Two-sided Markets: Mapping Social Welfare to Gain from Trade

Rica Gonen[1] and Ozi Egri[2]

1. Dept. of Management and Economics,
ricagonen@gmail.com
2. Dept. of Mathematics and Computer Science,
ozieg@hotmail.com
The Open University of Israel,
The Dorothy de Rothschild Campus,
1 University Road, Raanana 43107

Abstract. Though the definition of gain from trade extends the definition of social welfare from auctions to markets, from a mathematical point of view the additional dimension added by gain from trade makes it much more difficult to design a gain from trade maximizing mechanism. This paper provides a means of understanding when a market designer can choose the easier path of maximizing social welfare rather than maximizing gain from trade. We provide and prove the first formula to convert a social welfare approximation bound to a gain from trade approximation bound that maintains the original order of approximation. This makes it possible to compare algorithms that approximate gain from trade with those that approximate social welfare. We evaluate the performance of our formula by using it to convert known social welfare approximation solutions to gain from trade approximation solutions. The performance of all known two-sided markets solutions (that implement truthfulness, IR, BB, and approximate efficiency) are benchmarked by both their theoretical approximation bound and their performance in practice. Surprisingly, we found that some social welfare solutions achieve a better gain from trade than other solutions designed to approximate gain from trade.

1 Introduction

In recent years the established research on one-sided markets in economics and computer science was extended to two-sided markets where both buying and selling agents are strategic. A key difficulty that arises when moving from one-sided to two-sided markets is handling the additional intricacies that arise when approximating gain from trade (acronym GFT). The need to approximate gain from trade stems from Myerson and Satterthwaite's impossibility result [**?**]. [**?**] states that no two-sided market can simultaneously satisfy the economically desirable properties (of truthful reporting, participation without a loss and not running a deficit) while maintaining efficiency. The two-sided market literature largely chooses to approximating efficiency in order to maintain the other economic properties.

Given the intricacies involved in approximating gain from trade the literature often chooses to approximate social welfare (acronym SWF). Social welfare is the parallel of gain from trade in a one-sided market setting. In one-sided markets, where only the buying agents are strategic and there is a single selling agent, efficiency is measured by maximizing the sum of the buying agents' valuations, i.e., social welfare. Social welfare in one-sided markets extends to two-sided markets by summing the buying agents' valuations and subtracting the selling agents' costs, i.e., gain from trade [**?**,**?**]. Much of the literature on two-sided markets provides solutions that approximates social welfare efficiency maximization as opposed to approximating gain from trade efficiency, i.e., maximizing the sum of the buying agents' valuations plus the sum of the non sold commodities' costs held by selling agents at the end of trade [**?**,**?**,**?**,**?**,**?**,**?**,**?**,**?**].

One would expect that notionally gain from trade would be similar to social welfare in two-sided markets as gain from trade extends the definition of social welfare from auctions (one-sided markets) to two-sided markets [1]. Despite their conceptual similarity, it is much more complex to design a gain from trade maximizing approximation mechanism than a social welfare maximizing approximation mechanism. For example, if a buying agent has value \$10 for a commodity and a selling agent has cost \$7 and they trade. The social welfare is \$10, while the gain from trade is

[1] Illustratively, the broad explanation of maximizing social welfare in a two-sided market is that sellers who place a relatively higher cost on a given commodity should end up retaining that commodity while the broad explanation of maximizing gain from trade is that sellers who place a relatively lower cost on a given commodity should end up selling that commodity.

\$3. It is easy to see that any mechanism that maximizes social welfare also maximizes gain from trade. However, the two objectives are rather different in approximation. In the example above, if the buying agent and the selling agent do not trade, the mechanism achieves a social welfare of \$7 which is 70% of the optimal social welfare, however it achieves 0 gain from trade which is not within any constant factor of the optimal gain from trade. It can be observed that any constant factor approximation of a mechanism's gain from trade is necessarily a constant factor approximation of the mechanism's social welfare, however the other direction does not hold. Thus, gain from trade is a more difficult objective to approximate. Even so, gain from trade is an important market concept that accurately captures the value of the market to both sides; buyers and sellers (see [**?**,**?**] as an example). This paper provides a means of understanding when a designer can take the easier path of designing a market that approximates the maximization of social welfare instead of gain from trade.

Two-sided market research is motivated by numerous applications such as web advertising and securities trading, and indeed the literature contains multiple two-sided market designs. Some solutions present two-sided markets with a single commodity and unit-demand [**?**,**?**,**?**,**?**,**?**,**?**,**?**,**?**,**?**] while others are combinatorial markets with multiple commodities and demand for bundles [**?**,**?**,**?**,**?**]. Some of the two-sided markets are offline, i.e., optimize given all agents' bids in advance [**?**,**?**,**?**,**?**,**?**,**?**,**?**,**?**,**?**] while others are online, i.e., optimize as agents' bids arrive [**?**,**?**,**?**,**?**]. Lastly both deterministic [**?**,**?**,**?**,**?**] and randomized [**?**,**?**,**?**,**?**,**?**,**?**,**?**,**?**,**?**,**?**] solutions exist. The above literature seeks to maintain the desirable economic properties of truthfulness (agents dominant strategy is to report their true valuation/cost), IR (No agent should end up with a negative utility if the agent?s true valuation/cost is submitted to the mechanism.) and BB (The price paid by the buying agents is at least as high as the price received by the selling agents, i.e., the market does not run a deficit), while keeping as much as possible of the trade efficiency. However, the existing theoretical tools do not allow a designer to compare the efficiency of all the available solutions as some approximate social welfare [**?**,**?**,**?**,**?**,**?**] while others approximate gain from trade [**?**,**?**,**?**,**?**,**?**,**?**].

We provide and prove the first formula to convert a social welfare approximation bound to a gain from trade approximation bound that maintains the original order of approximation (under natural conditions). This makes it possible to compare solutions that approximate gain from trade with those that approximate social welfare. The conversion formula applies to the most general setting of two sided markets, i.e. each agent can buy or sell bundles of multiple distinct commodities which may have a different number of identical units from each. This is the most general combinatorial market setting and the conversion bound does not require any restrictions on the valuation functions. Indeed the formula can convert the bounds of single unit-single demand markets as well as combinatorial markets (see section 4). Moreover the formula does not change the mechanism's allocation nor the computed prices and therefore the economic properties remain the same.

We evaluate the performance of our formula by using it to convert social welfare approximating solutions in the literature to gain from trade approximating solutions. We compare the performance of all known two-sided market solutions (that implement truthfulness, IR, BB, and approximate efficiency) according to the theoretical approximation bound as well as in practice. With respect to the comparisons of theoretical bounds, we show that the converted bounds perform well even when our conditions are not met. More specifically, the converted bounds are guaranteed to maintain the competitiveness order of the original social welfare bound (roughly speaking) only when more commodities switched hands in market than not. However, the converted bounds perform well even when most commodities did not change hands and were not sold.

We also implement and run the various algorithms in practice using synthetic data to evaluate their relative performance at maximizing gain from trade. These results are compared to the converted theoretical bounds. Surprisingly, in the practical runs we found that some of the social welfare solutions achieve better gain from trade than solutions that were designed to approximated gain from trade. This even happens in cases where the social welfare solution was intended for a combinatorial market settings and the gain from trade solution was intended for single commodity and unit-demand settings.

Another interesting aspect of our conversion formula is that it can be used to indicate, without an actual practical run, the practical performance of a social welfare maximizing two-sided market compared to a gain from trade maximizing two-sided market. By converting the social welfare bound of a two-sided social welfare maximizing algorithm to a gain from trade maximizing bound and comparing it with another gain from trade maximizing two-sided market algorithm we found that one can estimate the practical performance of the two algorithms with respect to gain from trade maximization.

In summary, this paper's contributions are threefold. First, we provide the first means of comparing the performance of previously uncomparable solutions. Second, the experimental tests show that our conversion formula can be used to predict the practical performance difference between the compared mechanisms. Third, we show that for combinatorial two-sided markets it is better to use the known social welfare maximizing solutions than to use a solution that directly maximizes gain from trade.

2 Preliminaries

There are multiple commodities and each one comes in a number of units. Let A be the total number of units of all commodities in a market. There are l agents interested in selling commodities. These agents may also be interested in selling multiple units of each of their commodities. There are n agents who are interested in buying commodities. These agents may also be interested in buying multiple units of these commodities. An allocation in a two-sided market can be represented as a pair of vectors $(X,Y) = ((X_1,...,X_n),(Y_1,...,Y_l))$ such that sum of elements in the union of $X_1,...,X_n,Y_1,...Y_l$ is A, and $X_1,...X_n,Y_1,...,Y_l$ are mutually non-intersecting. Each buying agent i, $1 \leq i \leq n$ has a valuation function v_i that assigns a non-negative value to each allocation X_i. Each selling agent t, $1 \leq t \leq l$ has a bundle of commodities S_t that he initially owns and a cost function c_t that assigns a positive cost for each allocation Y_t. The auctioneer's goal in the one-sided auction is to partition the commodities by allocating each buying agent i, X_i, so as to maximize $\sum_{i=1}^{n} v_i(X_i)$. This goal is referred to as maximizing social welfare (SWF) (or efficiency).

In a two-sided market the market maker's goal is to change hands and partition the commodities by allocating each buying agent i, X_i and each selling agent t, Y_t, so as to maximize $\sum_{i=1}^{n} v_i(X_i) - \sum_{i=1}^{l} c_t(Y_t)$. This goal is referred to as maximizing gain from trade (acronym GFT) (efficiency).

As discussed in section 1 much of the literature on two-sided markets provides solutions that approximates social welfare efficiency maximization as opposed to approximating gain from trade efficiency. In a two-sided market SWF means maximizing the sum of the buying agents' valuations plus the sum of the unsold commodities' costs. The motivation behind this extension is that SWF accounts for all agents that end up with commodities at the end of the trade.

Let (X^o,Y^o) be the pair of vectors containing the optimal allocation in the two-sided market. Let (X,Y) be the pair of vectors containing the two-sided market algorithm's allocation solution. Let $V_{ALG} = \sum_{i=1}^{n} v_i(X_i)$ and let $V_{OPT} = \sum_{i=1}^{n} v_i(X^o{}_i)$ be the two-sided market algorithm's solution and the optimal SWF maximization solution computed only using the buying agents, and without accounting for the unsold commodities. Let $C_{ALG} = \sum_{t=1}^{l} c_t(Y_t)$ and let $C_{OPT} = \sum_{t=1}^{l} c_t(Y^o{}_t)$ be the two-sided market algorithm's solution and optimal's SWF minimization solution for the selling agents. Let $G_{ALG} = \sum_{t=1}^{l}(c_t(S_t) - c_t(Y_t))$ and let $G_{OPT} = \sum_{t=1}^{l}(c_t(S_t) - c_t(Y^o{}_t))$ be the two-sided market's solution and the SWF maximization solution computed using only the unsold commodities.

Let $W_{ALG} = V_{ALG} + G_{ALG}$ and let $W_{OPT} = V_{OPT} + G_{OPT}$.
Let $\gamma = \frac{V_{OPT}}{C_{OPT}}$, let $\delta = \frac{G_{OPT}}{C_{OPT}}$ and let $\mu \geq \frac{W_{OPT}}{W_{ALG}}$.

3 Converting Social Welfare to Gain from Trade

In this section we show how to convert SWF maximization approximation bound in two-sided markets into GFT maximization approximation ratio guarantee.

In the following theorem we assume non trivial market mechanisms, i.e. mechanisms where at least one trade occurs where the seller has a positive cost for that trade and the optimal GFT is strictly positive. That is $\gamma > 1$, $C_{OPT} > 0$ and $\mu > 0$. Let $H = \left(\frac{2\gamma + (-\mu+2)\delta - \mu - \frac{\mu W_{ALG}}{C_{OPT}}}{\mu(\gamma-1)} \right)$.

Theorem 1. *Any two-sided market mechanism, such that $\gamma > 1, C_{OPT} > 0$, that maximizes SWF[2] within a factor of $\mu > 1$, i.e., $W_{ALG} \geq \frac{1}{\mu} W_{OPT}$ is H-competitive with respect to the optimal GFT, i.e., $(V_{ALG} - C_{ALG}) \geq \frac{1}{H}(V_{OPT} - C_{OPT})$.*

[2] maximizes SWF of buying agents and remaining commodities.

Moreover if $\delta \leq 1$ and $W_{ALG} > G_{OPT} + C_{OPT}$[3] then $0 < H \leq 1$ and $\frac{1}{H}$ approximation factor maintains the original $\frac{1}{\mu}$ order of approximation.

Hence, for $W_{ALG} = \frac{1}{\mu} W_{OPT}$ the competitive ratio is $V_{ALG} - C_{ALG} \geq \left(\frac{\gamma+\delta}{\mu(\gamma-1)} - \frac{\delta+1}{\gamma-1} \right) (V_{OPT} - C_{OPT})$.

Intuitively the formula's outcome GFT bound maintains the μ order of approximation only in settings where the optimal solution SWF from sold commodities is at least as high as SWF from unsold commodities, i.e. δ less or equal 1. The performance condition is easier to understand when one considers a market where most trade can not occur and most commodities are left unsold. In such a market the cost of the unsold commodities will contribute to the SWF sum while the GFT will be unboundedly low including only the value of the few sold commodities in the few trades that will occur. Furthermore the larger the ratio of buyers' optimal SWF to sellers' optimal SWF, i.e., γ, is with respect to the converted SWF bound, the closer the converted approximation is to μ. Similarly to the above intuition, a market with a higher γ has a high GFT "potential" as there are high values of sold commodities compared to their costs.

Proof. The proof of Theorem 1 is composed of Lemma 1 and Lemma 2.

Lemma 1 that shows that the two-sided market mechanism is $\left(\frac{\gamma+\delta-\frac{\mu G_{ALG}}{C_{OPT}}}{\mu\gamma} \right)$-competitive with respect to the buying agents' optimal SWF, i.e., $V_{ALG} \geq \left(\frac{\gamma+\delta-\frac{\mu G_{ALG}}{C_{OPT}}}{\mu\gamma} \right) V_{OPT}$.

Lemma 2 that shows that the two-sided market mechanism is $\left(1+\delta(1-\frac{1}{\mu}) - \frac{\gamma}{\mu} + \frac{V_{ALG}}{C_{OPT}}\right)$-competitive with respect to the selling agents' optimal SWF, i.e., $C_{ALG} \leq \left(1+\delta(1-\frac{1}{\mu}) - \frac{\gamma}{\mu} + \frac{V_{ALG}}{C_{OPT}}\right) C_{OPT}$.

For simplicity of exposition, let $\alpha = \frac{\mu\gamma}{\gamma+\delta-\frac{\mu G_{ALG}}{C_{OPT}}}$ and let $\beta = 1+\delta(1-\frac{1}{\mu}) - \frac{\gamma}{\mu} + \frac{V_{ALG}}{C_{OPT}}$.

Combining the two Lemmas we have that

$$V_{ALG} - C_{ALG} \geq \frac{1}{\alpha} V_{OPT} - \beta C_{OPT} = \left[\frac{\gamma}{\alpha} - \beta \right] C_{OPT} \tag{1}$$

$$= \left[\frac{\frac{\gamma}{\alpha} - \beta}{\gamma - 1} \right] (V_{OPT} - C_{OPT}) \tag{2}$$

$$= \left(\frac{\gamma+\delta-\frac{\mu G_{ALG}}{C_{OPT}} - \mu - \delta(\mu-1) + \gamma - \frac{\mu V_{ALG}}{C_{OPT}}}{\mu(\gamma-1)} \right) (V_{OPT} - C_{OPT}) \tag{3}$$

$$= \left(\frac{2\gamma + (-\mu+2)\delta - \mu - \frac{\mu W_{ALG}}{C_{OPT}}}{\mu(\gamma-1)} \right) (V_{OPT} - C_{OPT})$$

Equalities (1) and (2) follow since $\gamma = \frac{V_{OPT}}{C_{OPT}}$. By substituting α and β in equality (2) we achieve equality (3).

It remains to show that the competitive ratio claimed in Theorem 1 is greater than zero and less or equal to one. Since $\mu > 0$ and $\gamma > 1$, in order for the competitive ratio to be greater than zero we assume that

$$2\gamma + (-\mu+2)\delta - \mu - \frac{\mu W_{ALG}}{C_{OPT}} > 0 \Rightarrow$$

$$2\frac{V_{OPT}}{C_{OPT}} - \mu\frac{G_{OPT}}{C_{OPT}} + \frac{2G_{OPT}}{C_{OPT}} - \mu - \frac{\mu W_{ALG}}{C_{OPT}} > 0 \Rightarrow$$

$$\frac{2V_{OPT} - \mu G_{OPT} + 2G_{OPT} - \mu C_{OPT} - \mu W_{ALG}}{C_{OPT}} > 0$$

[3] Note that the requirement for $W_{ALG} > G_{OPT} + C_{OPT}$ is trivial in the context of two-sided markets where the SWF resulting from unallocated commodities is included as the algorithm can at least gain the SWF resulting from not allocating any commodities.

Since $C_{OPT} > 0$ we only need to assume that

$$2V_{OPT} - \mu G_{OPT} + 2G_{OPT} - \mu C_{OPT} - \mu W_{ALG} > 0 \Rightarrow$$

$$2(V_{OPT} + G_{OPT}) > \mu(G_{OPT} + C_{OPT} + W_{ALG}) \Rightarrow$$

$$2W_{OPT} > \mu(G_{OPT} + C_{OPT}) + W_{OPT} \Rightarrow$$

$$W_{OPT} > \mu(G_{OPT} + C_{OPT}) \Rightarrow W_{ALG} > G_{OPT} + C_{OPT}$$

Note that the requirement for $W_{ALG} > G_{OPT} + C_{OPT}$ is trivial in the context of two-sided markets where the social welfare resulting from unallocated commodities is included as the algorithm can at least gain the social welfare resulting from not allocating any commodities.

For the competitive ratio to be less or equal 1 we need to assume that

$$2\gamma + (-\mu + 2)\delta - \mu - \frac{\mu W_{ALG}}{C_{OPT}} \leq \mu(\gamma - 1) \Rightarrow$$

$$(-\mu + 2)\delta - \mu - \frac{\mu W_{ALG}}{C_{OPT}} \leq \mu(\gamma - 1) - 2\gamma \Rightarrow$$

$$(-\mu + 2)\delta - \frac{\mu W_{ALG}}{C_{OPT}} \leq \mu\gamma - 2\gamma \Rightarrow$$

$$(2 - \mu)\delta - \frac{\mu W_{ALG}}{C_{OPT}} \leq \gamma(\mu - 2)$$

If $\mu \geq 2$ any positive γ will satisfy the above condition. Since we assume non trivial market where $\gamma > 1$ then in this case no additional assumption is needed. If $\mu < 2$ then it has to hold that

$$(2 - \mu)(\delta + \gamma) < \frac{\mu W_{ALG}}{C_{OPT}} = \gamma + \delta \Rightarrow$$

$$2 - \mu < 1 \Rightarrow \mu > 1$$

Note that the above requirement for $\mu > 1$ is natural since $\mu \geq \frac{W_{OPT}}{W_{ALG}}$ and is an approximation factor of a combinatorial problem.

In the case where $W_{ALG} = \frac{1}{\mu} W_{OPT}$ the expression $\mu \frac{W_{ALG}}{C_{OPT}}$ can be simplify to $\gamma + \delta$ and therefore

$$= \frac{2\gamma + (-\mu + 2)\delta - \mu - \gamma - \delta}{\mu(\gamma - 1)} (V_{OPT} - C_{OPT})$$

$$= \frac{\gamma + \delta(-\mu + 1) - \mu}{\mu(\gamma - 1)} (V_{OPT} - C_{OPT})$$

$$= \left(\frac{\gamma + \delta}{\mu(\gamma - 1)} - \frac{\mu(\delta + 1)}{\mu(\gamma - 1)} \right) (V_{OPT} - C_{OPT})$$

$$= \left(\frac{\gamma + \delta}{\mu(\gamma - 1)} - \frac{\delta + 1}{\gamma - 1} \right) (V_{OPT} - C_{OPT})$$

Lemma 1. *The two-sided market mechanism is* $\frac{\gamma + \delta - \frac{\mu G_{ALG}}{C_{OPT}}}{\mu\gamma}$*-competitive with respect to the buying agents' optimal social welfare, i.e.,* $V_{ALG} \geq \frac{1}{\alpha} V_{OPT}$*.*

Proof. (Proof of Lemma 1) From μ definition we know that $\frac{W_{ALG}}{W_{OPT}} \geq \frac{1}{\mu}$ or in other words that $\frac{V_{ALG} + G_{ALG}}{V_{OPT} + G_{OPT}} \geq \frac{1}{\mu}$. Therefore $V_{ALG} \geq \frac{V_{OPT} + G_{OPT}}{\mu} - G_{ALG}$, dividing by V_{OPT} we get that $\frac{V_{ALG}}{V_{OPT}} = \frac{1}{\alpha} \geq \frac{V_{OPT} + G_{OPT}}{\mu V_{OPT}} - \frac{G_{ALG}}{V_{OPT}} = \frac{1}{\mu} + \frac{G_{OPT} - \mu G_{ALG}}{\mu V_{OPT}}$. By multiplying $\frac{G_{OPT} - \mu G_{ALG}}{\mu V_{OPT}}$ numerator and denominator by $\frac{1}{C_{OPT}}$ we get that $= \frac{1}{\mu} + \frac{\frac{G_{OPT}}{C_{OPT}} - \frac{\mu G_{ALG}}{C_{OPT}}}{\frac{\mu V_{OPT}}{C_{OPT}}} = \frac{1}{\mu} + \frac{\delta - \frac{\mu G_{ALG}}{C_{OPT}}}{\mu\gamma} = \frac{\gamma + \delta - \frac{\mu G_{ALG}}{C_{OPT}}}{\mu\gamma}$.

Lemma 2. *The two-sided market mechanism is* $1+\delta(1-\frac{1}{\mu})-\frac{\gamma}{\mu}+\frac{V_{ALG}}{C_{OPT}}$*-competitive with respect to the selling agents' optimal social welfare, i.e.,* $C_{ALG} \leq \beta C_{OPT}$.

Proof. (Proof of Lemma 2) From μ definition we know that $\frac{W_{ALG}}{W_{OPT}} \geq \frac{1}{\mu}$ or in other words that $\frac{V_{ALG}+G_{ALG}}{V_{OPT}+G_{OPT}} \geq \frac{1}{\mu}$. Therefore $G_{ALG} \geq \frac{V_{OPT}+G_{OPT}}{\mu} - V_{ALG}$, dividing by G_{OPT} we get that

$$\frac{G_{ALG}}{G_{OPT}} =$$

$$\tilde{\beta} \geq \frac{V_{OPT}}{\mu G_{OPT}} + \frac{1}{\mu} - \frac{V_{ALG}}{G_{OPT}} = \frac{1}{\mu} + \frac{V_{OPT}-\mu V_{ALG}}{\mu G_{OPT}} \tag{4}$$

Now we need to convert $\tilde{\beta}$ the approximation ratio of $\frac{G_{ALG}}{G_{OPT}}$ to β the approximation ratio of $\frac{C_{ALG}}{C_{OPT}}$. Or in other words we need to convert the social welfare approximation resulting from the unallocated selling agents' commodities to social welfare approximation resulting from allocated selling agents at the market.

We know from (4) that $1-\left(1-\frac{1}{\mu}-\left(\frac{V_{OPT}-\mu V_{ALG}}{\mu G_{OPT}}\right)\right) \leq \tilde{\beta} \leq 1+\varepsilon$ holds, then we can say for β that

$$1-\delta\varepsilon \leq \beta \leq 1+\delta\left(1-\frac{1}{\mu}-\left(\frac{V_{OPT}-\mu V_{ALG}}{\mu G_{OPT}}\right)\right) \tag{5}$$

It follows from (5) that $\beta \leq 1+\delta(1-\frac{1}{\mu})-\left(\frac{V_{OPT}-\mu V_{ALG}}{\mu C_{OPT}}\right) = 1+\delta(1-\frac{1}{\mu})-\frac{\gamma}{\mu}+\frac{V_{ALG}}{C_{OPT}}$. It is easy to see that in order to keep $\tilde{\beta}$'s approximation quality one needs to assume that $\delta \leq 1$.

4 Experimental Results

We used simulations to empirically study the performance of our conversion formula. We investigated the questions, if and when a mechanism designer can replace the use of gain-from-trade maximizing algorithm with a social-welfare maximizing algorithm given the convergence formula. Our simulations involved two types of empirical evaluations. The first evaluation type reflected by Figures 1, 3, 6 and 7 studies the converted theoretical gain from trade (acronym GFT) approximation bounds on the simulated data as a function of δ and $\frac{1}{K_{\max}}$, where $\frac{1}{K_{\max}}$ is the maximal demand/supply number of units of any commodity by any agent. The second evaluation type reflected by Figures 2, 4, 5 and 8 shows the actual GFT approximation achieved by the benchmarked algorithms' runs on the same simulated data, shown as a function of δ and $\frac{1}{K_{\max}}$. In order to compute the actual GFT approximation achieved by the benchmarked algorithms we implemented the algorithms in [**?**,**?**,**?**,**?**,**?**,**?**,**?**,**?**,**?**,**?**,**?**].

Inputs were generated based on various random distributions and found minimal to no qualitative difference between distributions. In the figures the uniform distribution was used in the following manner: Agents' costs and values bids were selected as uniformly random independent values between 1 and 10^5. The supply/demand of different commodities the agents hold/desire were also selected as uniformly random independent values between 1 and 5000. For each supplied/demanded commodity we selected the number of units as uniformly random independent values between 1 and 10^6. All parameters showing in the plots such as δ, $K_{\max}$ are histograms based on the instances generated for empirical evaluation.

The literature we compare makes different assumptions and valuation distribution limitations under which their theoretical bounds are guaranteed. In our experiments, for all figures, when comparing two algorithms we construct markets that maintain the distribution assumptions of both algorithms by choosing markets that fulfill the most restrictive assumptions needed by the compared algorithms. It is important to note that some algorithms are more restrictive than others in which case the worst case of one may fall beyond the restriction of the other. For every comparison presented in the paper we also performed a comparison based on the less restrictive algorithm of the two. However, the results were not significantly different and therefore those figures were not included.

The results presented in Figures 5, 7, 6, and 8 were averaged over 30,000 trials and reflect millions of sellers, buyers and units of each of 5000 commodities. The comparisons were performed on markets with $K_{\max} = 1600$. While we could compute the theoretical bounds of all the two-sided markets in the literature using the above magnitude we could not do so with [**?**]'s theoretical bound as it gives a negative (N/A) bound in these cases[4]. Therefore for the comparisons with [**?**]'s algorithm and bound, presented in Figures 2, 1, 4, and 3 we used a different setting. Figures 2, 1, 4, and 3 were averaged over 15,000 trails per each column due to their very large size. These markets support millions of sellers, buyers and units but can only run on 4 commodities. The upper bound for a single trade is 0.3, $\delta = 0.5$ and every agent supplies/demands a single unit.

We first empirically evaluate the performance of our formula for converting from a SWF approximation bound to a GFT approximation bound by applying the formula to the known two-sided market mechanisms that provide a bound for the SWF approximation maximization ([**?**,**?**,**?**,**?**,**?**]). We convert the above five bounds to a GFT approximation bound and compare the resulting bounds with the known two-sided market mechanisms that provide a direct bound on their GFT approximation maximization ([**?**,**?**,**?**,**?**]). The figures comparing the theoretical GFT bound with the formula converted theoretical SWF bound illustrates the formula guaranty, i.e., even when the formula is used in some worst-case scenarios it is better for a designer to consider SWF maximizing algorithm over a GFT one.

Conversion of the bounds is accomplished by simply computing the values of γ and δ for the simulated data, after which the SWF bound μ of the SWF maximizing algorithms ([**?**,**?**,**?**,**?**,**?**]) is plugged into the formula of Theorem 1 and compared with ([**?**,**?**,**?**,**?**]) algorithms' bound on GFT. We show that though the converted bounds are guaranteed to maintain the competitiveness order of the original SWF bounds only if $\delta \leq 1$, the converted bounds perform well even when δ is as large as 5 and most commodities did not change hands and were not sold (see Figure 6 and Figure 7). This result holds across all converted bounds whether they bound a single-unit single-commodity setting or a combinatorial market setting.

It is important to note that for conducting Figures 1, 3, 6 and 7 one does not need to compute γ directly from an algorithm's run. A bound on γ can be concluded without running an algorithm to compute GFT. This results from much of the current literature assuming a bound on the maximum valuation bid and minimum cost bid from which one can conclude a bound on V_{OPT} and C_{OPT}.

In addition to the theoretical bound comparison we empirically compare the various algorithms in practice using synthetic data to evaluate their relative performance at maximizing GFT. These results are compared to the converted theoretical bounds. Surprisingly, in the practical runs, we found that some of the SWF solutions achieve better GFT than other solutions designed to approximated GFT. More specifically we found that Colini-Baldeschi et al. 2017 [**?**] that originally approximates SWF achieves a higher GFT than McAfee 2008 [**?**], which approximates GFT though Colini-Baldeschi et al. 2017, is intended for combinatorial market settings and McAfee 2008 is intended for single commodity and unit-demand settings (see Figure 5). Another example is the work by Blum et al. [**?**] that originally approximates SWF in the single-commodity unit-demand setting and performs better at GFT approximation in practice than Segal-Halevi et al. [**?**] in settings where Segal-Halevi et al.'s algorithm runs instances with unit-demand (see Figure 2). The above observation is particularly interesting given the fact that Blum et al. is an online algorithm (i.e., computes the SWF optimization function on an ongoing input steam) while Segal-Halevi et al. is an offline algorithm (that computes the GFT optimization function given all input agents' bids in advance). Gonen & Egri [**?**] is similarly interesting in that it originally approximates SWF in an online combinatorial market environment and in practice performs better at approximating GFT than Segal-Halevi et al. [**?**]. This is achieved despite the fact that Segal-Halevi et al. is an offline algorithm (see Figure 4).

Another interesting observation is that the practical runs appear to show that for combinatorial markets one is better off using the known social-welfare maximizing solutions [**?**,**?**] than using the known GFT maximizing solution [**?**], even if the designer is interested in maximizing GFT. This observation was made from figures 8 and 4 comparing [**?**] to [**?**] and [**?**] which show the social-welfare maximizing solutions consistently outperforming the GFT solutions.

An interesting aspect of our conversion formula is that it can be used to estimate, without an actual practical run, the practical performance of a SWF maximizing two-sided market compared to a GFT maximizing two-sided market. By converting the SWF bound of a two-sided SWF maximizing algorithm to a GFT maximizing bound and comparing it with another GFT maximizing two-sided market algorithm we found that one can evaluate the practical performance of the two algorithms with respect to GFT maximization. For example see Figure 3 showing that in the worst case analysis running Gonen and Egri will result in almost the same performance as running Segal-Halevi et al. and indeed

[4] [**?**]'s theoretical bound is negative unless markets are very large as the bound is not tight and the algorithm only performs well on very large markets.

in practice (see Figure 4) Gonen and Egri performs slightly better. We found similar prediction in the case of Blum et al. with respect to Segal-Halevi et al. See Figures 1 and 2.

A detailed table summarizing the comparative practical performance of the known social-welfare maximizing two-sided markets against the known GFT maximizing two-sided market can be found in Figure 9.

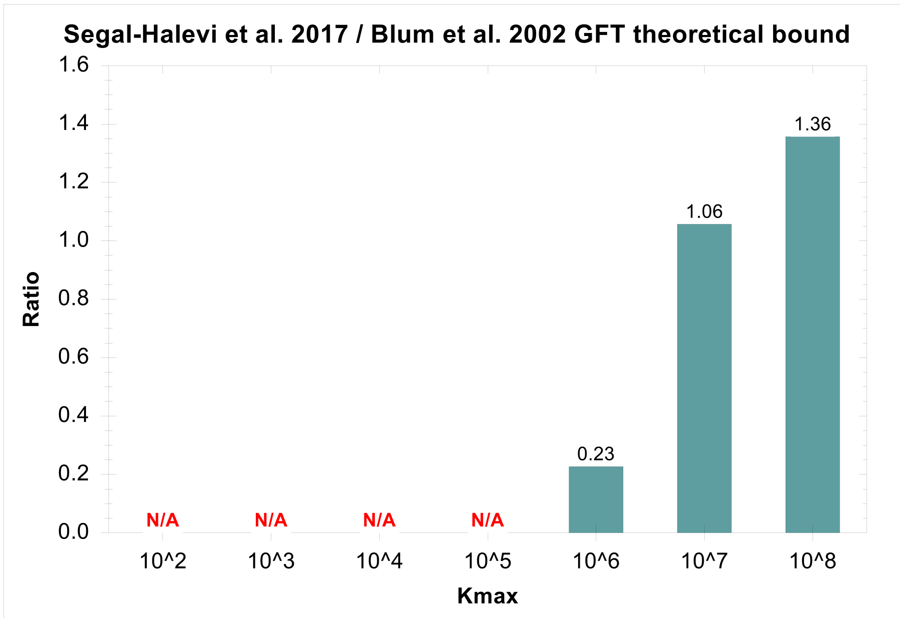

Fig. 1: Segal-Halevi et al. 2017's theoretical bound on gain from trade competitive ratio vs. Blum et al. 2002's converted theoretical bound on gain from trade competitive ratio. When under 10^6 *units are traded Segal-Halevi et al. 2017's theoretical bound results in a negative value. For very large markets where over* 10^8 *units are traded one might consider using Segal-Halevi et al. 2017's solution if no online aspect is required from the market.*

5 Conclusion and Discussion

In this paper we provided and proved the first formula to convert a SWF approximation bound for two-sided markets into a bound on a GFT approximation that maintains the original order of approximation. This conversion makes it possible to compare solutions that approximate GFT with those that approximate SWF. We evaluate the performance of our conversion formula by using it to convert SWF approximation solutions in the literature to GFT approximation solutions.

The experimental results showed that our conversion formula can be used to estimate (without an actual practical run) the practical performance (in GFT) of a social-welfare maximizing two-sided market compared to a (directly computed) GFT maximizing two-sided market.

We found that in some cases the solutions designed for SWF maximization perform better at maximizing GFT than algorithms designed for directly maximizing GFT. This is true in particular in the case of combinatorial markets where most known social-welfare maximizing solutions consistently perform better at maximizing GFT than the known combinatorial GFT (directly) optimizing solution.

We would like to see our conversion formula used in future work in multi-sided market design as a means for evaluation and comparison of new solutions to existing ones.

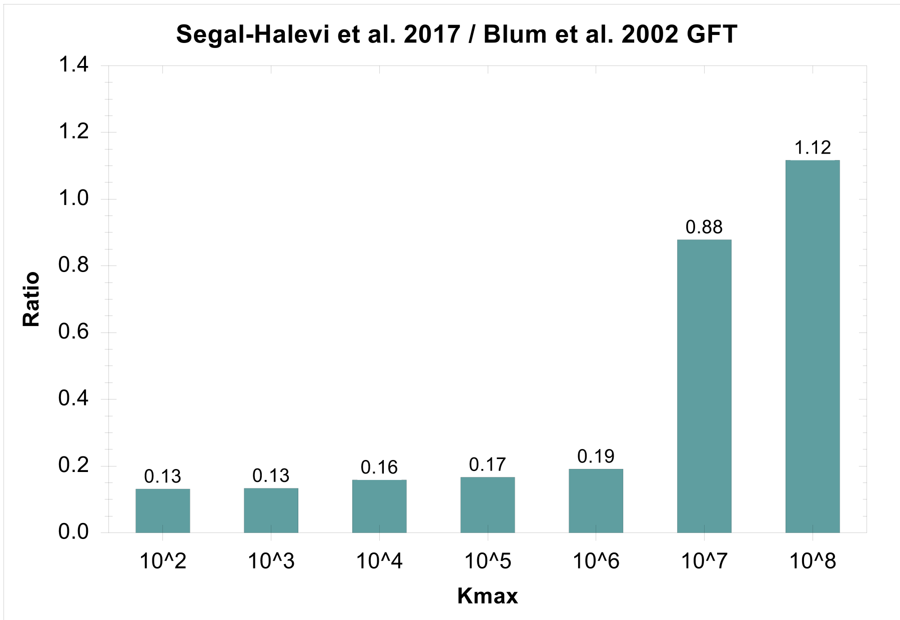

Fig. 2: Segal-Halevi et al. 2017's practical gain from trade competitive ratio vs. Blum et al. 2002's practical gain from trade competitive ratio. Similar to Figure 1 it can be seen that for very large markets where over 10^8 *units are traded one might consider using Segal-Halevi et al. 2017's solution if no online aspect is required from the market.*

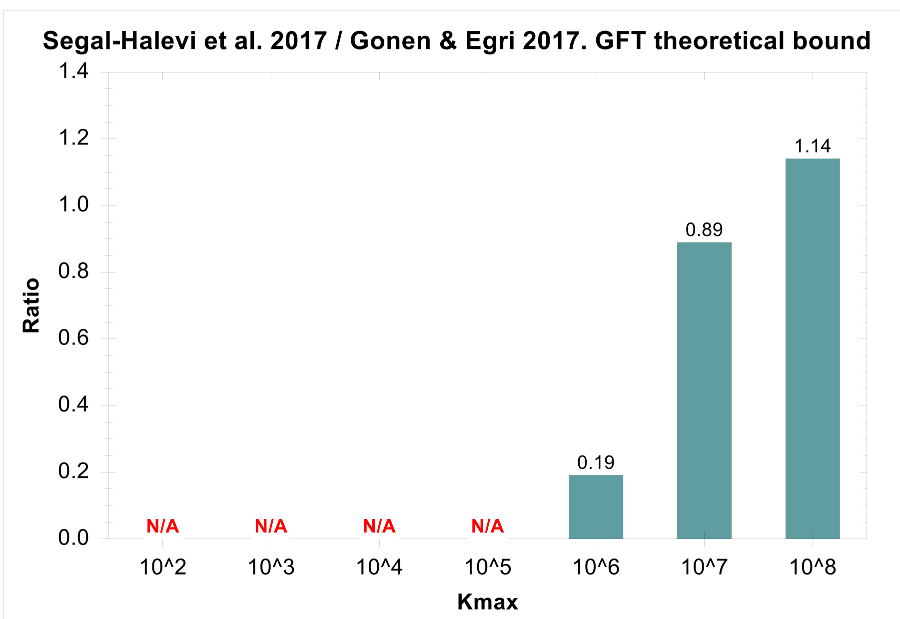

Fig. 3: Segal-Halevi et al. 2017's theoretical bound on GFT competitive ratio vs. Gonen & Egri 2017's converted theoretical bound on GFT competitive ratio. When under 10^6 *units are traded Segal-Halevi et al. 2017's theoretical bound results in a negative value. For very large markets, where over* 10^8 *units are traded, one might consider using Segal-Halevi et al. 2017's solution if no online aspect is required from the market.*

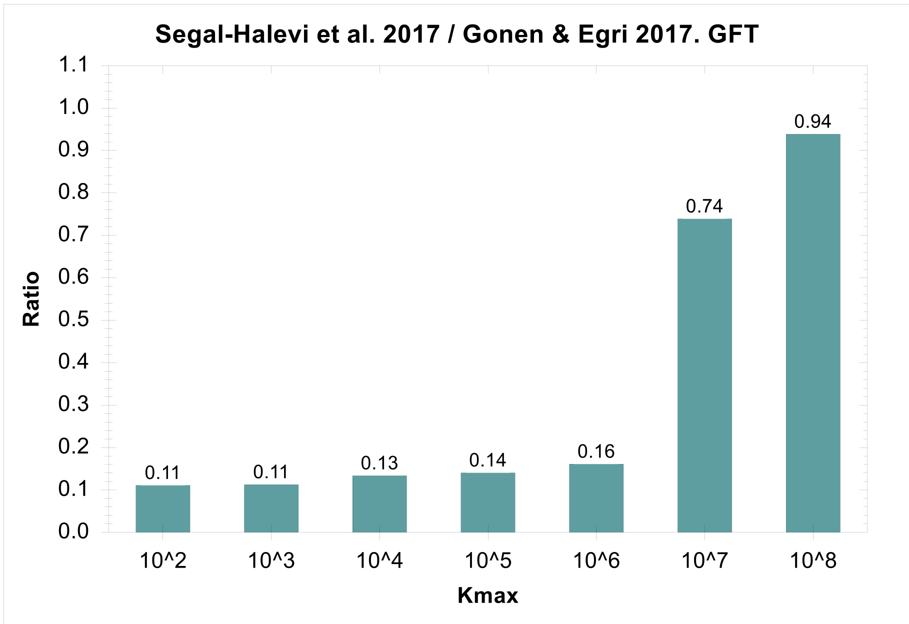

Fig. 4: Segal-Halevi et al. 2017's practical GFT competitive ratio vs. Gonen & Egri 2017's practical GFT competitive ratio. It seems that in practice Gonen & Egri 2017 performs better than Segal-Halevi et al. 2017 despite the fact that Gonen & Egri 2017 provides an online solution as opposed to Segal-Halevi et al. 2017's offline solution and Gonen & Egri 2017 are designed to maximize SWF as opposed to Segal-Halevi et al. 2017's which is designed to maximize GFT.

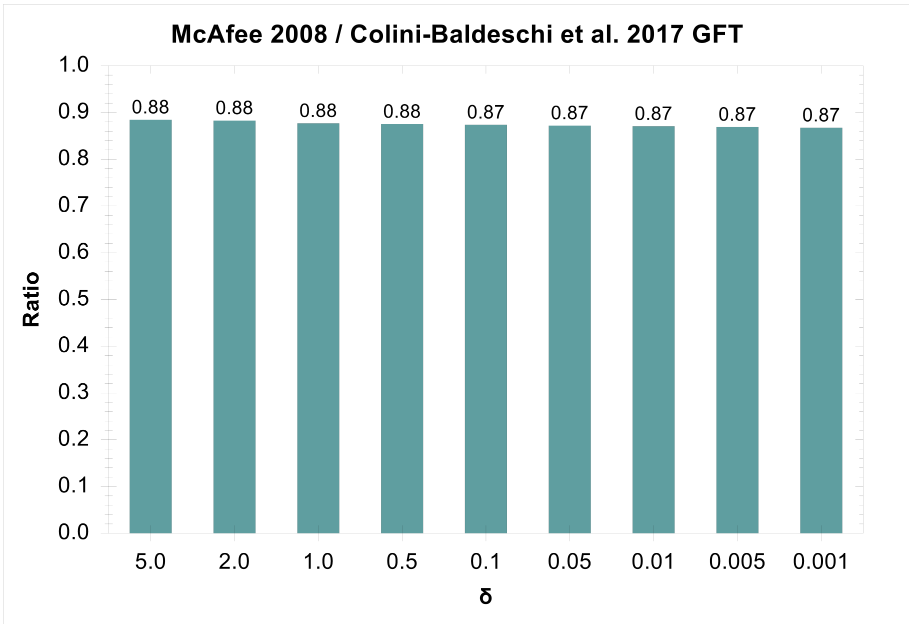

Fig. 5: McAfee 2008's practical GFT competitive ratio vs. Colini-Baldeschi et al. 2017's practical GFT competitive ratio. It appears that in practice Colini-Baldeschi et al. 2017 performs better than McAfee 2008 despite the fact that Colini-Baldeschi et al. 2017 provides a combinatorial solution as opposed to McAfee 2008's single commodity single-unit demand solution and Colini-Baldeschi et al. 2017 is designed to maximize SWF as opposed to McAfee 2008's which is designed to maximize GFT.

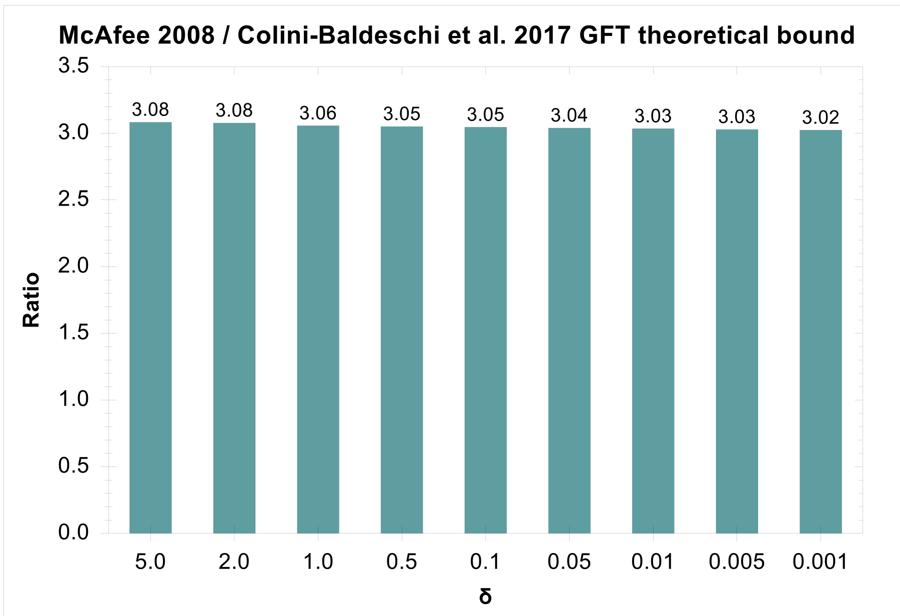

Fig. 6: McAfee 2008's theoretical bound on gain from trade competitive ratio vs. Colini-Baldeschi et al. 2017's converted theoretical bound on gain from trade competitive ratio. The theoretical bounds indicate that McAfee 2008 should perform better than Colini-Baldeschi et al. 2017. However the figure also indicates that the gap in performance might not be very large in particular if one wishes to design a combinatorial market (as Colini-Baldeschi et al. 2017) as opposed to a single-commodity unit-demand mechanism (as McAfee 2008 provides). Another interesting aspect of the figure is the effect of δ *on the converted bound. One can see that the relative difference in the converted theoretical bound between a market where most commodities are sold (*$\delta = 0.001$*) to a market where most commodities are not sold (*$\delta = 5$*) is less than 2%.*

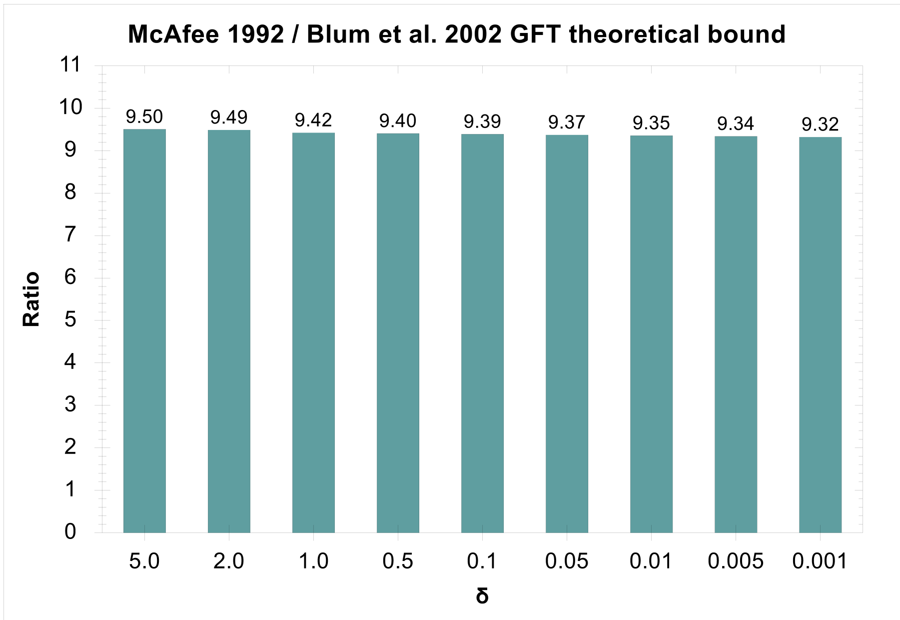

Fig. 7: McAfee 1992's theoretical bound on GFT competitive ratio vs. Blum et al. 2002's converted theoretical bound on GFT competitive ratio. Based on the theoretical bounds McAfee 1992 should perform better than Blum et al. 2002. However, the figure also indicates another interesting aspect which is the effect of δ on the converted bound. One can see that the relative difference in the converted theoretical bound between a market where most commodities are sold ($\delta = 0.001$) to a market where most commodities are not sold ($\delta = 5$) is less than 2%.

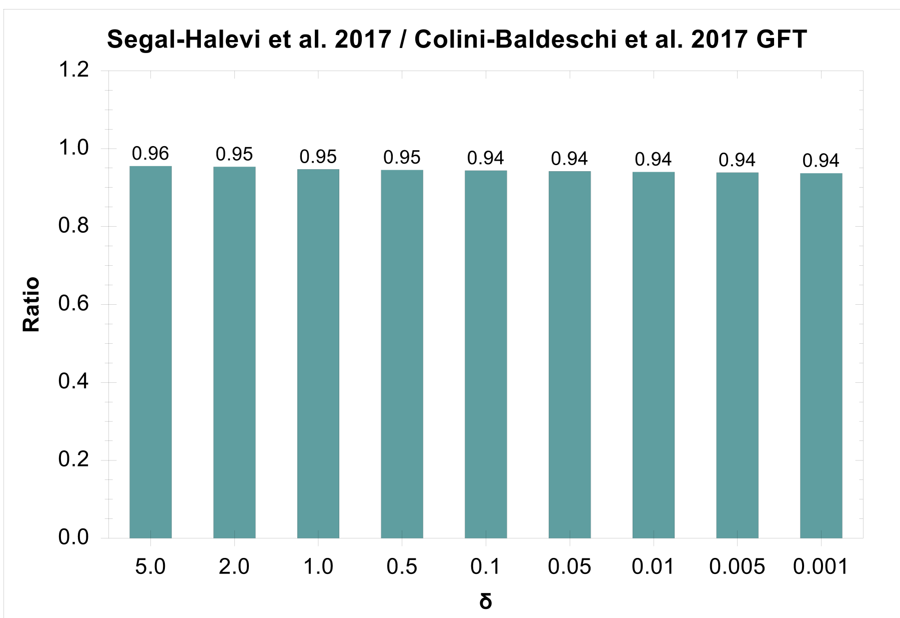

Fig. 8: Segal-Halevi et al. 2017's practical gain from trade competitive ratio vs. Colini-Baldeschi et al. 2017's practical gain from trade competitive ratio. In practice Colini-Baldeschi et al. 2017 performs better than Segal-Halevi et al. 2017 despite the fact that Colini-Baldeschi et al. 2017 is geared to maximize social welfare as opposed to Segal-Halevi et al. 2017's result which is geared to maximize gain from trade.

GFT / SWF	**Blumrosen & Dobzinski 2014** Off/Comb/Rand	**Colini-Baldeschi et al. 2016** Off/Sngl/Rand	**Berdin et al. 2007** On/Sngl/Rand	**Wurman et al. 1998** On/Sngl/Rand	**Gonen & Egri 2017** On/Comb/Detr
Segal-Halevi et al. 2017 Off/Comb/Rand	1.45*	0.14*	3.34	2.24	0.16*
Segal-Halevi et al. 2016 Off/Sngl/Detr	16.92	2.27	3.61	2.42	2.21
McAfee 2008 Off/Sngl/Rand	14.54	1.95	3.10	2.08	1.90
McAfee 1992 Off/Sngl/Detr	16.95	2.27	3.61	2.42	2.21

Off : offline , **On**: online
Comb: combinatorial market , **Sngl** : Single good market
Rand: algorithm is random , **Detr** : deterministic algorithm
* Were run in markets where Segal-Halevi et al. 2017 has a valid theoretical bound.

Fig. 9: This table demonstrates the practical performance of most known two-side market mechanisms with respect to GFT. The results suggest that it is better for a two-sided market designer who wishes to maximize GFT to use algorithms that directly maximize GFT. However, if the designer wishes to design a GFT maximizing two-sided combinatorial market than he/she should consider using a SWF maximizing algorithm as opposed to an algorithm that directly maximizes GFT.

References

1. Blum, A., Sandholm, T., Zinkevich, M.: Online algorithms for market clearing. In: SODA. pp. 971–980 (2002)
2. Blumrosen, L., Dobzinski, S.: Reallocation mechanisms. In: EC. pp. 617–640 (2014)
3. Bredin, J., Parkes, D., Duong, Q.: Chain: A dynamic double auction framework for matching patient agents. Journal of Artificial Intelligence Research **30**, 133–179 (2007)
4. BRUSTLE, J., CAI, Y., WU, F., ZHAO, M.: Approximating gains from trade in two-sided markets via simple mechanisms. In: In Proceedings of the 18th ACM Conference on Economics and Computation (EC). 589–590 (2017)
5. Colini-Baldeschi, R., Goldberg, P., de Keijzer, B., Leonardi, S., Roughgarden, T., Turchetta, S.: Approximately efficient two-sided combinatorial auctions. In: EC, 591-608 (2017)
6. Colini-Baldeschi, R., Goldberg, P., de Keijzer, B., Leonardi, S., Turchetta, S.: Fixed price approximability of the optimal gain from trade. In: Wine (2017)
7. Colini-Baldeschi, R., de Keijzer, B., Leonardi, S., Turchetta, S.: Approximately efficient double auctions with strong budget balance. In: SODA. pp. 1424–1443 (2016)
8. Gonen, R., Egri, O.: DYCOM: A dynamic truthful budget balanced double-sided combinatorial market. In: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017. pp. 1556–1558 (2017)
9. McAfee, R.P.: dominant strategy double auction. Journal of Economic Theory **56**, 434–450 (1992)
10. McAfee, R.P.: The gains from trade under fixed price mechanisms. Applied Economics Research Bulletin **1** (2008)
11. Myerson, R.B., Satterthwaite, M.A.: Efficient mechanisms for bilateral trading. Journal of Economic Theory **29**, 265–281 (1983)
12. Segal-Halevi, E., Hassidim, A., Aumann, Y.: SBBA: A strongly-budget-balanced double-auction mechanism. In: Algorithmic Game Theory - 9th International Symposium, SAGT 2016, Liverpool, UK, September 19-21, 2016. Proceedings. pp. 260–272 (September 2016)
13. Segal-Halevi, E., Hassidim, A., Aumann, Y.: Muda: A truthful multi-unit double-auction mechanism. In: Proceedings of AAAI (2018)
14. Wurman, P., Walsh, W., Wellman, M.: Flexible double auctions for electronic commerce: Theory and implementation. Decision Support Systems **24**, 17–27 (1998)