

Balanced Trade Reduction for Dual-Role Exchange Markets

Dengji Zhao, Sarvapali D. Ramchurn, Enrico H. Gerding and Nicholas R. Jennings

Electronics and Computer Science
University of Southampton
Southampton, SO17 1BJ, UK
{d.zhao, sdr, eg, nrj}@ecs.soton.ac.uk

Abstract

We consider dual-role exchange markets, where traders can offer to both buy and sell the same commodity in the exchange but, if they transact, they can only be *either* a buyer *or* a seller, which is determined by the market mechanism. To design desirable mechanisms for such exchanges, we show that existing solutions may not be *incentive compatible*, and more importantly, cause the market maker to suffer a significant deficit. Hence, to combat this problem, following McAfee's trade reduction approach, we propose a new trade reduction mechanism, called *balanced trade reduction*, that is incentive compatible and also provides flexible trade-offs between efficiency and deficit.

Introduction

Exchange markets (aka double auctions) are the most important institutions for modern economy, which are centralized markets consisting of exchange rules for traders to buy and sell commodities, e.g. stock exchanges. Most existing studies of exchanges are for the environments where a trader is either a buyer or a seller, but not both, of certain commodities (Myerson and Satterthwaite 1983; McAfee 1992; Wurman, Walsh, and Wellman 1998; Blum, Sandholm, and Zinkevich 2006; Bredin, Parkes, and Duong 2007; Parsons, Rodriguez-Aguilar, and Klein 2011).

In this paper, we study another type of exchanges. We call it *dual-role exchange*, where a trader offers to both buy and sell of certain commodities but when she transacts, she can be only a seller or a buyer, which is determined by the market maker. There exist many real-world applications of dual-role exchanges. For example, in ridesharing/carpooling (Kamar and Horvitz 2009; Zhao et al. 2014), a driver would like to sell his empty seats to riders at a reasonable price and may also be willing to ride with other drivers (i.e., buy seats from others) if the price they charge is lower than his driving cost. Also, in electric vehicle (EV) charging (Kempton and Tomi 2005; Han, Han, and Sezaki 2010), an EV owner might buy more electricity in order to finish a trip if the price is reasonable, otherwise, she might sell the power stored in the car and use other transportations instead. All these applications share a common feature: an agent cannot be both buying

and selling the same commodity at the same time, although she be able to *either* buy *or* sell the same commodity. In the ridesharing scenario, a commuter cannot drive and ride at the same time, and in the EV charging scenario, an EV cannot be charged and discharged at the same time. Another important feature of these domains is that a trader do not decide to buy or sell by herself if she is flexible to either buy or sell, which also gives us more opportunity to optimize social welfare. For example, in ridesharing, if every commuter has to decide to either drive or ride before she participates in the market, then if everyone decides to drive, we will end up with no sharing. However, if every (or some) driver is willing to ride with others, then we will be able to let some drivers ride with others, which will then make the sharing possible to save costs, reduce congestion and etc..

In designing an exchange mechanism, it is important to achieve a number of desirable properties, namely: maximizing social welfare (i.e., *efficient*), preventing manipulations of agents (i.e., *truthful*), an agent never pays more than what she gets (i.e., *individually rational*) and the market maker should not run the mechanism with a deficit (i.e., *budget balanced*). It is well known that designing an exchange mechanism that is efficient, truthful, individually rational and budget balanced is impossible (Myerson and Satterthwaite 1983). Since a loss-making mechanism does not make economic sense in reality, we seek to find budget balanced mechanisms that can be applied in the real-world. In order to achieve budget balance (i.e. removing deficit), there have been several attempts via trading efficiency (Parkes, Kalagnanam, and Eso 2001; Babaioff and Walsh 2005). One approach that stands out from classical exchange mechanisms is McAfee's trade reduction and its generalizations, which are more relevant to this study (McAfee 1992; Gonen, Gonen, and Pavlov 2007). The essential idea of McAfee's trade reduction is to reduce efficiency as little as possible to get budget balanced, truthful, and individually rational mechanisms.

Following McAfee's approach, we design budget balanced mechanisms for a dual-role exchange. We limit to the situation where each trader buys or/and sells a single-unit of a commodity. We first show that McAfee's trade reduction does not satisfy truthfulness in this type of exchange, although it is truthful under a special setting, called marginally decreasing valuation domain. We therefore propose a new

trade reduction to fix the incentive problem of McAfee's trade reduction. However, this fix might induce some deficit. In order to further control the deficit, we generalize our proposed reduction mechanism to provide the flexibility to trade off between efficiency and deficit.

The remainder of the paper is as follows. We first introduce the dual-role exchange model and the desirable properties for designing mechanisms for the model. Then we show the performance (good and bad) of McAfee's trade reduction in this model. Following that, we propose our new trade reduction mechanism. Finally, we generalize the new mechanism for more flexible deficit control.

The Model

We consider an exchange where a set of **traders** N exchange one type of commodity. Each trader $i \in N$ has a valuation vector (v_i^b, v_i^s) where $v_i^b \geq 0$ is i 's valuation for buying the commodity and $v_i^s \leq 0$ is i 's valuation for selling the commodity. For example, in ridesharing, a driver with one empty seat is willing to sell it for \$2, and she is also willing to buy one seat from other drivers for \$3. Therefore, the driver's valuation is $(3, -2)$. If trader i does not want or does not have the commodity to sell, then $v_i^s = -\infty$. If i does not want to buy, then $v_i^b = 0$. We say i 's valuation is **marginally decreasing** if $v_i^b < -v_i^s$. In this setting, we consider that no trader is allocated to sell and buy at the same time as discussed in the introduction, otherwise the problem can be solved as a classical exchange by splitting the trader into a seller and a buyer. Let θ_i denote the valuation of trader i , θ be the valuation profile of all traders and θ_{-i} be the valuation profile of all traders except i .

We study auction mechanisms requiring each trader to report her valuation (not necessarily the truthful valuation) to the mechanisms. In the rest of this paper, let $\theta_i = (v_i^b, v_i^s)$ indicate i 's true valuation and $\hat{\theta}_i$ be i 's reported valuation. A **mechanism** consists of an **allocation** π and a **payment** x . Given traders' valuation report profile $\hat{\theta}$, allocation $\pi_i(\hat{\theta}) \in \{\text{buyer, seller, unmatched}\}$ denotes the transaction of i : i receives one commodity if $\pi_i(\hat{\theta}) = \text{buyer}$, gives one commodity out if $\pi_i(\hat{\theta}) = \text{seller}$, and does no trade if $\pi_i(\hat{\theta}) = \text{unmatched}$. $x_i(\hat{\theta}) \in \mathcal{R}$ is the payment for trader i and $x_i(\hat{\theta}) < 0$ means that i receives $-x_i(\hat{\theta})$ from the mechanism. An allocation π is **feasible** if the number of buyers is the same as the number of sellers and the sellers' selling valuation is not $-\infty$. Note that even if a trader's buying valuation is 0, she can still be allocated as a buyer (i.e. we assume *free disposal*). In the rest of the paper, we only consider feasible allocations. Let us first define several key properties for the mechanism to satisfy in the following.

Given traders' valuation report profile $\hat{\theta}$ and allocation π , we define the **valuation** of i as:

$$v(\theta_i, \pi_i(\hat{\theta})) = \begin{cases} v_i^b & \text{if } \pi_i(\hat{\theta}) = \text{buyer,} \\ v_i^s & \text{if } \pi_i(\hat{\theta}) = \text{seller,} \\ 0 & \text{if } \pi_i(\hat{\theta}) = \text{unmatched.} \end{cases}$$

Note that the valuation here for holding some initial endowment is zero and becomes negative if one sells it, which is

the same as if we have a positive valuation for holding the initial endowment and becomes zero after one sells it.

Given an allocation, the sum of valuations of all traders is referred to as the **social welfare**. We say an allocation is **efficient** if it maximizes the social welfare for any valuation report profile.

Definition 1. An allocation π is **efficient** if for all θ , $\pi(\theta) \in \arg \max_{\pi' \in \Pi} \sum_{i \in N} v(\theta_i, \pi'_i(\theta))$, where Π is the set of all feasible allocations.

Given traders' valuation report profile $\hat{\theta}$ and mechanism (π, x) , the **utility** of i is quasilinear and is defined as:

$$u(\theta_i, \hat{\theta}, (\pi, x)) = v(\theta_i, \pi_i(\hat{\theta})) - x_i(\hat{\theta}).$$

We say mechanism (π, x) is **individually rational** if $u(\theta_i, (\theta_i, \hat{\theta}_{-i}), (\pi, x)) \geq 0$ for all $i \in N$, all θ_i , and all $\hat{\theta}_{-i}$, i.e. no trader will get negative utility via truthfully reporting no matter what others do. We say (π, x) is **truthful** or **incentive compatible** if i 's utility is maximized via truthfully reporting, i.e. reporting truthfully is a dominant strategy for each trader.

Definition 2. A mechanism (π, x) is **truthful** if $u(\theta_i, (\theta_i, \hat{\theta}_{-i}), (\pi, x)) \geq u(\theta_i, (\hat{\theta}_i, \hat{\theta}_{-i}), (\pi, x))$ for all $i \in N$, all θ_i , all $\hat{\theta}_i$, and all $\hat{\theta}_{-i}$.

Lastly, we say mechanism (π, x) is **budget balanced** if $\sum_{i \in N} x_i(\hat{\theta}) = 0$ for all valuation report profile $\hat{\theta}$, or **weakly budget balanced** if $\sum_{i \in N} x_i(\hat{\theta}) \geq 0$. If $\sum_{i \in N} x_i(\hat{\theta}) < 0$, the mechanism runs a **deficit** $|\sum_{i \in N} x_i(\hat{\theta})|$.

The well-known Vickrey-Clarke-Groves (VCG) mechanism is applicable in this model, which is efficient, truthful, individually rational but not budget balanced. In VCG, each agent pays the harm he causes to others. Formally, given valuation report profile θ and efficient allocation π , the VCG payment for trader i is defined as:

$$x_i^{vcg}(\theta) = V(\theta_{-i}, \pi) - V_{-i}(\theta, \pi) \quad (1)$$

where

- $V(\theta_{-i}, \pi) = \sum_{j \in N \setminus \{i\}} v(\theta_j, \pi_j(\theta_{-i}))$, i.e. the social welfare for all traders, excluding i , of the efficient allocation without considering i 's report.
- $V_{-i}(\theta, \pi) = V(\theta, \pi) - v(\theta_i, \pi_i(\theta))$, i.e. the social welfare for all traders, excluding i , of the efficient allocation considering all traders' reports θ .

As Myerson and Satterthwaite (1983) showed, it is impossible to design an exchange mechanism that is efficient, truthful and individually rational without outside subsidies. Since a mechanism running a significant deficit is hardly applicable in real-world applications, we will search for mechanisms with deficit control. There have been many attempts to circumvent the impossibility by giving up efficiency for budget balance under different settings (McAfee 1992; Babaioff and Nisan 2001; Babaioff, Nisan, and Pavlov 2004; Gonen, Gonen, and Pavlov 2007). In the next section, we show that the well-known McAfee's trade reduction mechanism cannot be applied in general in this model and therefore, we need to design new solutions for this model.

McAfee's Trade Reduction

In order to get budget balance, McAfee (1992) proposed a truthful trade reduction mechanism for the setting where a trader can either sell or buy, but not both. The idea is to remove the pair with the lowest buying and selling valuations from the efficient allocation, if necessary, to set up the payments to the other buyers and sellers. McAfee's trade reduction can be extended to this model as follows:

McAfee's trade reduction $\mathcal{M}^{McAfee}$

1. Given traders' valuation report profile, compute the efficient allocation, and let v_0^b, v_0^s be the lowest valuations for buying and selling of all (matched) buyers and sellers respectively (it is evident that $v_0^b \geq -v_0^s$).
2. Let v_{-1}^b, v_{-1}^s be the highest valuations for buying and selling respectively of all unmatched traders and let $p = \frac{v_{-1}^b - v_{-1}^s}{2}$.
3. If $p \in [-v_0^s, v_0^b]$, the payment for each buyer is p and $-p$ for each seller,
4. Otherwise, remove the buyer (seller) with buying (selling) valuation v_0^b (v_0^s), breaking ties independently of their reports, and the payment for each remaining buyer is v_0^b and v_0^s for each remaining seller.

McAfee (1992) showed that $\mathcal{M}^{McAfee}$ is truthful, individually rational, and budget balanced in the classical exchange setting. Furthermore, the social welfare of $\mathcal{M}^{McAfee}$ approaches to the optimal in the large market limit, since it removes at most one pair of traders, who give the lowest social welfare increase, from the efficient allocation. However, we show in Theorem 1 that $\mathcal{M}^{McAfee}$ loses the truthfulness property in our general setting.

Theorem 1. $\mathcal{M}^{McAfee}$ is not truthful.

Proof. The proof is by example. The following example shows that $\mathcal{M}^{McAfee}$ is not truthful for a buyer that is able to benefit from switching to being a seller by misreporting.

	θ	π^{VCG}	π^{McAfee}	x^{VCG}	x^{McAfee}
1:	(9, 0)	<i>s</i>	<i>s</i>	-5.5	-4
2:	(7, -2.5)	<i>b</i>	<i>b</i>	4	6
3:	(6, -6)	<i>b</i>	<i>um</i>	4	0
4:	(2, -4)	<i>s</i>	<i>um</i>	-5	0
5:	(2, -5)	<i>um</i>	<i>um</i>	0	0

The above table shows the outcome of the VCG and $\mathcal{M}^{McAfee}$ for 5 traders, where *b* stands for *buyer*, *s* stands for *seller* and *um* stands for *unmatched*. Traders 3 and 4 are removed from the efficient allocation by $\mathcal{M}^{McAfee}$, and the payment for trader 1 is -4 and 6 for trader 2.

If trader 2 reported a valuation (7, > -2), then she will be allocated as a seller in $\mathcal{M}^{McAfee}$ (see the following table) and receive 4 (with utility 1.5), which is better than being allocated as a buyer with payment 6 (with utility 1).

	θ	π^{VCG}	π^{McAfee}	x^{VCG}	x^{McAfee}
1:	(9, 0)	<i>b</i>	<i>b</i>	-	6
2:	(7, > -2)	<i>s</i>	<i>s</i>	-	-4
3:	(6, -6)	<i>b</i>	<i>um</i>	-	0
4:	(2, -4)	<i>s</i>	<i>um</i>	-	0
5:	(2, -5)	<i>um</i>	<i>um</i>	-	0

Note that, trader 2 can also report (< 6.5, -2.5) to be allocated as a seller. $\square$

As McAfee (1992) proved, $\mathcal{M}^{McAfee}$ is truthful if no trader offers to buy and to sell at the same time. Therefore, the incentive problem of $\mathcal{M}^{McAfee}$ is caused by the traders who offer both buying and selling as we have seen in the proof of Theorem 1. However, Theorem 2 shows that traders are not incentivized to cheat in $\mathcal{M}^{McAfee}$, if their valuations are *marginally decreasing*, i.e. for all trader i , $v_i^b < -v_i^s$, although they might offer both to buy and to sell.

$-v_{b_m}^s >$	$v_{b_m}^b$ $\vee$	$-v_{s_m}^s$ $\wedge$	$> v_{s_m}^b$
$\vdots$	$\vdots$	$\vdots$	
$-v_{b_i}^s >$	$v_{b_i}^b$ $\vee$	$-v_{s_i}^s$ $\wedge$	$> v_{s_i}^b$
$\vdots$	$\vdots$	$\vdots$	
$-v_{b_0}^s >$	$v_{b_0}^b (v_0^b)$ $\vee$	$-v_{s_0}^s (-v_0^s)$ $\wedge$	$> v_{s_0}^b$
	$v_{b_{-1}}^b$	$\leq -v_{s_{-1}}^s$	

Figure 1: Efficient allocation example with marginally decreasing valuations for the proof of Theorem 2.

Theorem 2. $\mathcal{M}^{McAfee}$ is truthful if for all trader i , $v_i^b < -v_i^s$, i.e., i 's valuation is marginally decreasing.

Proof. Given the efficient allocation, let v_0^b, v_0^s be the lowest buying and selling valuations of all buyers and sellers respectively (see Figure 1 for an example). We first show that a buyer (seller) in the efficient allocation cannot misreport to be allocated as a seller (buyer) with positive utility in $\mathcal{M}^{McAfee}$.

Assume a buyer b_i , in Figure 1, misreported to be allocated as a seller, then the new efficient allocation will add more buyers or remove some sellers. Before b_i switches, we have that all unmatched traders' buying valuations are $\leq v_0^b$ because of the efficient allocation, and all sellers' buying valuations are $< -v_0^s \leq v_0^b$ because of their marginally decreasing valuations and efficiency. If after b_i switches, new buyers are added, then b_i 's payment as a seller is bounded by the buyer with the lowest buying valuation, which is $> -v_{b_0}^b > v_{b_i}^s$. If after b_i switches, some sellers are removed, then the payment for b_i as a seller is bounded by the highest selling valuation of unmatched traders, which is $\geq v_0^s > v_{b_i}^s$. Thus the utility for b_i being a seller is negative. Similarly, we can prove that a seller in the efficient allocation cannot misreport to be allocated as a buyer with positive utility in $\mathcal{M}^{McAfee}$.

Since no buyer/seller can gain anything by switching, the only way that a buyer/seller can manipulate its report to benefit in $\mathcal{M}^{McAfee}$ is to lower her buying/selling valuation to reduce her payment. However, McAfee (1992) already proved that a buyer/seller cannot gain anything by decreasing her buying/selling valuation. Therefore, for an unmatched trader, it is evident that if she misreported to be allocated as a buyer/seller, she will pay more than her valuation. $\square$

The Balanced Trade Reduction

As discussed in the last section, McAfee's trade reduction might not be truthful unless all traders' valuation is marginally decreasing. Moreover, in real-world applications, the buying price is not necessarily smaller than the selling price (i.e. not marginally decreasing). For instance, in ridesharing, a driver would sell an empty seat with a very low price if the costs of adding one more person in her car is very small, and she is also willing to pay as much as her driving costs, which can be very high, to buy a seat to ride. In this section, we propose a new trade reduction mechanism to handle the situation where not all traders' valuations are marginally decreasing. The new mechanism defines every trader's payments by adding same amount to the VCG payments such that buyers will pay more and sellers receive less compared to VCG. The intuition behind our mechanism is that VCG maximizes every trader's utility and if we increase a trader's VCG payments by the same amount independently of her valuation report, we will prevent her from switching between the buyer and seller sides (i.e. truthful) as well as reduce deficit.

Before we describe the new mechanism, let us first introduce the general payment setting used in the mechanism, which defines every trader's payment via increasing the VCG payments until they reach some thresholds.

Balanced payment setting $x(\underline{v}^b, \underline{v}^s)$

Given $\underline{v}^b$ and $\underline{v}^s$ the two thresholds for the payment increase ($\underline{v}^b$ for buying and $\underline{v}^s$ for selling),

1. For trader i , let x_i^{vcg} and $\hat{x}_i^{vcg}$ be the VCG payments for i being a buyer and a seller respectively in VCG, which can be calculated by assuming i 's report is $(\infty, -\infty)$ and $(0, 0)$ respectively.
2. The payment for i being a buyer is

$$x_i^b = \begin{cases} x_i^{vcg} & \text{if } x_i^{vcg} \geq \underline{v}^b \\ \min(\underline{v}^b, x_i^{vcg} + (\underline{v}^s - \hat{x}_i^{vcg})) & \text{otherwise.} \end{cases}$$

3. The payment for i being a seller is

$$x_i^s = \begin{cases} \hat{x}_i^{vcg} & \text{if } x_i^{vcg} \geq \underline{v}^b \\ \min(\underline{v}^s, \hat{x}_i^{vcg} + (\underline{v}^b - x_i^{vcg})) & \text{otherwise.} \end{cases}$$

$x(\underline{v}^b, \underline{v}^s)$ does not increase the VCG payments if one of the

VCG payments is above the corresponding threshold, otherwise, increases the VCG payments by the same amount of $\min(\underline{v}^b - x_i^{vcg}, \underline{v}^s - \hat{x}_i^{vcg})$. That is, the increments in a trader's VCG payments for being a buyer and a seller are balanced and limited by the thresholds $\underline{v}^b, \underline{v}^s$, which is essential to prevent a buyer/seller from switching to the other side.

Given the above payment setting mechanism, we define the new trade reduction mechanism which applies $x(\underline{v}^b, \underline{v}^s)$ by defining the two thresholds for each trader and then uses the payment for the allocation.

Balanced trade reduction $\mathcal{M}^{btr}$

1. Given the traders' valuation report profile θ , compute the efficient allocation $\pi^e(\theta)$, and let b_0, s_0 be the buyer and the seller with the lowest buying and selling valuations of all (matched) buyers and sellers respectively (breaking ties independently of their reports).
2. For each buyer $i \neq b_0$ in $\pi^e(\theta)$, i is allocated as a buyer and the payment is defined by $x(\underline{v}^b, \underline{v}^s)$ with the parameters $\underline{v}^b = v_{b_0}^b$ and $\underline{v}^s$ defined as follows:

$$\begin{cases} \min_{j \neq i, \pi_j^e(\theta_{-i}, \hat{\theta}_i) = seller} v_j^s & \text{if } |\{j : \pi_j^e(\theta_{-i}, \hat{\theta}_i) = seller\}| > 1, \\ \infty & \text{otherwise.} \end{cases} \quad (2)$$

where $\hat{\theta}_i$ is any report of i s.t. $\pi_i^e(\theta_{-i}, \hat{\theta}_i) = seller$.

3. For each seller $i \neq s_0$ in $\pi^e(\theta)$, i is allocated as a seller and the payment is defined by $x(\underline{v}^b, \underline{v}^s)$ with the parameters $\underline{v}^s = v_{s_0}^s$ and $\underline{v}^b$ defined as follows:

$$\begin{cases} \min_{j \neq i, \pi_j^e(\theta_{-i}, \hat{\theta}_i) = buyer} v_j^b & \text{if } |\{j : \pi_j^e(\theta_{-i}, \hat{\theta}_i) = buyer\}| > 1, \\ \infty & \text{otherwise.} \end{cases} \quad (3)$$

where $\hat{\theta}_i$ is any report of i s.t. $\pi_i^e(\theta_{-i}, \hat{\theta}_i) = buyer$.

4. For each $i \in \{b_0, s_0\}$, define i 's payments by $x(\underline{v}^b, \underline{v}^s)$ with the parameters

- $\underline{v}^b$ defined as Equation (3).
- $\underline{v}^s$ defined as Equation (2).

Then, b_0 is allocated as

$$\begin{cases} \text{buyer} & \text{if } v_{b_0}^b > x_{b_0}^b, \\ \text{unmatched} & \text{otherwise.} \end{cases}$$

and s_0 is allocated as

$$\begin{cases} \text{seller} & \text{if } v_{s_0}^s > x_{s_0}^s, \\ \text{unmatched} & \text{otherwise.} \end{cases}$$

5. All unmatched traders in $\pi^e(\theta)$ are unmatched with payment zero.

Let us first discuss the key difference between $\mathcal{M}^{McAfee}$ and $\mathcal{M}^{btr}$. For the sake of simplicity, $\mathcal{M}^{btr}$ does not consider the case where $\mathcal{M}^{McAfee}$ does not need to remove a buyer and a seller from the efficient allocation, i.e., when the




Figure 2: Price setting example for a buyer in $\mathcal{M}^{btr}$

price given by unmatched traders does not violate individual rationality, which can be easily included in $\mathcal{M}^{btr}$ without violating the following analysis.

From the definition of $\mathcal{M}^{btr}$, we have that given traders' report profile, a buyer (seller) allocated in $\mathcal{M}^{McAfee}$ is also allocated as a buyer (seller) in $\mathcal{M}^{btr}$. That is, the efficiency of $\mathcal{M}^{btr}$ is at least that of $\mathcal{M}^{McAfee}$, which approaches to the optimal in the large market limit (McAfee 1992).

Proposition 1. $\mathcal{M}^{btr}$ is as efficient as $\mathcal{M}^{McAfee}$.

For buyer i other than b_0 , i will pay $v_{b_0}^b$ in $\mathcal{M}^{McAfee}$, while in $\mathcal{M}^{btr}$, i 's payment might be $\leq v_{b_0}^b$ depending on i 's payment setting for being allocated as a seller if i misreported. For example, in Figure 2, x_i^{vcg} and $\hat{x}_i^{vcg}$ are the VCG payments for i being a buyer and a seller in VCG, and $\underline{v}^b$ and $\underline{v}^s$ are the thresholds defined in $\mathcal{M}^{btr}$ for i (it is evident that $\underline{v}^b = v_{b_0}^b$). If $\underline{v}^b - x_i^{vcg} > \underline{v}^s - \hat{x}_i^{vcg}$, then the payment for i being a buyer in $\mathcal{M}^{btr}$ is $x_i^b = x_i^{vcg} + (\underline{v}^s - \hat{x}_i^{vcg}) < \underline{v}^b$. Because of this balanced payment setting, $\mathcal{M}^{btr}$ guarantees that a buyer (seller) $i \notin \{b_0, s_0\}$ allocated in VCG is still better off for being a buyer (seller) in $\mathcal{M}^{btr}$. The threshold $\underline{v}^b$ ($\underline{v}^s$) is ∞ , if the trader is the only buyer (seller) allocated, otherwise is the lowest buying (selling) valuation of other buyers (sellers).

It is worth mentioning that for a buyer/seller i allocated in both $\mathcal{M}^{McAfee}$ and $\mathcal{M}^{btr}$, i 's payment in $\mathcal{M}^{btr}$ can be also bigger than her payment in $\mathcal{M}^{McAfee}$, if the VCG payment is already above the threshold, which does exist.

Unlike $\mathcal{M}^{McAfee}$, $\mathcal{M}^{btr}$ does not simply remove buyer b_0 and seller s_0 from the efficient allocation, because they might be able to switch/misreport to trade on the other side with a positive utility. See Figure 3 as an example, if we completely removed b_0 and s_0 in $\mathcal{M}^{btr}$, trader 3, i.e. b_0 , with true valuation $(10, -3)$ is incentivized to misreport to trade on the other side with a utility at least 1. In Theorem 3, we formally show that $\mathcal{M}^{btr}$ is truthful.

Theorem 3. $\mathcal{M}^{btr}$ is truthful.

Proof. Since $x(\underline{v}^b, \underline{v}^s)$ is based on the VCG payment, the payment for each buyer/seller in $\mathcal{M}^{btr}$ is at least the VCG payment. Therefore, for all traders who are unmatched in VCG, their utility is maximized when they are also unmatched in $\mathcal{M}^{btr}$.

	θ	π^{VCG}	$\hat{\pi}^{btr}$	x^{VCG}	x^{btr}
1:	(11, -3)	b	b	—	—
2:	(8, -2)	s	s	—	—
3:	(10, -3)	b	um	6	0
4:	(4, -4)	s	um	—	—
5:	(6, -7)	um	um	—	—
6:	(5, -6)	um	um	—	—

	θ	π^{VCG}	$\hat{\pi}^{btr}$	x^{VCG}	x^{btr}
1:	(11, -3)	b	b	—	—
2:	(8, -2)	s	s	—	—
3:	(0, -3)	s	s	-5	≤ -4
4:	(4, -4)	s	um	—	—
5:	(6, -7)	b	b	—	—
6:	(5, -6)	b	um	—	—

Figure 3: A buyer with the lowest buying valuation switches to the other side with positive utility.

For all buyer $i \neq b_0$, threshold $\underline{v}^b = v_{b_0}^b \leq v_i^b$ which cannot be changed by i , as no matter what i reports, $v_{b_0}^b$ is still the lowest buying valuation of other buyers, as soon as i is allocated as a buyer. Similarly, we can show that $\underline{v}^s$ is not determined by i 's report as soon as i is allocated as a seller. We know that the VCG payments $x_i^{vcg}, \hat{x}_i^{vcg}$ are independent of i 's report and $v_i^b - x_i^{vcg} \geq v_i^s - \hat{x}_i^{vcg}$. Thus, $x_i^b = x_i^{vcg} + \delta$ and $x_i^s = \hat{x}_i^{vcg} + \delta$ are independent of i 's report and $v_i^b - x_i^b \geq v_i^s - x_i^s$, where $\delta \geq 0$. We can further show that $v_i^b - x_i^b \geq 0$, no matter $x_i^{vcg} \geq \underline{v}^b$ or $x_i^{vcg} < \underline{v}^b$. Therefore, i 's utility is maximized for being allocated as a buyer in $\mathcal{M}^{btr}$.

For buyer b_0 , threshold $\underline{v}^b \geq v_{b_0}^b$ is the second lowest buying valuation of all buyers when i is not the only buyer allocated in the efficient allocation. Otherwise, $\underline{v}^b = \infty$. In both situations, b_0 cannot change $\underline{v}^b$ as soon as b_0 is allocated as a buyer. Also b_0 can not change $\underline{v}^s$. Since $\underline{v}^b \geq v_{b_0}^b$, payment $x_{b_0}^b$ might be greater than $v_{b_0}^b$. If $x_{b_0}^b > v_{b_0}^b$, then b_0 is unmatched because $0 > v_{b_0}^b - x_{b_0}^b \geq v_{b_0}^s - x_{b_0}^s$, otherwise b_0 is allocated as a buyer. Therefore, b_0 's utility is maximized in $\mathcal{M}^{btr}$.

Similarly, we can show that for all sellers in the efficient allocation, their utility is maximized in $\mathcal{M}^{btr}$. $\square$

Since $\mathcal{M}^{btr}$ does not always remove b_0, s_0 from the efficient allocation as $\mathcal{M}^{McAfee}$ does, it might end up with that the numbers of buyers and sellers are not balanced. See Figure 4 for the example from the proof of Theorem 1, if we apply $\mathcal{M}^{btr}$, trader 2's payment is reduced from 6 to 5 and trader 3 is reallocated as a buyer, which solves the incentive problem we had in $\mathcal{M}^{McAfee}$, but also allocates an additional buyer.

To handle the extra buyer/seller allocated by $\mathcal{M}^{btr}$, we cannot simply choose an unmatched trader to fix the unbalance, because if we use the same payment policy for the newly allocated trader as for other buyers/sellers, the trader will receive a negative utility (i.e., violates individual rationality), and if the payment is too large for her, then others might be incentivized to be unmatched (i.e., violates incentive compatibility). Therefore, one simple solution is to as-

	θ	π^{McAfee}	π^{btr}	x^{McAfee}	x^{btr}
1:	(9, 0)	<i>s</i>	<i>s</i>	-4	-4
2:	(7, -2.5)	<i>b</i>	<i>b</i>	6	5
3:	(6 , -6)	<i>um</i>	b	0	5
4:	(2, -4)	<i>um</i>	<i>um</i>	0	0
5:	(2, -5)	<i>um</i>	<i>um</i>	0	0

Figure 4: A running example for $\mathcal{M}^{McAfee}$ and $\mathcal{M}^{btr}$.

sume that the mechanism (market maker) has a *backup*, i.e., the mechanism is able to buy one unit and is also able to sell one unit. The backup is only used for fixing the unbalance when necessary. In real-world applications, this assumption is not hard to implement. For example, in ridesharing, the mechanism can use taxi companies as backups, and they are only called up if necessary. In electric vehicle charging, this can be easily implemented by adding some capability for storing the energy in the system.

Deficit Control and the Generalization

Unlike $\mathcal{M}^{McAfee}$, the payment increase in $\mathcal{M}^{btr}$ is limited by balancing traders' utilities for being a buyer and a seller. Although not every trader is capable of switching, we are not able to use their valuations to decide the payment increase. Therefore, even in the situation where no trader offers both buying and selling or all traders' valuations are marginally decreasing, the outcomes of applying $\mathcal{M}^{btr}$ and $\mathcal{M}^{McAfee}$ are not the same. In general, $\mathcal{M}^{btr}$ cannot guarantee budget balance. At the same time, we need to consider the costs of using the backup of the mechanism, e.g., the costs of using taxis in ridesharing, which can be modelled as the mechanism's valuation for buying and selling. For the example in Figure 4, if the mechanism's selling valuation is ≥ -6 , then there is no deficit for applying $\mathcal{M}^{btr}$.

In order to further reduce or completely remove deficit, we can further increase traders' payments by pushing up the thresholds of the payment increase, i.e. v^b and v^s in $\mathcal{M}^{btr}$. In this section, we propose a generalization of $\mathcal{M}^{btr}$, to further increase the payments by using the k -th lowest buying and selling valuations as the payment increase thresholds. We call this generalization *k-balanced trade reduction*, where $k \geq 1$ is chosen independently of traders' reports.

k-balanced trade reduction $\mathcal{M}^{btr,k}$

1. Given the traders' valuation report profile θ , compute the efficient allocation $\pi^e(\theta)$.
2. For each buyer/seller i in $\pi^e(\theta)$, define i 's payments by $x(\underline{v}^b, \underline{v}^s)$ with the parameters

- $\underline{v}^b$ defined as follows:

$$\begin{cases} \min_{j \neq i, \pi_j^e(\theta_{-i}, \hat{\theta}_i) = seller} v_j^s & \text{if } |\{j : \pi_j^e(\theta_{-i}, \hat{\theta}_i) = seller\}| > k, \\ \infty & \text{otherwise.} \end{cases}$$

where $\hat{\theta}_i$ is any report of i s.t. $\pi_i^e(\theta_{-i}, \hat{\theta}_i) = seller$, and $\min^k \{\dots\}$ denotes the k -th lowest.

- $\underline{v}^s$ defined as follows:

$$\begin{cases} \min_{j \neq i, \pi_j^e(\theta_{-i}, \hat{\theta}_i) = seller} v_j^s & \text{if } |\{j : \pi_j^e(\theta_{-i}, \hat{\theta}_i) = seller\}| > k, \\ \infty & \text{otherwise.} \end{cases}$$

where $\hat{\theta}_i$ is any report of i s.t. $\pi_i^e(\theta_{-i}, \hat{\theta}_i) = seller$.

Then, give i the allocation maximizing her utility.

Given Theorem 3, it is not hard to show that $\mathcal{M}^{btr,k}$ is truthful. Note that, to apply $\mathcal{M}^{btr,k}$, the system will need k backups for buying and selling. As discussed in the last section, backups (e.g., extra drivers) are usually available in real-world applications. Depending on the costs of using the backups, it is not always the case that the deficit generated by $\mathcal{M}^{btr,k_1}$ is less than the deficit generated by $\mathcal{M}^{btr,k_2}$ if $k_1 > k_2$. Moreover, it is easy to see the following proposition. In the extreme case, for a sufficient large k^* , we will have all traders are unmatched which is obviously budget balanced. It is challenging to pick a desirable k in general, because k should be independent of traders' reports. However, we may be able to provide some guidelines for selecting k if we know that traders' valuations satisfy certain distributions.

Proposition 2. *Given the traders' valuation report profile θ , we can always find a $k^* \geq 1$ such that the outcome of applying $\mathcal{M}^{btr,k^*}$ on θ is budget balanced.*

Lastly, in real-world applications, if we know some traders can only buy or sell, then we do not need to consider their ability of switching, which will further reduce the deficit. For instance, in ridesharing, for a commuter who does not have a car, then it is clear that she is not able to sell seats to others, i.e., we can increase her buying payment to the threshold.

Conclusion

We studied a type of exchanges, called dual-role exchanges, where a trader can offer to both buy and sell for one unit of a commodity, and when she transacts in the allocation, she can either buy or sell, but not both. Dual-role exchanges are very different from more commonly studied exchanges where a trader has to decide whether to buy or to sell before participating in the exchanges. In dual-role exchanges, the decision is made by the mechanism instead of the traders (traders just need to report their valuations for buying and selling). We showed that designing budget balanced mechanisms in such exchanges can be very challenging. Especially, we showed that McAfee's trade reduction mechanism, which performs very well in classical exchanges, is not truthful in general in dual-role exchanges. Therefore, we proposed a balanced trade reduction based on VCG and McAfee's trade reduction, which is truthful and as efficient as McAfee's trade reduction, but not budget balanced. To further reduce the deficit, we generalize the balanced trade reduction to trade off between efficiency and deficit.

We note that, when computing the efficient allocation in the reduction mechanisms, since a trader might be able to both buy and sell, we cannot directly apply the allocation

algorithm from McAfee’s original trade reduction which orders traders’ reported valuations with respect to buying or selling and matches them if buying prices are greater than selling prices. In this model, we can apply an efficient allocation based on maximum-weighted bipartite matching, which starts with an empty allocation/matching and adds one pair of unmatched traders giving the “maximum social welfare increase” into the matching iteratively until the social welfare cannot be increased. The allocation problem here is slightly different from the classical bipartite matching, because it does not have clear bipartite sets of buyers and sellers. Therefore, at this stage, it is not clear if we can apply the well-known Hungarian algorithms for the allocation, if so, then it can be solved in polynomial time.

It is also worth mentioning that the balanced trade reduction mechanisms are Groves mechanisms, where the payment is set by adding certain value to each trader’s VCG payment independently of the trader’s valuation report. Given the modified VCG payments, the mechanisms compute an allocation that maximises each trader’s utility (which sometimes also lead to the unbalanced numbers of buyers and sellers). Regarding the backups used in the proposed mechanisms, we have tried to substitute unmatched (or pre-selected) traders for backups, but these substitutions violate either truthfulness or individual rationality with the payments we defined. However, we do have truthful, individually rational and (weakly) budget-balanced mechanisms that do not require backups. For example, a mechanism first separates all traders into two disjoint sets independently of their reports: traders in one set can only be allocated as buyers or unmatched and traders in the other set can only be allocated as sellers or unmatched, and then applies, e.g., McAfee’s trade reduction mechanism. These kinds of mechanisms essentially shrink the allocation space via reducing the dual-role exchange to a single-role exchange and therefore are also very inefficient. We conjecture that there is some impossibility for achieving such a mechanism without backups, if some trader has the ability to switch between the seller and buyer sides, i.e., a trader is allocated as a buyer in one situation and as a seller in another situation with the same report.

To apply the mechanisms to broader applications, in future, we aim to investigate how our mechanisms can be applied to other domains. For example, in EV charging, every EV requires/supplies different amounts of electricity. In ridesharing, each trip consists of a set of route sections and each commuter wants to sell/buy different bundles of sections. It would be also very interesting to extend the model to dynamic settings, where traders come and leave at different points in time (Blum, Sandholm, and Zinkevich 2006; Parkes 2007).

Acknowledgments

The authors are especially grateful to David Parkes for his helpful comments and advices. We also acknowledge funding from the EPSRC-funded International Centre for Infrastructure Futures (ICIF) (EP/K012347/1).

References

- Babaioff, M., and Nisan, N. 2001. Concurrent auctions across the supply chain. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*, EC ’01, 1–10.
- Babaioff, M., and Walsh, W. E. 2005. Incentive-compatible, budget-balanced, yet highly efficient auctions for supply chain formation. *Decis. Support Syst.* 39(1):123–149.
- Babaioff, M.; Nisan, N.; and Pavlov, E. 2004. Mechanisms for a spatially distributed market. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, EC ’04, 9–20.
- Blum, A.; Sandholm, T.; and Zinkevich, M. 2006. Online algorithms for market clearing. *J. ACM* 53(5):845–879.
- Bredin, J.; Parkes, D. C.; and Duong, Q. 2007. Chain: A dynamic double auction framework for matching patient agents. *Journal of Artificial Intelligence Research* 30:133–179.
- Gonen, M.; Gonen, R.; and Pavlov, E. 2007. Generalized trade reduction mechanisms. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, EC ’07, 20–29.
- Han, S.; Han, S.; and Sezaki, K. 2010. Development of an optimal vehicle-to-grid aggregator for frequency regulation. *Smart Grid, IEEE Transactions on* 1(1):65–72.
- Kamar, E., and Horvitz, E. 2009. Collaboration and shared plans in the open world: studies of ridesharing. *IJCAI’09*, 187–194.
- Kempton, W., and Tomi, J. 2005. Vehicle-to-grid power fundamentals: Calculating capacity and net revenue. *Journal of Power Sources* 144(1):268 – 279.
- McAfee, R. P. 1992. A dominant strategy double auction. *Journal of Economic Theory* 56(2):434–450.
- Myerson, R. B., and Satterthwaite, M. A. 1983. Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29(2):265–281.
- Parkes, D. C.; Kalagnanam, J.; and Eso, M. 2001. Achieving budget-balance with vickrey-based payment schemes in exchanges. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence - Volume 2*, IJCAI ’01, 1161–1168.
- Parkes, D. C. 2007. Online mechanisms. In *Algorithmic Game Theory*. Cambridge, UK: Cambridge University Press.
- Parsons, S.; Rodriguez-Aguilar, J. A.; and Klein, M. 2011. Auctions and bidding: A guide for computer scientists. *ACM Comput. Surv.* 43(2):10:1–10:59.
- Wurman, P. R.; Walsh, W. E.; and Wellman, M. P. 1998. Flexible double auctions for electronic commerce: theory and implementation. *Decis. Support Syst.* 24:17–27.
- Zhao, D.; Zhang, D.; Gerding, E. H.; Sakurai, Y.; and Yokoo, M. 2014. Incentives in ridesharing with deficit control. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS ’14, 1021–1028.