

Gains from Trade in Two-Sided Markets

Today we study a new setting, the two-sided market setting, and a new objective, gains from trade.

Definition 1. In *two-sided markets*, we have n buyers and m sellers, and a platform facilitating trade. Each seller owns some item(s) and has private values for them \mathbf{s}_j and will not sell below these values. Each buyer has private values for item(s) \mathbf{b}_i and will not buy above these values.

Let's simplify significantly for now and focus on the simplest possible setting: one buyer and one seller for one item. This setting is called *bilateral trade*.

Definition 2. In the *bilateral trade* setting, there is one seller with one item for sale, and item $s \sim F_S$ for their own item. There is also one buyer with item $b \sim F_B$ for the item. The platform's job is to determine a price for the buyer to pay, p^B , and a payment to the seller p^S .

We need to review our standard concepts in this new setting and make sure that we understand them, and see if anything additional is necessary.

Utility. If trade occurs at these prices, the buyer and seller respectively get utility

$$b - p^B \quad \text{and} \quad p^S - s.$$

If trade does not occur, the buyer and seller respectively get utility 0 each.

Budget Balance. Since the mechanism designer is the platform here, facilitating trade between the platform, it's natural to ask that the designer shouldn't have to put in any money to make the mechanism work (just like the designer doesn't have to pay any money in VCG). That is, we ask that the payment of the buyer covers the payment to the seller: $p^B \geq p^S$. We call this *weak budget balance* because it is only an inequality. Sometimes, mechanisms ask for *strong budget balance*—that the platform does not take any money off the table. Or, we relax these conditions to be *in expectation* over randomness of the mechanism or of the draw of values.

Many Single-Dimensional Buyers and Sellers. Now, we consider the setting with m identical sellers, each seller j with one item and one value $s_j \sim F_S$ for their item. There are n buyers, each with a value b_i for any item, where $b_i \sim F_B$. And we are in the i.i.d. case, so $F_B = F_S$.

Welfare. The welfare is just the sum of the values of those who hold the items. That is, if seller j did not trade their item, we count their value for the item, s_j , as they are holding the item. If they *did* trade their item, we instead count the value of the buyer they traded it to. We could write this a few different ways.

$$\sum_i b_i \cdot \mathbb{1}[i \text{ has item}] + \sum_j b_j \cdot \mathbb{1}[j \text{ has item}]$$

or, if M is the matching of buyers and sellers that trade, then

$$\sum_j s_j + \sum_{(i,j) \in M} b_i - s_j.$$

Observe that $\sum_j s_j$ is the *pre-trade welfare*—the welfare before trade occurs.

Gains from Trade. In contrast, the *Gains from Trade (GFT)* measures exactly the (aggregate) value that the platform adds from the actual trade. While the platform can still get $\sum_j s_j$ in welfare without making a single trade, the platform is not serving a purpose in this case. If M is again the matching of buyers and sellers that *actually trade*, then

$$\text{GFT} = \sum_{(i,j) \in M} b_i - s_j.$$

We put these notions together to observe that

$$\text{WELFARE} = \text{GFT} + \sum_j s_j$$

where the pre-trade welfare (the sum of the sellers' values) is *independent* of the mechanism. Then for any mechanism we design, we notice that approximating GFT will approximate welfare, but approximating welfare will not approximate GFT, so GFT is a harder objective.

OPT vs. Constrained-OPT. Our goal is to maximize GFT, and we would like the mechanism that does so to be

1. Dominant-Strategy Incentive-Compatible
2. Ex-Post Individually Rational
3. Weakly Budget-Balanced

In economics, they call the allocation that is the solution to the unconstrained optimization problem of maximizing GFT “first-best.” They call the mechanism that is the solution to the constrained optimization problem of maximizing GFT *subject to* (1-3) “second-best.”

Theorem 1 (Myerson Satterthwaite [3]). *Even for 1 buyer, 1 seller, and 1 item, the allocation that maximizes GFT (and thus welfare) may not be implementable by any mechanism satisfying (1-3). That is, first-best is not always attainable.*

The Optimal (First-Best) Allocation. Sort buyer values in decreasing order, seller values in decreasing order, and match them off while buyer values exceed seller values. The sum is independent of specific matchings, and adding any other pairs here would add something negative to the sum, so we do not want to add any of the other agents into the trade.

- We let q denote the number of trades that occur in the optimal allocation.
- Let $\text{OPT}(n, m)$ denote the GFT from the optimal allocation (which is not necessarily a DSIC/IR/BB mechanism) in a market with n buyers and m sellers.

decreasing		increasing
$b^{(1)}$	\geq	$s^{(1)}$
$b^{(2)}$	\geq	$s^{(2)}$
\vdots		\vdots
$b^{(q)}$	\geq	$s^{(q)}$
$b^{(q+1)}$	\leq	$s^{(q+1)}$
\vdots		\vdots
$b^{(n)}$	\leq	$s^{(m)}$

Figure 1: The optimal allocation.

Claim 1. The (post-trade) welfare is equal to the sum of the highest m values in the population.

Proof. The agents with items are in the green boxes.

The buyers with items are among the top m highest-valued agents. The buyers with items (in the green box) have higher values than those without as the buyers are ordered in decreasing order. They also have higher values than the sellers without because of the ordering of buyers and sellers, and $b^{(q)} \geq s^{(q)}$ as it is in the optimal allocation.

The sellers with items are among the top m highest-valued agents. The sellers with items (in the green box) have higher values than those without as the sellers are ordered in increasing order. They also have higher values than the buyers without because of the ordering of buyers and sellers, and $b^{(q+1)} < s^{(q+1)}$ as it is not in the optimal allocation.

Then the m agents with items have higher values than the n agents without, so the (post-trade) welfare is the sum of the values of those with items, which are the m highest values. □

The Buyer Trade Reduction (BTR) Mechanism [1]. The simple prior-free mechanism we will use is as follows, inspired by McAfee’s Trade Reduction mechanism [2]:

1. Solicit all buyer and seller values.
2. Compute the optimal allocation on the reported values.
3. Buy items at some p^S ; sell items to buyers at some p^B .
 - (a) Try to use the value of the highest-valued unmatched buyer as $b^{(q+1)} = p^B = p^S$ if sellers will accept. (They will, if $b^{(q+1)} \geq s^{(q)}$).
 - (b) If sellers will not accept (so $b^{(q+1)} < s^{(q)}$), then **reduce** the last trade between $b^{(q)}$ and $s^{(q)}$. Buy at $p^S = s^{(q)}$ and sell at $p^B = b^{(q)}$.

Let $BTR(n, m)$ denote the GFT from this mechanism in a market with n buyers and m sellers.

Observation 2 (DSIC+IR). *This mechanism is DSIC and ex-post IR because we set prices only using the values of non-winning agents, so winning agents pay prices lower than their values that they cannot impact.*

Observation 3 (Budget Balance). *Setting prices according to (3a) or (3b) satisfies weak budget balance.*

Claim 2. BTR reduces if and only if the $m + 1^{\text{st}}$ highest-valued agent is a seller.




Figure 2: The optimal allocation.

Proof. We are in case (3a) where $b^{(q+1)}$ works as a price for sellers exactly when $b^{(q+1)} \geq s^{(q)}$; otherwise we are in case (3b) and we must reduce a trade.

Because the number of trades in the optimal allocation is q by definition, then there are the q highest-valued buyers and $m - q$ highest-valued sellers with items among the top m highest-valued agents.

Who is next, as the $m + 1^{\text{st}}$? Either a buyer or a seller. If it is a buyer, then the $m + 1^{\text{st}}$ highest-valued agent is specifically $b^{(q+1)}$, and thus $b^{(q+1)} \geq s^{(q)}$, so we do not need to reduce.

However, if it is a seller, then the $m + 1^{\text{st}}$ highest-valued agent is specifically $s^{(q)}$, and thus $s^{(q)} > b^{(q+1)}$, so we are in case (3b) and must reduce the trade between $b^{(q)}$ and $s^{(q)}$ in order to set prices. \square

Theorem 4 (Babaioff G. Gonczarowski [1]). *When buyers and sellers are drawn i.i.d. from some distribution F , given an initial market with n buyers and m sellers, running Buyer Trade Reduction on a market with 1 additional buyer yields at least as much GFT as the optimal GFT in the initial market.*

$$BTR(n + 1, m) \geq \text{OPT}(n, m).$$

Proof. Approach: Aim to show that

$$\text{OPT}(n + 1, m) - \text{OPT}(n, m) \geq \text{OPT}(n + 1, m) - BTR(n + 1, m).$$

First we consider how OPT changes with augmentation of the market: $\text{OPT}(n + 1, m) - \text{OPT}(n, m)$.

A way to compare the larger and smaller markets is as follows. We draw $n + 1$ (or n) buyers and m sellers in the following way:

- Draw $n + 1 + m$ samples i.i.d. from F .
- (For a market with only n buyers, exclude 1 sample UAR.)
- Label m samples UAR as sellers.

Case 1: The excluded sample, to cut down to n , is among the lowest $n + 1$ values. Then there is no change in the highest m values between the two markets, thus, there is no change in the welfare or GFT between the two markets. This case occurs with probability $\frac{n+1}{n+1+m}$.

Case 2: The excluded sample, to cut down to n , is among the highest m values. This happens with probability $\frac{m}{n+1+m}$.

In this event, the change in OPT is the difference in the sum of the top m values—that is, the difference between the removed sample and what was previously the $m + 1^{\text{st}}$ sample which got pushed up to replace it.

$$\begin{aligned} &= \mathbb{E}[\text{sample removed} \mid \text{from top } m] - \mathbb{E}[m + 1^{\text{st}}] \\ &= \mathbb{E}[1 \text{ to } m \text{ UAR}] - \mathbb{E}[m + 1^{\text{st}}]. \end{aligned}$$

Now, we compare $\text{OPT}(n + 1, m) - BTR(n + 1, m)$ —how our mechanism does compared to OPT in the larger market.

Case 1: The $m + 1^{\text{st}}$ highest-valued agent is a buyer, in which case, BTR doesn't reduce, and there is no difference between the mechanism and OPT. This case occurs with probability $\frac{n+1}{n+1+m}$.

Case 2: The $m + 1^{\text{st}}$ highest-valued agent is a seller, in which case, BTR does reduce. This happens with probability $\frac{m}{n+1+m}$.

In this event, the difference between BTR and OPT is exactly the GFT from the reduced pair in expectation, $\mathbb{E}[b_q - s_q]$. But we know that $s^{(q)}$ is the $m + 1^{\text{st}}$ highest-valued agent, and $b^{(q)}$ is the lowest valued buyer among the top m valued agents, so the change is exactly

$$\mathbb{E}[\text{min buyer} \geq m] - \mathbb{E}[m + 1].$$

Putting this all together: We compare $\text{OPT}(n + 1, m) - \text{OPT}(n, m)$ and $\text{OPT}(n + 1, m) - \text{BTR}(n + 1, m)$, splitting both into 2 cases where we can *couple* the probabilities of these cases.

Case 1 occurs with probability $\frac{n+1}{n+1+m}$ and there is no change in either quantity.

Case 2 occurs with probability $\frac{m}{n+1+m}$ and we see

$$\mathbb{E}[1 \text{ to } m \text{ UAR}] - \mathbb{E}[m + 1^{\text{st}}]$$

as the change in OPT as the market is augmented, compared to the loss due to BTR in the large market which is

$$\mathbb{E}[\text{min buyer} \geq m] - \mathbb{E}[m + 1]$$

where the second quantity is clearly smaller, as it is the smallest buyer value in the top m vs. a random value in the top m . Hence

$$\text{OPT}(n + 1, m) - \text{OPT}(n, m) \geq \text{OPT}(n + 1, m) - \text{BTR}(n + 1, m)$$

and thus, we have proven the claim:

$$\text{BTR}(n + 1, m) \geq \text{OPT}(n, m).$$

(This type of argument is called a coupling argument.) □

Additional Results. This paper also investigates environments beyond when buyers and sellers are i.i.d. from the same distribution. They can be summarized below, where “ F_B FSD F_S ” refers to the buyer distribution F_B first order stochastically dominating F_S .

Bulow-Klemper-style results: $\text{BTR}(m_B + \text{sufficient } \# \text{ buyers}, m_S) \geq \text{OPT}(m_B, m_S)$.

#B, #S	Condition	Sufficient #B	Insufficient #B
m_B, m_S	i.i.d. ($F_B = F_S$)	1	0 [MS '83]
m_B, m_S	arbitrary F_B, F_S	Impossible (\implies)	0 any finite number
1, 1	F_B FSD F_S	4	1
$m_B, 1$	F_B FSD F_S	$4\sqrt{m_B}$	$\lceil \log_2 m_B \rceil$
m_B, m_S	F_B FSD F_S	$m_S(m_B + 4\sqrt{m_B})$	above

Sample-style results: The approximation that BTW with a single sample from the buyer distribution gives to $\text{OPT}(1, 1)$.

Condition	Approx	LB
i.i.d. ($F_B = F_S$)	1/2	$> 1/2$
F_B FSD F_S	1/4	$> 7/16$

References

- [1] Moshe Babaioff, Kira Goldner, and Yannai A. Gonczarowski. Bulow-klemperer-style results for welfare maximization in two-sided markets. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2452–2471. SIAM, 2020.
- [2] R Preston McAfee. A dominant strategy double auction. *Journal of economic Theory*, 56(2):434–450, 1992.
- [3] Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of economic theory*, 29(2):265–281, 1983.