Not quite the best response: Truth-telling, strategy-proof matching, and the manipulation of others[*]

Pablo Guillen

The University of Sydney

Rustamdjan Hakimov

WZB, Berlin Social Science Center

November, 2016

Abstract
Following the advice of economists, school choice programs around the world have lately been adopting strategy-proof mechanisms. However, experimental evidence presents a high variation of truth-telling rates for strategy-proof mechanisms. We crash test the connection between the strategy-proofness of the mechanism and truth-telling. We employ a within-subjects design by making subjects take two simultaneous decisions: one with no strategic uncertainty and one with some uncertainty and partial information about the strategies of other players. We find that providing information about the out-of-equilibrium strategies played by others has a negative and significant effect on truth-telling rates. That is, most participants in our within-subjects design try and fail to best-respond to changes in the environment. We also find that more sophisticated subjects are more likely to play the dominant strategy (truth-telling) across all the treatments. These results have potentially important implications for the design of markets based on strategy-proof matching mechanisms.

[*] We are grateful to Dorothea Kübler, Frank Heinemann, Bettina Klaus and Glenn Harrison for their kind comments and support; to Nina Bonge for helping us programming the experiments and conducting the sessions; and to Jennifer Rontganger for the copy editing.

1. Introduction

Matching theory has been extremely successful in providing the mechanisms used for the design of markets in the real world. Matching mechanisms are currently used for markets such as medical schools for graduates (Roth, 1984; Roth and Peranson, 1999), housing for students (Chen and Sönmez, 2002; Abdulkadiroglu and Sönmez, 2003), school choice (Abdulkadiroglu et al., 2005), and kidney exchange (Roth et al., 2004). When it comes to its practical application, one of the most important advantages of any mechanism is its strategy-proofness. That is, if participants could be convinced of the impossibility to manipulate, they would then devote their energy to discovering their own preferences. For instance, investigating which schools are best suited for them, rather than devising strategies to game the system.

There are two leading strategy-proof matching mechanisms recommended by market designers for school choice: Top Trading Cycles (TTC), see Abdulkadiroglu and Sönmez (2003) and Deferred Acceptance (DA), see Gale and Shapley (1962). TTC and DA have competing properties. That is, other than being strategy-proof TTC is Pareto optimal, but not envy-free. DA is envy-free but not Pareto optimal. Both mechanisms have been adopted by school boards. The Boston Public School system chose to use DA although market designers recommended TTC, see Abdulkadiroglu et al (2005). The New Orleans Recovery District adopted TTC in 2012 (Vanacore, 2012).

There is an ongoing debate on whether strategy-proofness can be safely assumed for the real-life implementation of a matching algorithm. Early matching experiments (i.e., Chen and Sönmez, 2002; Chen and Sönmez, 2006)[1] suggest truth-telling rates are higher for strategy-proof mechanisms than for non-strategy-proof mechanisms. However, this result might be driven by the fact that the non-strategy-proof mechanism used for comparison (immediate acceptance) is easy to manipulate, in the sense that it is easy to find a seemingly good or satisfactory way to manipulate them. Conversely, the low manipulation rates found for strategy-proof mechanisms may not be caused by the participants' understanding of strategy-proofness, but by them being unable to find a satisfactory manipulation strategy, thus leading them to report a default option—the induced preference order. Guillen and Hing (2014) give

[1] Other noticeable experimental papers in the literature are Kagel and Roth (2000), Haruvy and Unver (2007), Echenique et al. (2016), Niederle and Yariv (2009), Featherstone and Mayefsky (2011), Chen and Kesten (2016), Hugh-Jones et al. (2014), Klijn et al. (2013), Hakimov and Kesten (2014), Niederle et al (2013), Featherstone and Niederle (2013), Braun et al. (2014).

some support to these ideas by showing how manipulation becomes modal when wrong advice is introduced. In a similar vein, Pais and Pintér (2008) and Pais et al. (2011) find that manipulation rates increase when more information about the underlying preferences of other participants is introduced. To sum up, truth-telling rates vary as a response to theoretically irrelevant changes in the environment. Thus, among the previously safely assumed strong connection of observed truth-telling rates in the lab theoretical strategy-proofness is very much in doubt. One of the main aims of the current study is to crash-test this connection.

Strategy-proofness implies that truth-telling is the best response to any strategy chosen by other players. Some evidence indicates manipulation attempts in the real-life implementation of theoretically strategy-proof mechanisms. For instance, Guillen and Hing (2014) cite popular blogs that encourage manipulation in the Boston Public School (BPS) deferred-acceptance-based system. Fisher (2009) elaborates on the general dysfunctionality of the quasi strategy-proof National Resident Matching Program (NRMP) system, Nagarkar and Janis (2012) point out the fact that "*Advisors occasionally tell* [NRMP] *applicants to realistically consider their chances of matching at a program when determining its position on their rank lists.*" In the same vein a survey-based study, Rees-Jones (2015), presents evidence of attempts at misrepresentation in the NRMP match. Another recent study, Hassidim et al. (2015), reaches similar conclusions when studying the market for graduate psychologists in Israel. Given the evidence, it is reasonable to hypothesize that the actions of participants may be influenced by the likely manipulation of other participants.[2] We use this conjecture as a motivation to test the connection between truth-telling and the understanding of strategy-proofness by varying, within subjects, the information available on the strategies chosen by other players.

We make use of an individual decision-making set-up to precisely control the amount of information on the preferences, underlying and/or submitted, by computer-simulated players (computer players) to human participants. All the subjects in our experiment played two treatments simultaneously:[3] a full information deterministic baseline and one out of four treatments with different amounts of information on the strategies of other participants. The

[2] In a recent paper, Ding and Schotter (2014a, 2014b) show that in a repeated environment chatting and intergenerational advice decrease the truth-telling of subjects. This evidence only supports our hypothesis of conditioning the strategies on the behavior of other players.

[3] Subjects are asked to fulfill rankings for both baseline and treatment in the same decision screen. Both rankings are submitted together.

latter four treatments are: uncertain misrepresentation (UMT), in which participants know the underlying preferences of the computer players and they know that at least one of them will not report truthfully; certain blocking misrepresentation (CBMT), in which the underlying preferences are known and human players are informed that a computer player is misrepresenting its preferences in a particular manner such that the human's first preference is blocked under the assumption of the truth-telling of other computer players. However, we do not point this out to subjects but show only that one computer player misrepresents her preferences; certain unblocking misrepresentation (CUMT), in which one of the computer players misrepresents her underlying preferences in a way that does not affect the human player's chances of getting the top choice; and underlying preferences (UPT), in which subjects know the underlying preferences of computer players, but nothing about how they are reported, other than that the computer players will maximize their profit.

The baseline allows participants to use the TTC algorithm to find the best response to the perfectly-known behavior of computer-simulated agents. Over 62% of the subjects truthfully report the full preference list and 78% report their true top choice (sufficient to maximize the payoff) in the baseline. Truthful preference revelation decreases greatly and significantly in each of the misrepresentation treatments. Nevertheless, there is no significant difference within subjects between truth-telling rates in the Baseline and UPT. This leads us to draw the conclusion that the high truth-telling rates in Baseline or UPT cannot be attributed to subjects' understanding of strategy-proofness: information about the misrepresentation by computer players leads human subjects to misrepresent more often. We observe how the majority of subjects, 69%, behave as if they try and fail to best-respond to changes in the environment.

We cannot reject the understanding of the dominant strategy property of TTC for 31% of the subjects as they submitted their true preference orders in the two treatments they played. The percentage grows to 34% among subjects who solved the allocation task correctly (note that this difference is not significant). Additional tests allowed us to conclude that these subjects are more likely to achieve a higher score in the Cognitive Reflection Test (Frederick, 2005), the Wonderlic IQ test and be more successful in answering multiple-choice questions about the mechanism's properties. In contrast to Klijn et al. (2013) we find no significant difference between the risk-aversion of subjects who played optimal and defensive strategies.

The rest of the paper is structured as follows: in section 2 we justify the experimental design and treatments, while we explicitly formulate our hypotheses in section 3. Section 4 presents the results. It is followed by the concluding remarks in section 5.

2. Experimental design

We design an experiment to compare the individual decisions of participants in matching markets in the lab under the Top Trading Cycles mechanism (TTC). We use TTC for the school choice problem with a preliminary assignment by Abdulkadiroğlu and Sönmez (2003). In a school choice problem, a certain number of students are to be assigned to a certain number of schools. Each school has a certain number of available slots (capacities), and the total number of slots is no less than the number of students. Let $I = \{i_1; i_2; ::: ; i_n\}$ denote the set of students, and $S = \{s_1; s_2; ::: ; s_n\}$ denote the set of schools. Each student has strict preferences over all schools. A strict priority order of all students for each school is exogenously given.

Then TTC works as follows (Abdulkadiroğlu and Sönmez, 2003):

"**Step 1:** Assign a counter for each school which keeps track of how many seats are still available at the school. Initially set the counters equal to the capacities of the schools. Each student points to her favorite school under her announced preferences. Each school points to the student who has the highest priority for the school. Since the number of students and schools are finite, there is at least one cycle. (A cycle is an ordered list of distinct schools and distinct students (s_1 , i_1, s_2,..., s_k, i_k) where s_1 points to i_1, i_1 points to s_2 .... s_k points to i_k, i_k points to s_1.) Moreover, each school can be part of at most one cycle. Similarly, each student can be part of at most one cycle. Every student in a cycle is assigned a seat at the school she points to and is subsequently removed. The counter of each school in a cycle is reduced by one and if it is reduced to zero, the school is also removed. The counters of all the other schools stay put.

In general, at **Step k:**

Each remaining student points to her favorite school among the remaining schools and each remaining school points to the student with the highest priority among the remaining students. There is at least one cycle. Every student in a cycle is assigned a seat at the school that she points to and is subsequently removed. The counter of each school in a cycle is reduced by one and if it is reduced to zero the school is also removed. The counters of all the other schools remain in place. The algorithm terminates when all students are assigned a seat. Note that there can be no more steps than the cardinality of the set of students."

For the aim of the experiment instructions we used the common formulation of TTC by Chen and Sönmez (2006).

We do not aim to simulate the complexity of the real-world school allocation problem, but rather to create a simple artificial environment in which we can test the consistency of the decision of subjects and the effect of the amount of information on the reported preferences of others. The preference profiles of participants are fixed across all treatments, as are the priorities of students in schools. An experimental subject represents one-out-of-four students in a market. The other three students are played by the computer. We choose a small market to keep things as simple as possible. So there are four schools in the market with one slot each. The preferences of players are designed in such a way as to ensure the decisive power of the human player. A misrepresentation of preferences will cause a suboptimal outcome in all treatments but one. The priorities of students in schools are generated through the district school priority, in which each player has a priority only to the school in its own district. The preferences of students and the priorities of the school for all environments are as follows (tables 1 and 2).

Table 1. Student priorities

Home school	Student
A	Computer 1
B	Computer 2
C	Human
D	Computer 3

Table 2. Underlying preferences

	Human	Computer 1	Computer 2	Computer 3
Top choice	A	B	D	C
2nd choice	B	C	C	D
3rd choice	D	A	A	A
4th choice	C	D	B	B

In all treatments, subjects received 10 euros if they were assigned to their top choice (school A), 7 euros if they were assigned to their second choice (school B), 4 euros if assigned to their third choice (school D), and 1 euro if assigned to their last choice (school C). Allocations were implemented through a centralized clearinghouse (see the instructions in the online supplementary material). We insure that the home school is the least preferred choice for the human player in order to make the typical district-school bias manipulation as costly as possible. This structure of the priorities and preferences is common for all five treatments. The baseline is a fully deterministic game as participants know that the computer players will send their true preferences to the clearing house. The baseline treatment is played by all subjects *simultaneously* with one of the other four treatments.

2.1. Treatments

The treatments are described below:

*1. Baseline treatment***.** In the baseline treatment subjects know the underlying preferences of the computer players and are aware that the computer players are submitting their true preferences. The game is deterministic. Subjects know the exact inputs in the mechanism and should be able to calculate the outcome. Subjects are not required to understand strategy-proofness to behave optimally because the top trading cycle of exchanges of the top choices of all participants should be obvious (there are no conflicts of the top choices).

2. Uncertain misrepresentation treatment (UMT). In this treatment the participants are aware of the underlying preferences of the computer, and they know that at least one of the computer players did not report its preferences truthfully. They do not know the way in which the preferences were misrepresented. In this treatment subjects need a deeper understanding of the mechanism to make the optimal decision, truth-telling, as the mechanical calculation of outcomes is no longer an option.

3. Certain blocking misrepresentation treatment (CBMT). In this treatment subjects know the underlying preferences of the computer players and are aware that computer player 1 will submit A-B-C-D instead of its true preference B-C-A-D and other computer players behave to maximize their payoffs. This misrepresentation by computer player 1 blocks the top choice of the human player. In this treatment there is more than one payoff-maximizing strategy. As their top choice is blocked, subjects can swap their first and second preferences. Subjects with an understanding of the dominant strategy property of the TTC should not invest time in calculating the outcomes under different strategies and should still submit the true list.

4. Certain unblocking misrepresentation treatment (CUMT). In this treatment subjects are aware of the underlying preferences of the computer players. They also know that computer player 2 will submit C-D-A-B instead of his true preference D-C-A-B and other computer players behave to maximize their payoff. This misrepresentation, however, does not influence the possibility of the participants getting their top choice. The only payoff-maximizing strategy in this treatment is to send the true list (strictly speaking, the true top choice).

5. Underlying preferences treatment (UPT). The participants only know the underlying preferences of the computer players. They are also informed that the computer players will state their preferences in such a way so as to maximize their payoffs. This informational structure makes this treatment the most similar to the usual implementation of a matching experiment in an incomplete information environment.

As was previously mentioned, every subject played two of the above treatments: the baseline plus one of the other four treatments. Subjects had to submit their preferences *simultaneously* for the two treatments they played. Only the resulting allocation of one of the two treatments, randomly chosen, was payoff relevant. Subjects could influence which treatment was to be chosen for the payoff.[4]

2.2. Additional controls

The following tasks were performed in the lab by every subject. Subjects had to complete two tasks before the main body of the experiment:

- Immediately after reading the instructions subjects were asked to use the TTC algorithm to solve an example of the allocation problem. The structure of this task is very similar to the example in the experimental instructions (see the online supplementary material). Subjects received 2 euros for finding the correct allocation, but they only learn the result at the end of the experiment. [Allocation]
- After Allocation, subjects were also asked to provide answers to two multiple-choice questions about features of the TTC mechanism (see online supplementary

[4] More precisely, subjects had three options: 55% chance of baseline to be payoff relevant and 45% chance of treatment, 50% chance for baseline and treatment and 45% chance for baseline and 55% chance for treatment. The vast majority of subjects chose 50-50.

material for exact formulations). They received 50 cents for each correct answer. [MC]

The following tasks were run after subjects submitted their decisions for the main experiment:

- First of all, we ran two cognitive ability tests.
    a. The well-known three-question Cognitive Reflection Test (Frederick, 2005). All three questions were put on one screen and subjects had 2.5 minutes to submit their answers. They received 50 cents for each correct answer. [CRT]
    b. 10-question Wonderlic cognitive ability test (Wonderlic and Hovland, 1939). Subjects had three minutes to provide answers to 10 questions, one after another. They received 30 cents per correct answer. [Wonderlic]
- Finally, we elicited subjects' risk-aversion. We used The Bomb Task (Crosetto and Filippin, 2013). The details are presented in the online supplementary material. [Risk]

2.3. Procedures

Nine experimental sessions were run in the laboratory for the economic experiments of the Technical University Berlin between November 2012 and June 2015. In total, 214 experimental subjects participated in the experiment. Most of them were at the time students at Berlin universities. Twelve subjects were not able to submit their ranking lists within the 10 minutes provided. So, only 202 data points were used in the subsequent analysis. The average length of the session was 80 minutes, and subjects earned 15.07 euros on average.

3. Hypotheses

As the TTC mechanism is strategy-proof it is at least a weakly-dominant strategy to state the truth in all treatments. Note that truth-telling is not the only payoff-maximizing strategy in CBMT, as the top choice is blocked by the computer players.

Hypothesis 1: Subjects should reveal their true preferences in all treatments.

Hypothesis 1 is based on the strategy-proofness of TTC. However, from the previous experimental literature on matching we know that it is highly unlikely that all experiment participants will be able to follow the dominant strategy. We expect – inspired by the anecdotal evidence cited in the introduction – that the amount of information about the behavior of the other participants would be a factor explaining the misrepresentation of the preferences. Thus, we state an alternative hypothesis in the following way:

Hypothesis 1a: The rate of truthful preference revelation should be higher in the baseline than in other treatments.

We form Hypothesis 1a, because the baseline does not require an understanding of strategy-proofness due to the deterministic nature of the treatment.

Additionally, we plan to explore the result for subsamples of subjects who managed to find the allocation correctly or not at the beginning of the experiment. As it is typical to attribute some degree of manipulation to a poor understanding of the mechanics of the mechanism, we formulate the following hypothesis:

Hypothesis 2: Subjects who are able to solve the allocation task reveal their true preferences more often and are more consistent than other subjects with their decisions.

4. Results

Result 1 (Truthful preference revelation, treatment effect.): Truthful revelation is higher in the Baseline than in UMT, CBMT, and CUMT. The proportion of the submitted list that is compatible with the best response in the Baseline is higher than in UMT, CUMT, and UPT.

Support: Table 3 (panel A) shows the truth-telling rates by treatment. Panel B shows the proportion of submissions compatible with best-response[5] behavior by treatment (*p*-values for two-sided Fisher's exact tests are presented in Table 4). Misrepresentation is quite common across all treatments and thus there is little support for hypothesis 1. Hypothesis 1a cannot be rejected. The highest truth-telling rates are in the baseline treatments where, on average, 62% of subjects report the full list truthfully and 78% report the top choice truthfully. Note that we observe some variation of the proportions of truthful reporting of the full list and the top

5 Note that the true top choice is a best response in the Baseline. Reporting the true top choice in any other position of the list is irrelevant for a payoff in the CBMT, and thus all submissions with this kind of manipulation are compatible with best response reporting. In all other treatments best response requires the full list to be reported truthfully.

choice within the baseline, depending on the other treatment, played by the same subject, however, there is no significant difference in these proportions (the minimum two-sided Fisher's exact *p*-value are for the comparison of truthful reporting in Baseline, given UMT, versus Baseline, given CUMT is 0.15; the minimum two-sided Fisher's exact p-value are for the comparison of truthful top choice reporting in Baseline, given UMT, versus Baseline, given CUMT is 0.14; all other *p*-values for pairwise comparisons are higher than 0.30).

In summary, we find evidence of subjects conditioning their behavior on the information about the strategies of others. That is, most subjects who submitted truthfully in the baseline do not do so because they understand strategy-proofness.

There is no significant difference in reporting full truthful lists when comparing, Baseline and UPT within subjects. The difference is substantial, though insignificant. Note that this difference becomes significance when we pool the baselines from all subjects and compared them with UPT (two-sided Fisher exact *p*-value 0.03), see Full Baseline in Table 4.[6]

Also, note that subjects best-respond in Baseline more often than in all the other treatments, with the exception of CBMT. In CBMT, due to the fact that the true top choice is blocked by the misrepresentation of computer player 3, it is sufficient to keep the relative order of three other schools, while putting the top choice anywhere in the list. We tend to interpret the "success" of subjects in the CBMT with caution, as it might to a large extent be driven by a random attempt to manipulate the submitted list. Note, that the best response in CBMT includes a switch of the first and second choices in the reported rankings – the most common way of preference manipulation in other misrepresentation treatments (we observe it happening 13 times in UMT and 20 in CUMT).

The decrease in truth-telling rates from the baseline to each of the limited information treatments indicates that there is a high proportion of subjects who could previously be falsely classified as subjects who understand the incentive properties of the mechanism. Irrational manipulation becomes the modal behavior across the limited information treatments. This result shows that without any additional explanation of the properties, subjects in the lab tend to misreport their preferences. It could also be evidence of the fact that the high percentage of

[6] In Table 4 we report pairwise treatment comparisons. Note that the Baseline vs Treatment comparison are the within-subjects. Treatment vs Treatment comparisons are between-subjects.

truth-telling in the no-information matching experiments[7] (i.e., Chen and Sönmez, 2006) may be driven by the default option of truthful reporting, rather than an understanding of the incentive properties of the mechanisms (Guillen and Hing, 2014).

Table 3. Preference reporting

Panel A (Truthful reporting of preferences)			
N	Treatment	Truth (Baseline)	Truth (Treatment)
51	UMT	36/51 (70%)	21/51 (41%)***
50	CBMT	31/50 (62%)	15/50 (30%)***
50	CUMT	28/50 (56%)	14/50 (28%)***
51	UPT	31/51 (61%)	23/51 (45%)
Panel B (Best responses)			
N	Treatment	Best response (Baseline)	Best response (Treatment)
51	UMT	44/51 (86%)	21/51 (41%)***
50	CBMT	38/50 (76%)	36/50 (72%)
50	CUMT	37/50 (74%)	14/50 (28%)***
51	UPT	39/51 (76%)	23/51 (45%)***

***- significant at 1% level, **- significant at 5% level, *-significant at 10% level.

Finally, there are a few subjects who report truthfully in the limited information treatments and do not report truthfully in the baseline (three subjects in each of the limited information treatments and two subjects in UPT). Clearly, these subjects do not understand how the first step of TTC works, let alone its strategy-proofness. Excluding them from the truth-telling sample of the limited information treatment would only strengthen the significance of the result.

How do subjects misrepresent their preferences? In the baseline, after truth-telling, the second most common report is A-B-C-D (20 out of 202 subjects). That is, moving the district school from the fourth to the third rank in the list, a district school bias. The third most common misreporting is the switch of the first and the second choice (thus report B-A-D-C), 18 out of 202 subjects report it in the baseline. That is also the most common attempt of

[7] The highest truth-telling rates are found in experiments in which no information is provided about either the stated or underlying preferences.

manipulation across all the treatments. That accounts for 13 out of 51 reports in UMT, 16 out of 50 reports in CBMT, 20 out of 50 reports in CUMT, and 14 out of 51 reports in UPT.

Table 4. Pairwise treatment comparisons

Panel A	Truthful reporting (full list)			
	UMT	CBMT	CUMT	UPT
Baseline	0.005	0.002	0.008	0.165
UMT	-	0.300	0.211	0.842
CBMT	-	-	1.000	0.151
CUMT	-	-	-	0.099
Full Baseline*	0.007	0.000	0.000	0.027
Panel B	Best-response reporting			
	UMT	CBMT	CUMT	UPT
Baseline	0.002	0.820	0.000	0.000
UMT	-	0.003	0.211	0.842
CBMT	-	-	0.000	0.009
CUMT	-	-	-	0.099
Full Baseline*	0.000	0.353	0.000	0.000

*Full Baseline considers all 202 Baseline observations

Next we analyze how market outcomes are influenced by individual behavior. We consider two outcomes: payoff of human subject and the efficiency of the allocation. Note that we ignore the fairness aspects of the allocation, as our schools do not have priorities other than district school priority.

Result 2 (Efficiency): The average payoff of subjects in the baseline is significantly higher than in every other treatment. Allocations in the baseline are, on average, significantly more efficient than in every other treatment.

Support: Table 5 (panel A) presents the average payoff of subjects in the experiment (human players) by treatments. Subjects have the highest average payoff from the allocation in the baseline treatment, and it equals 9.05 euros (note that the maximum payoff is 10 euros). In all other treatments, the payoff is significantly smaller. Note that this happens by design in CBMT. The maximum possible payoff for the subject is only 7 euros because the top choice is blocked by computer player 1. Conversely, the high misreporting rates in CUMT do not have proportional impact on payoff differences, which are only marginally different (10%

level) to the baseline. This is explained the fact that in 10 out of 36 misreported submitted lists the reported top choice was the true top choice, and that allowed subject earn 10 euros.

Table 5. Subjects payoffs and overall efficiency

Panel A: Average payoffs				
N	Treatment	Baseline	Treatment	*p*-value
51	UMT	9.47	7.88***	0.00
50	CBMT	9.04	6.22***	0.00
50	CUMT	8.92	8.20*	0.055
51	UPT	8.78	7.65***	0.00
Panel B: Average efficiency				
N	Treatment	Baseline	Treatment	*p*-value
51	UMT	97.4%	89.4%***	0.00
50	CBMT	95.2%	70.6%***	0.00
50	CUMT	94.6%	73.0%***	0.00
51	UPT	91.7%	86.7%***	0.00

***- significance under 1% level, **- significance under 5% level, *-significance under 10% level, Wilcoxon matched pairs test for equality of the value for option A and option B.

Table 5 (panel B) shows the efficiency of allocations. Efficiency is calculated as the sum of payoffs of all players divided by the sum of the payoffs in the Pareto efficient allocation. Note that for computer players we use the same cardinal payoffs for the first, second, third, and fourth choices as for the human player. In the Pareto efficient allocation, all players receive their top choice, and thus the sum of payoffs is 40 euros. The average efficiency in the baseline is 94.7%. That is high, and is significantly higher than in every other treatment. Note, however, that in CBMT and CUMT part of the efficiency loss is driven by misrepresentations of computer players and thus appear “by design.”

Overall, we can conclude that, as expected, individual misreporting has a strong effect both on the payoff of the human player and on the efficiency of the allocation.

Next we turn back to individual behavior and analyze misreporting subjects in detail. First of all we check whether the understanding of the mechanism shown by the subjects who are able to solve the allocation task has an effect on truth-telling in the Treatment.

Result 3 (Allocation task and truth-telling)**:** Subjects who solved the allocation task correctly report truthfully in Baseline significantly more often than subjects who failed to reach the correct solution in the allocation task. Nevertheless, they do not report their truthful preferences more often in the Treatment.

Support: In total, 112 out of 202 subjects solved the allocation task correctly. The average percent of truthful reported lists in Baseline is 70% for those who solved the task correctly and 53% for those who failed. The difference is significant (two-sided Fisher's exact test p-value 0.02). Table 6 presents the truthful reporting rates by subjects who solved the allocation task correctly. The truth-telling rates in Baseline are significantly higher than in UMT, CBMT, and CUMT (two-sided Fisher's exact test p-values are presented in the last column of Table 6). Thus, we observe though that a "poor" understanding of the mechanism can explain some percent of manipulation. Critically, it cannot explain the inconsistency of choices across the two decisions. Even for the subjects who used their understanding of the mechanism to solve the allocation task, truth-telling cannot be attributed to an understanding of strategy-proofness.

Table 6. Truthful reporting conditional on correct allocation task

N	Treatment	Baseline	Treatment	*p*-value
34	UMT	27/34 (79%)	14/34 (41%)***	0.003
25	CBMT	17/25 (68%)	8/25 (32%)**	0.023
28	CUMT	17/28 (61%)	8/28 (29%)**	0.031
25	UPT	17/25 (68%)	13/25 (52%)	0.387

Our design allows us to classify subjects into the following categories (see Table 7):

- **Dominant strategy:** Subjects who played *as if* they understood strategy-proofness in both the baseline and the other treatment. These are the subjects who submitted the full true lists in both the baseline and the treatment they played.
- **Best response:** Subjects who were able to play best response to the market, but failed to report truthfully. (For instance, subjects who submitted only their true top choice in the baseline, or subjects who skipped the blocked top choice in CBMT.)
- **Bias:** Subjects who played best response in the baseline, but manipulated in a limited information treatment.
- **Limited ability:** Subjects who failed to best-respond (reveal at least their top choice truthfully) in the deterministic, full information baseline.

Result 4 (Behavior in line with strategy proofness): Only 31% of subjects behaved as if they understood strategy-proofness.

Support: For 31% of subjects we cannot reject the understanding of the dominant strategy concept. We emphasize the relatively strong requirement for this categorization, as even in case of certainty, where only the true top choice matters for allocation, we require the full truthful list to be submitted. In addition, 9% of subjects (40% in total, dominant strategy + best response) were able to maximize their payoffs. Almost all of the "best response" subjects (17 out of 19) played CBMT, which is most likely driven by random luck, as discussed above, due to higher number of strategies compatible with best response. The proportion of subjects who are categorized as understanding the dominant strategy is not significantly different for the subsample of subjects who understood the mechanics of the mechanism, which is in line with Result 2.

Table 7. Distribution of subjects between categories

	Full sample	Correct allocation task
Dominant strategy	62 (31%)	38 (34%)
Best response	19 (9%)	10 (9%)
Bias	77 (38%)	41 (37%)
Limited ability	44 (22%)	23 (21%)

The modal category is Bias. It includes subjects who would be treated as though they understood the dominant strategy if only the baseline was played. The fact that this category is modal and includes 38% of the subjects emphasizes the importance of a cautious interpretation of truth-telling in one-shot experiments as a sign of understanding strategy-proofness.

The allocation task is not a good predictor of truth-telling in the treatment. Is there any other control predicting truth-telling? We now proceed to analyze the effect of the MC, Wonderlic, CRT, and Risk tasks.

After trying to solve the allocation task, subjects were asked to answer the following multiple-choice question[8] about the mechanism and received 50 cents for each correct answer:

Which of the following statements about the mechanism is correct?:

a. *Before selecting what ranking to choose, students should be careful to avoid applying to the most popular school.*
b. *Knowing the preferences and ranking of the others is crucial for choosing your own ranking list.*
c. *The mechanism is constructed in such a way that the ranking list should always coincide with your true preferences.*
d. *You should only state your true preferences if you are certain that the other participants will also state their true preferences.*

Once participants submitted their preferences in the main experimental task, they were given 2.5 minutes to answer the three CRT questions. Then they went over a 10-question, three-minute Wonderlic test (see the online supplementary material). No subject was able to finish 10 questions within the given time. Participants received 50 cents for answering any CRT or Wonderlic question correctly. Finally, after the CRT and Wonderlic, we elicited risk preferences using the so-called "bomb task" by Crosetto and Filippin (2013).

Result 5 (Dominant strategy categorization and performance in side tests): Subjects are more likely to behave as if they understand the dominant strategy property of TTC if they are successful in answering the mechanism-related multiple choice question or if their CRT performance and Wonderlic test performance are higher.

Support: Table 8 shows the marginal effects of the Probit regression for the dominant strategy category dummy. Note that CRT and Wonderlic were highly correlated, so we generated the variable which is the sum of the scores in these tests.

---

[8] In the experiment, subjects were also asked the following question:

The allocation procedure is constructed in such a way to guarantee students an assignment at least as good as their district school, according to the ranking list: True or False. However, we do not include the answers in our analysis, as a lot of subjects did not understand the formulation. In the post-experiment questionnaires subjects complained (those who answered incorrectly), that in fact you can get a worse school than your district school if you listed the worse school higher in the ranking lists. Thus, we concluded that the question was not clearly formulated and excluded it from our analysis.

Table 8. Probit regression, marginal effects

Predict (Dominant strategy category dummy)	Marginal effects
CTR+Wonderlic test score	0.038*** (0.014)
Correct solution of allocation task	0.007 (0.069)
Correct answer in mechanism multiple question	0.206*** (0.067)
Number of bombs collected	0.002 (0.001)
Male	-0.001 (0.072)

Standard errors are in brackets. ***- significant at 1% level, **- significant at 5% level, *- significant at 10% level.

Result 5 shows that seemingly smarter subjects are more likely to understand that TTC has a dominant strategy, no matter the information given about the behavior of other market participants.

5. Conclusion

Strategy-proofness is often cited as being one of the most important properties regarding the practical implementation of matching mechanisms. This line of thought has been even further encouraged by laboratory experiments in which, often, the majority of subjects behaved *as if* they understood strategy-proofness. However, more recent experimentation indicates that the rates of truthful revelation decrease when subjects are given advice or a certain amount of information. Our experiment is also inspired by real evidence of misrepresentation in the application of theoretical strategy-proof mechanisms. The idea behind our design is that, for a boundedly-rational individual, misrepresentation may feel compelling just because others do it. Our experiment gives strong support to this idea: 69% of participants react to the environment by trying, and failing, to best respond. A behavior compatible with the dominant strategy play can only be found for the remaining 31% of the subjects. Additionally, we also find out that more able subjects are more likely to play the dominant strategy.

A key question is, can our result be attributed to an experimenter demand effect? According to Zizzo (2010) a demand effect is a change in behavior by experimental subjects due to cues about what constitutes appropriate behavior (behavior demanded from them). Our experimental instructions (see the online supplementary material) clearly state, at the very beginning: *"The instructions are simple, and if you follow them carefully and make good decisions you might earn a considerable amount of money which will be paid to you in cash at the end of the experiment."* That is, if anything and in line with a vast body of experimental research, we are cueing a profit-maximizing behavior. This cue is further stressed by the inclusion of a solved example in the instructions and an incentive-based allocation task. Even further, once our subjects went through such a thorough training,[9] the baseline treatment has an easy-to-calculate, profit-maximizing best response. An experimenter demand effect implies that subjects are willing to give up profits in order to please the experimenters. It could be argued that this is precisely the way subjects behave when reacting to computer misrepresentation, in spite of our profit-maximizing cues. Result 4, however, suggests that this is not the case. Indeed, the more sophisticated subjects are, the less prone they are to copying computer misrepresentation. If a demand effect affects *only* unsophisticated players it can hardly be called a demand effect anymore.

The fact that cleverer subjects are able to perform significantly better suggests there may be some room to improve the education of participants in real-world markets. That is very much in line with the insight in Chen and Sönmez (2006). We have unveiled yet another potentially serious shortcoming of designed matching markets, see also Guillen and Hing (2014) and Ding and Schotter (2014b) for more disappointing news. Market designers should take note: mechanisms with good theoretical properties like strategy-proofness may be the most adequate in many circumstances, but behavior stemming from good properties cannot be taken for granted. We believe that those problems can only be addressed by a combination of education and/or effective and credible advice coming from a trustworthy source (see Guillen and Hakimov, 2015, for some encouraging results on the effects of advice). That is the goal of the ongoing investigation.

[9] Thorough in relative terms: in the typical economic experiment subjects get just 15minutes to read experimental instructions before they are asked to participate in incentive-based tasks. The training of participants in real-life market design tasks based on TTC is far less intense.

References

Abdulkadiroğlu, A. and Sönmez, T. (2003). "School Choice: A Mechanism Design Approach." American Economic review 93(3): 729–747.

Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., & Sönmez, T. (2005). "The Boston public school match." American Economic Review, 368-371.

Braun, S., Dwenger, N., Kübler, D., and Westkamp, A. (2014). "Implementing Quotas in University Admissions: an Experimental Analysis." Games and Economic Behavior, 85, pp. 232-251.

Chen, Y., and Kesten, O. (2016). "College and High School Admissions Reforms in China: An Experimental Analysis." Working paper..

Chen, Y. and Sönmez, T. (2002). "Improving Efficiency of On-Campus Housing: An Experimental Study." The American Economic Review 92(5): 1669–1686.

Chen, Y. and Sönmez, T. (2006). "School Choice: an Experimental Study." Journal of Economic Theory 127(1): 202–231.

Crosetto, P. and Filippin, A. (2013). "The Bomb Risk Elicitation Task." Journal of Risk and Uncertainty, 47 (1): 31–65.

Ding, T, and Schotter, S. (2014a). "Matching and Chatting: An Experimental Study of the Impact of Network." Working Paper.

Ding, T, and Schotter, S. (2014b) "Intergenerational Advice and Matching: An Experimental Study." Working Paper.

Echenique, F., Wilson, A. and Yariv, L. (2016). "Clearinghouses for Two-Sided Matching: An Experimental Study." Quantitative Economics, Volume 7(2), 449–482.

Featherstone, C and Mayefsky, E. (2011). "Why Do Some Clearinghouses Yield Stable Outcomes? Experimental Evidence on Out-of-Equilibrium Truth-Telling." Working paper.

Featherstone, C., & Niederle, M. (2013). "Improving on strategy-proof school choice mechanisms: An experimental investigation." Working Paper.

Frederick, S. (2005). "Cognitive Reflection and Decision Making." Journal of Economic Perspectives, 19(4), 25–42.

Fisher, C. (2009). "Manipulation and the Match." JAMA: The Journal of the American Medical Association 302.12: 1266–1267.

Gale, D., & Shapley, L. S. (1962). "College admissions and the stability of marriage." The American Mathematical Monthly, 69(1), 9-15.

Guillen, P. and Hakimov R. (2015). “The Effectiveness of Top-Down Advice in Strategy-Proof Mechanisms: a Field Experiment”. Working paper.

Guillen, P. and Hing, A. (2014). “Lying through their teeth: Third party advice and truth telling in a strategy proof mechanism.” European Economic Review, 70(October), 178-185.

Hakimov, R. and Kesten, O. (2014). “Equitable Top Trading Cycles Mechanism for school choice: theory and experiment.” Working paper.

Haruvy, E. and Unver, U. (2007). “Equilibrium Selection and the Role of Information in Repeated Matching Markets,” Economic Letters 94, 284–289.

Hassidim, A., Marciano-Romm, D., Schorrer, R. and Romm, A. (2015). “Strategic Players in a Strategy-Proof Environment.” Working paper.

Hugh-Jones, David, Morimitsu Kurino, and Christoph Vanberg. "An experimental study on the incentives of the probabilistic serial mechanism." Games and Economic Behavior 87 (2014): 367-380.

Kagel, J. and Roth, A. (2000): “The Dynamics of Reorganization in Matching Markets: A laboratory Experiment Motivated by A Natural Experiments,” Quarterly Journal of Economics, 115(1), 201–235.

Klijn, F., Pais, J. and Vorsatz, M. (2013). “Preference intensities and risk aversion in school choice: a laboratory experiment.” Experimental Economics 16 (1): 1–22.

Nagarkar, P. and Janis, P. (2012). “Fixing the ‘Match’: How to Play the Game.” Journal of Graduate Medical Education 4 (2): 142.

Niederle, M., Roth A. and Unver, U. (2013). “Unraveling Results from Comparable Demand and Supply,” *Games*, Special Issue: Games and Matching Markets, 4(2), 243–282.

Niederle, M. and Yariv, L. (2009). “Decentralized Matching with Aligned Preferences,” Working paper.

Pais, J. and Pintér, A. (2008). “School choice and information: An experimental study on matching mechanisms.” Games and Economic Behavior 64(1): 303–328.

Pais, J., Á. Pintér and Veszteg, R. (2011). “College Admissions and the Role of Information: An Experimental Study.” International Economic Review 52(3): 713–737.

Rees-Jones, A. (2015). “Suboptimal behavior in strategy-proof mechanisms: Evidence from the residency match.” Working paper.

Roth, A. E. (1984). “The evolution of the labor market for medical interns and residents: a case study in game theory.” The Journal of Political Economy, 991-1016.

Roth, A. E., & Peranson, E. (1999). "The redesign of the matching market for American physicians: Some engineering aspects of economic design." No. w6963. National bureau of economic research.

Vanacore, A. (2012). "Centralized enrollment in Recovery School District gets first tryout." http://www.nola.com/education/index.ssf/2012/04/centralized_enrollment_in_reco.html.

Wonderlic, E. and Hovland, C. (1939). "The Personnel Test: a restandardized abridgment of the Otis S-A test for business and industrial use." Journal of Applied Psychology 23 (6): 685–702.

Zizzo, D. J. (2010). "Experimenter demand effects in economic experiments." Experimental Economics 13: 75–98