount of surplus.
- (4) In equilibrium, the market maker selects a parameter (e.g. C) for the payment scheme to achieve average-case budget-balance. Payment rules are computed *before* agents bid, and the parameters are known to bidding agents.

Agent l has a quasi-linear utility function, $u_l = v_l - p$, for submitting the highest bid where p is its payment to the exchange, i.e. $p = b_l - \Delta_l$. Figure 1 illustrates each payment rule in this simple model, plotting bid price b_l against adjusted price $b_l - \Delta_l$; e.g., in Vickrey the agent pays only x for any bid $b_l \geq x$, in Threshold the agent pays $b_l = x + C$ for $b_l \geq x + C$, and b_l for $x \leq b_l < x + C$, given parameter C , etc.

For each payment scheme we determine: (a) an agent's optimal bidding strategy as a function of the parameters of the rule, e.g. C or μ ; and (b) the equilibrium parameterization of the rule, e.g. value for C , that leads to budget-balance given that agents follow this optimal bidding strategy. The analysis leads to a relationship between the *available surplus and the degree of manipulation* for each payment rule (see Figure 4).

One can be critical of our assumptions. We leave undefined both the valuation distribution functions $F_l(v)$ and the distribution that defines the item an individual agent values. It is quite likely that there are no $F_l(v)$ that are consistent with our assumption of a uniformly distributed second-highest bid in equilibrium. In addition, we adopt average-case budget-balance and compute payment rules before agents bid, but ignore any effect that rules have on surplus via agents' bidding strategies.

However, we believe that this analysis has significant value. Its main success is that it clearly demonstrates the effect that different types of budget-balanced Vickrey-based payment rules can have on agent manipulation. We leave a

full equilibrium analysis for future work.

Graphical Intuition. Manipulation 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 2 we plot the *utility* for a particular bid, $b = v - \theta$, as the value of the outside bid x varies, for payment schemes Vickrey, No-Discount, Threshold and Fractional. Each subplot is for a single scheme, with individual curves corresponding to different bids.⁶




Figure 2: Utility of bids $b = v - \theta$ with $\theta = \{0, 0.3, 0.5\}$, $v = 1$, as the best outside bid x varies between 0 and 1. $C = 0.4$ in Threshold, and $\mu = 0.5$ in Fractional.

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 outside bids (closer than C 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).

Making our assumption about the distribution of x around an agent's value v_l , we can compute the expected utility for different levels of manipulation under each scheme as the area under a particular curve in a plot like Figure 2.⁷ The expected-utility maximizing bid corresponds to the curve with maximum area. In Figure 3 we plot the expected gain in utility (in comparison with truthful bidding), $Eu(\theta) - Eu(0)$, for bid $b = v - \theta$ in each payment rule. Rule parameters are set to give BB with surplus $\alpha = 0.1$ at optimal agent strate-

⁶Although 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.

⁷It is at this stage that an equilibrium analysis would need to use a derived expression for the distribution of x .




Figure 3: Expected Gain in Utility for different bids $b = v - \theta$ under each payment scheme, with rule parameters set to give BB with surplus $\alpha = 0.1$.

gies. Notice that the level of manipulation, θ^* , that maximizes the agent's gain in utility is smallest in the Threshold scheme for this value of surplus.

The results (below) show that the Large and Threshold rules perform well in this model, and lead to the following intuitive remarks about payment rules (see Figure 1):

1. A large flat section for bids close to the agent's true value is useful, i.e. with adjusted bid price independent of the agent's bid price.
2. Nowhere should the adjusted bid price be greater than the agent's bid price (for IR with truthful bidding), which constrains the line to lie to the right of the "no-discount" line.
3. It is more important to implement the flat section for values, v , far from the highest outside bid, x , than values close to the highest outside bid (i.e. *Large* rather than *Small*), because manipulation is already more risky for true values close to v than far from x .⁸

It is useful to think about the "degree of manipulation freedom" available to an agent, which in this simple single-bundle model is the difference between an agent's value v and the highest outside bid x . In general, this is simply measured by the Vickrey discount to an agent that bids truthfully, i.e. the amount by which it could have reduced its bid price and still participated in the same trades. The Large and Threshold schemes are effective because they make manipulation more difficult and less useful for an agent with a large degree of manipulation freedom, while leaving the ability to manipulate of agents with a small degree of manipulation freedom unchanged. This is a good incentive strategy because it attacks the "low risk" manipulation opportunities, but leaves the "high risk" opportunities. Agents are uncertain about the bids from other agents and always run the risk of bidding too low and forfeiting a profitable trade.

Results. Table 3 summarizes the analytical results, giving an agent's optimal bidding strategy, θ^* , as a function of the parameter in each scheme, and the expected discount per-

⁸Note that in terms of efficiency the picture is mixed. While we can stand more manipulation from agents with large values compared to x , without changing the trades that we implement, if the bids from those agents does change the final implementation the effect on efficiency is likely to be quite large.

Rule	Optimal Manipulation, θ^*	Expected Discount
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	$\min[C, \delta/2]$	$\max\left[0, \frac{(\delta-2C)^2}{4\delta}\right]$
Equal	$\frac{\delta-D}{2}$	$\frac{D(D+D)}{4\delta}$
Small	$\max[0, \min(\delta/2, \delta - C)]$	$\min[\delta/4, C^2/4\delta]$
Large	0, if $C \leq \delta/\sqrt{2}$ $\delta/2$, otherwise	$-C^2/4\delta + \delta/4$, if $C \leq \delta/\sqrt{2}$ 0, otherwise
Reverse	$\max[0, \frac{\delta-C}{2}]$	$\min[\delta/4, C/4]$

Table 3: Analytical results.




Figure 4: Optimal agent manipulation, $\theta^*(\alpha)$, (as a proportion of δ) under each payment scheme as the amount of available surplus varies from 0 to $\delta/2$ per-agent.

agent given that optimal strategy.⁹ We present an example derivation, for the Threshold rule, below.

In Figure 4 we enforce BB, computing parameters in the payment schemes to set the expected discount equal to surplus $\alpha\delta$, and plot the equilibrium manipulation performed in each payment scheme as the amount of surplus varies. The Vickrey payment scheme can be implemented with surplus $\delta/4$ per-agent, so all schemes except Equal and No-Discount prevent manipulation completely for $\alpha \geq 0.25$. For smaller amounts of surplus the market maker is forced to deviate from Vickrey, and move left in Figure 4. At $\alpha = 0$ no schemes can provide any discount, and the agent manipulates by $\delta/2$.

First, notice that the simple minded Equal scheme appears to have 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. Agent uncertainty coupled with the risk of bidding too low and either falling from the flat section or under-bidding the highest outside bid lead agents to bid truthfully.

⁹It is useful to confirm that all expressions reduce to that for the Vickrey and No-Discount rules at extreme parameter values (e.g. $\mu = \{0, 1\}$ in Fractional, $C = \{0, \delta/2\}$ in Threshold, etc.)

Example Derivation: Threshold Rule. Each agent receives discount $\max(0, b - (x + C))$, for some constant $C > 0$. The agent's utility given bid $v - \theta$, value v , highest outside bid x , and Threshold C , is computed as:

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

Assume that $C \leq \delta$, so that the agent will receive a discount for some choice of $\theta < \delta$. Consider three cases.

Case (1), $\theta \leq \delta - C$. The expected utility for bid θ given C , $\text{EU}(\theta, C)$ is:

$$\begin{aligned} & \int_{x=v-\delta}^{v-\theta-C} [v - (x + C)] f(x) dx + \int_{x=v-\theta-C}^{v-\theta} [v - (v - \theta)] f(x) dx + \int_{x=v-\theta}^{v+\delta} 0 f(x) dx \\ &= \frac{(\delta - \theta - C)}{2\delta} (v - C) - \frac{1}{4\delta} [(v - \theta - C)^2 - (v - \delta)^2] + \frac{\theta}{2\delta} (\theta + C - \theta) \end{aligned}$$

In *case (2)*, $\delta - C < \theta \leq \delta$, $\text{EU}(\theta, C) = \int_{x=v-\delta}^{v-\theta} [v - (v - \theta)] f(x) dx = \theta(\delta - \theta)/2\delta$. In *Case (3)*, $\delta < \theta$, then $\text{EU}(\theta, C) = 0$. From this, the agent's optimal bidding strategy, denoted $\theta^*(C)$, by differentiation w.r.t. θ and case analysis, is:

$$\theta^*(C) = \min [C, \delta/2]$$

The discount to the agent for bid $b = v - \theta$ is: $\Delta(v - \theta, v, x, C) = \max[0, v - \theta - (x + C)]$. The expected discount, first in the case that $\theta \leq \delta - C$, is:

$$\begin{aligned} E\Delta(\theta, C) &= \int_{x=v-\delta}^{v-\theta-C} [v - \theta - (x + C)] f(x) dx \\ &= (v - \theta - C) \frac{(\delta - \theta - C)}{2\delta} - \frac{1}{4\delta} [(v - \theta - C)^2 - (v - \delta)^2] \end{aligned}$$

or $E\Delta(\theta, C) = 0$ in the case $\theta > \delta - C$. Substituting for the agent's optimal bidding strategy $\theta^*(C)$ we have: $E\Delta(\theta^*(C), C) = \max\left[0, \frac{(\delta - 2C)^2}{4\delta}\right]$. Now, with per-agent surplus $\alpha\delta$ and budget-balance, such that $E\Delta^*(C) \leq \alpha\delta$, the exchange should set $C^*(\alpha) = \min\left[0, \frac{\delta}{4}(2 - \sqrt{16\alpha})\right]$ to minimize manipulation.

Experimental Analysis

In this section we provide an experimental analysis of the payment schemes in a set of combinatorial problem instances. Agents are either buyers or sellers, and values are assigned to agents for *bundles* following the Random, Weighted Random, Decay, and Uniform distributions from Sandholm (1999), adapted in this case to a combinatorial exchange. Each agent submits bids (asks) for multiple bundles, with exclusive-or constraints across bids (asks). We test problems with 5, 10, and 20 agents, a total of 100 bids and asks (evenly distributed across agents), 50 goods, and with different proportions of buyers and sellers.¹⁰

¹⁰Results are averaged over 80 problem instances, for numbers of Buyer/Sellers $\in \{5/5, 7/3, 2/3, 4/1, 10/10, 15/5\}$.




Figure 5: 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.

In our theoretical model we adopted average-case budget-balance to make the analysis tractable. We now revert to the more natural worst-case (or every-time) budget-balance in which the market maker distributes exactly the available surplus every time the exchange is cleared. Payment rules are now computed *after* bids are received.

We perform a limited strategic analysis. First, we assume that the strategy of agent l is to adjust *all* its bids and asks by the same fractional amount, $y_l\%$, i.e. submitting bid prices $y_l\%$ below value and ask prices $y_l\%$ above value. Second, we look for a *symmetric* Nash equilibrium in which every agent follows the same strategy, for some $y\%$. Finally, we compute an *approximation* to this equilibrium for computational tractability. We compute the average utility gain to a single agent for 0% vs. $y\%$ manipulation, given that every other agent manipulates by $y\%$, and determine the amount of manipulation, y^* , that maximizes this utility gain. We assume that this is also the optimal strategy for an individual agent against a population of agents with *fixed* strategies y^* , and therefore the Nash equilibrium.¹¹

Given this, we read off the symmetric Nash equilibrium under a particular payment rule as the peak of a plot such as that in Figure 5, which plots the gain in utility for strategy $y\%$ vs. 0% in a system in which every agent follows strategy $y\%$, in this case for the 5 buyers/5 sellers problem set. In this case, notice that the equilibrium manipulation level in Large and Threshold is less than under the other rules, in this case 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 Table 4 we summarize the results of experiments across all problem sets. We compare: the average utility gain, and the correlation with Vickrey discounts, at manipulation levels of 10%, 20% and 30% in each scheme; and the average optimal degree of manipulation by agents in each scheme,

¹¹One benefit of this technique is that we have a method to measure the degree of manipulation even when there is in fact *no* symmetric pure Nash equilibrium.

	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 4: Experimental results. Utility gain and Correlation with Vickrey discounts computed for manip. 10%, 20% and 30%, and averaged over all problem instances (for 5–20 agents).

and the corresponding allocative efficiency. The allocative efficiency in the Large and Threshold schemes is considerably higher than in the other schemes.

Discussion

The partial ordering {Large, Threshold } \succ Fractional \succ Reverse \succ {Equal, Small } from the experimental results is remarkably consistent with the results of our theoretical analysis. Although the Large scheme generates slightly less manipulation and higher allocative efficiency than Threshold, the correlation between discounts and Vickrey discounts is much greater in Threshold than Large. An agent’s discount in Large is very sensitive to its bid, and we expect Large to be less robust than Threshold in practice because of this all-or-nothing characteristic.

As discussed earlier, we have made a number of assumptions, both in the analytic models of agent manipulation and also in the manipulation structure considered experimentally. In addition to understanding the effects of these assumptions, in future work we would also like to: quantify worst-case and average-case utility gains from manipulation in each payment scheme, given a particular amount of surplus; and derive *optimal* payment schemes, for example minimizing worst-case gains from manipulation. One avenue is to ask how bad would the efficiency get if every agent was perfectly informed about the other agents, and followed a best-possible bidding strategy given the payment rules. Finally, we suspect that stochastic payment rules might prove to have interesting incentive properties.

Conclusions

We constructed budget-balanced payment schemes to minimize different distance functions to Vickrey payments, and showed analytically and experimentally that a simple Threshold rule has better incentive properties than other payment schemes. The effect of the payment scheme is to implement a distribution of manipulation-preventing discounts across a population of agents to exploit an agent’s inherent uncertainty about bids from other agents and the degree to which manipulation can be useful. The Threshold rule increases the amount by which an agent with a large degree of manipulation freedom must adjust its bid to have a useful effect on the price it finally pays, while leaving unaffected

the manipulation properties for agents with a small degree of manipulation freedom.

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.

Acknowledgments

We thank William Walsh, and the anonymous reviewers, for their helpful comments and suggestions. This research was funded in part by National Science Foundation Grant SBR 97-08965. The first author gratefully acknowledges financial support from an IBM Research Fellowship.

References

- Barbera, S., and Jackson, M. O. 1995. Strategy-proof exchange. *Econometrica* 63(1):51–87.
- Clarke, E. H. 1971. Multipart pricing of public goods. *Public Choice* 11:17–33.
- de Vries, S., and Vohra, R. 2000. Combinatorial auctions: A brief survey. Tech. report, MEDS Department, Kellogg Graduate School of Management, Northwestern University.
- Groves, T. 1973. Incentives in teams. *Econometrica* 41:617–631.
- Kalagnanam, J. R.; Davenport, A. J.; and Lee, H. S. 2000. Computational aspects of clearing continuous double auctions with assignment constraints and indivisible demand. IBM Research Report RC 21660 (97613). To appear in *Electronic Commerce Research Journal*.
- McAfee, R. P. 1992. A dominant strategy double auction. *J. of Economic Theory* 56:434–450.
- Myerson, R. B., and Satterthwaite, M. A. 1983. Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 28:265–281.
- Rothkopf, M. H.; Pekeč, A.; and Harstad, R. M. 1998. Computationally manageable combinatorial auctions. *Management Science* 44(8):1131–1147.
- Sandholm, T. 1999. An algorithm for optimal winner determination in combinatorial auctions. In *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, 542–547.
- Varian, H., and MacKie-Mason, J. K. 1995. Generalized Vickrey auctions. Tech. report, University of Michigan.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16:8–37.
- Walsh, W.; Wellman, M.; and Ygge, F. 2000. Combinatorial auctions for supply chain formation. In *Proc. ACM Conference on Electronic Commerce*, 260–269.
- Wurman, P. R.; Walsh, W. E.; and Wellman, M. P. 1998. Flexible double auctions for electronic commerce: Theory and implementation. *Decision Support Systems* 24:17–27.