Supplementary Appendix
Punishment strategies in repeated games: Evidence from experimental markets

Julian Wright

May 2013

1 Introduction

This supplementary appendix provides further details, results and alternative specifications for the paper "Punishment strategies in repeated games: Evidence from experimental markets". Sections 2-5 provide further details about the experimental design, including the instructions, the payoff table and screen shots. Section 6 provides further details for some results discussed in the main paper, as well as results corresponding to various additional specifications.

2 Experimental design: Full details

The experiment was conducted at the National University of Singapore (NUS) and through the Internet from August 2007 to March 2008. 219 subjects were recruited from third and fourth year undergraduate classes across NUS. The experiment was conducted in English, the common language of Singapore and the medium of instruction in Singapore's education system, including at NUS.

Students were handed a sign-up flyer after entering or exiting one of the 50 identified classes, from which they received a unique code which allowed them to enrol through the Internet for a briefing session. Enrolled subjects attended one of 13 separate briefing sessions. The briefing sessions, which lasted about 45 minutes, involved giving out instructions (which were read out aloud) as well as having subjects go through some practice rounds on computers in the lab and a short test designed to check that subjects could determine how payoffs were generated from different combinations of prices. The experiment (including the briefing session) was run through a custom-built website. Subjects took away copies of the instructions and the payoff table (see Section 3) to use for the actual experiment, which was carried out online, along with a payment of ten Singapore dollars (S\$10 $\simeq$ US\$7 at the time) in cash for attending the briefing session. Given there was excess demand for the 200 places in the experiment, subjects were selected to participate in the ongoing experiment on the basis of having given the correct answers to the questions in the short test.

The 200 selected subjects were randomly paired into 100 markets subject to the constraint that no two subjects from the same major were matched. Although it is possible subjects knew other subjects from their classes in the experiment, it seems very unlikely any of them would know the subject they were matched to, especially given a total student population of around 29,000.

All subjects received the same payoff table which showed their payoff (in lab dollars L\$) for each combination of their own price and the other subject's price. The payoffs are derived from a standard one-shot symmetric Bertrand duopoly game with differentiated products and linear demands. Specifically, inverse demand functions are given by $p_i = \alpha - \beta(q_i + \gamma q_j)$, where $0 < \gamma < 1$ serves as a measure of the degree of product substitutability. Inverting the inverse demand functions implies firm i's demand function is (for price $0 < p_i < \alpha$)

$$\begin{aligned} q_i &= \frac{\alpha(1-\gamma) - p_i + \gamma p_j}{\beta(1+\gamma)(1-\gamma)} && \text{if} && \frac{-\alpha(1-\gamma)+p_j}{\gamma} < p_i < \alpha(1-\gamma) + \gamma p_j \\ &= \frac{\alpha - p_i}{\beta} && \text{if} && 0 < p_i \leq \frac{-\alpha(1-\gamma)+p_j}{\gamma} \\ &= 0 && \text{if} && p_i \geq \alpha(1-\gamma) + \gamma p_j. \end{aligned}$$

The payoff for subject i is $\pi_i(p_i, p_j) = p_i q_i$ (to the nearest lab cent), where $\alpha = 46$, $\beta = 1$, $\gamma = 21/22$ and p_i is any whole number from 1 to 25. The one-shot Nash equilibrium price $p^n = \alpha(1-\gamma)/(2-\gamma)$ is exactly L\$2, which is also the unique Nash equilibrium price of the discretized game. The corresponding payoff is denoted π^n. The jointly optimal cooperative price (i.e. the monopoly price) $p^m = \alpha/2$ is exactly L\$23. In addition to the S\$10 for participating in the lab session, subjects received S\$1 for every 50 lab dollars (L\$) they obtained (paid into their bank account from the university after their market ended). Thus, their payoff was L\$45.02 or S\$0.90 per round at the one-shot Nash equilibrium price, L\$270.65 or S\$5.41 per round at the monopoly price, and L\$525 or S\$10.50 in any round in which they chose L\$21 while the other subject chose the monopoly price (i.e. they chose the myopic best response to the monopoly price). By way of contrast, the standard hourly wage for undergraduate research assistants at NUS was S$8.74 during the same period. Given that subjects could enter their prices at any time up to a fixed cutoff time for each round, and could do so whenever they were already logged onto the Internet, the experiment required minimal time (e.g. a minute or two to log-in and enter their price for a round). On the other hand, unlike a lab session, subjects were free to spend as long as they wanted making or changing their choice (up to the cut off time for each round, which was at 8pm on a Tuesday, Thursday and Sunday of each week).

In the instructions, subjects were told that they represent a "seller" of a product or service and that they have to decide what price to set. They were matched throughout with one other subject, known to them only as the "other seller". In the first round and most subsequent rounds, subjects had to decide which price to set (from L$1 to L$25). A screen shot from the experiment for a typical price round is given in Section 5 (the name and numbers are fictitious). Subjects can see their total payoff to date in lab dollars and can click on a link to see past choices (theirs and the other seller's) as well as their own past payoffs.

They were informed that after round 12, there was a 5% chance each round that a market would end, upon which both subjects would be paid their cumulative earnings. It was also possible that a subject did not enter their price by the cutoff time, despite reminders. In this case, as indicated in the instructions, the affected market was closed, the offending subject received no payment, and the other subject was paid in a way which neither particularly rewarded or punished him for this event, as had been stated in the briefing session. This avoids the problem, that a subject may exit the experiment in order to impose a more severe punishment on the other subject, albeit at a personal cost (which Camera and Casari, 2009, show can have important implications for cooperation).[1]

Subjects were also informed that sometimes a round may differ in that instead of entering a single price, they would have to enter a pricing strategy. This, together with the price the other subject actually set in the previous round, determined their actual price for the round. A screen shot from a typical strategy round (the name and numbers are fictitious) is given in Section 5. Unlike prices, a subject's past strategy is not observed by the other subject they are matched to (only the price that results from it is observed in subsequent rounds).

In theory, it would make no difference to a subject whether a round is a strategy round or a price setting round. The available choices and associated payoffs are essentially the same. However, asking for strategies in every round would likely cause confusion among subjects since they would never see the actual prices set by their opponent in the preceding

[1]In actual practice, the other subject was paid according to projected earnings, calculated assuming she continued to earn the market average earnings for the remaining rounds in which the market would have existed.

round.[2] Instead, strategies were only asked starting from round 10, after subjects had first experienced setting prices. Specifically, in rounds 10, 15, 20, ... each market which was still open had a 25% chance of being selected as a strategy round if the market was not going to end within the subsequent five rounds and a 75% chance of being selected as a strategy round if the market was going to end within the subsequent five rounds. To implement this approach, the length of each market was fixed prior to the beginning of the experiment based on a simulation using the 5% chance each market will end each round (after round 12) and entered into the computer program that ran the experiment over the Internet. This design ensured that subjects would not get asked for strategies too many times, especially during the earlier part of their market's existence, while at the same time increasing the chance of eliciting their strategies once their cooperative behavior had stabilized. Subjects were not informed of this particular formula; in the lab briefing they were only told that "sometimes" a round will be a strategy round. By only asking for strategies in round numbers which are multiples of five, and keeping randomness in the decision to ask for strategies, subjects would unlikely be able to work out any linkage between being asked for a strategy and the increased likelihood their market would end.[3]

The types of strategies subjects may use to sustain cooperative outcomes could be different if subjects can communicate their intentions to each other. To study the effect of communication on strategies, subjects in one quarter of the markets were informed that they could post text messages to each other through the website (one per round in each of the rounds). These 25 markets (subject-pairs) were randomly selected before the experiment started and held fixed throughout.

Since messages in round t are observed at the same time as prices from round t (i.e. when round $t+1$ prices are being set) communication is limited and there is no scope for simultaneous two-way communication. Messages were screened to ensure identities were not revealed in any way. For the remaining markets, communication between matched subjects seems highly unlikely given subjects from the same major were not matched and given subjects would have to find the one other person at NUS they were randomly matched to out of a total student population of around 29,000. Subjects were not shown the outcomes from markets other than their own.

[2]This was the main feedback from a pilot experiment that was run earlier with a group of 88 business and economics students for 20 rounds, in which one-period-ahead strategies were elicited in every round.

[3]Indeed, as I note in Section 6.2 below, I do not find any such linkage in their pricing decisions.

3 Instructions for NUS Experiment

Welcome to this briefing session. Please read the following instructions carefully together with me. Before explaining the detailed rules of the experiment, let me emphasize a few important points.

- It is important you remain quiet during this session. If you have a question, please raise your hand and somebody will answer your question discretely.
- This experiment will involve you receiving actual money. To receive this money you must complete the experiment. This means you must enter the required information in the given website by the stated deadlines. Failure to do so will likely involve removal from the experiment.
- This experiment does not involve us deceiving you in any way.
- You will be matched with one other participant throughout the experiment. Your identity will not be revealed to this or any other participant. Your responses will only be used in an anonymous way.
- These instructions and all other information you see (including the payoff table you have been given) are identical for all participants.

I will now explain the detailed rules of the experiment.

Rules

Over the course of the coming weeks and months, you will be expected to make some decisions which will earn you real money. In this experiment, you represent a firm that sells a product or a service. We will call you a "seller". You will have to decide what price to set. You will be anonymously matched with one other participant, who we will call the "other seller". The other seller you are matched with will remain the same throughout the experiment.

In the first round (and most subsequent rounds) you will have to decide a price for the round. This involves logging into a website and entering a price between 1 and 25. Your payoff will be determined by your own price and that chosen by the other seller. All prices and payoffs will be in a fictitious currency, which we will call lab dollars or L$ for short. Note 50 lab dollars equals 1 Singapore dollar.

There will be 3 such rounds per week. After 12 rounds (4 weeks), we will introduce a 5% chance each round that the experiment will end for you (as well as the other seller you are matched with). The amount of Singapore dollars we will pay you is determined by the sum of all the lab dollars you obtain in the experiment divided by 50. We will pay this to you after you have finished the experiment. An additional S$10 will be paid to you in cash after this briefing session for your participation.

If you withdraw from the experiment or fail to meet one of the deadlines, this will also constitute the end of the experiment for you (as well as the other seller you are matched with). In this case you will not receive any payment from the experiment other than the participation fee of S$10. The other affected participant will receive their earnings to date (calculated as above) plus an additional payment designed so that the other participant will neither be particularly advantaged or disadvantaged due to this event.

Payoff table

Please now look at the payoff table you have been given (we have given you two copies). This payoff table is important. It shows you how to determine your payoff for a round. You should keep this payoff table and use it throughout the experiment. It will also be loaded on the website for your convenience.

Different rows correspond to different prices you might choose. Different columns correspond to different prices the other seller might choose. The intersection of a row and column tells you how much you will earn given your price and that of the other seller. The other seller is given exactly the same payoff table as you.

We will now go through a few example rounds so you can see how the experiment will work. Keep in mind this is a practice session, so your decisions here are purely for practice and will not be used in the actual experiment in any way. Moreover, the other seller you are matched to in the actual experiment will be different from the one you are matched to today.

Example round 1 (Pricing round)

Now please go to the website *www.NUSexperiment.sg/experiment* and login using your username and password. After considering the payoff table, please choose a price for round 1 as instructed on the screen. This should be a whole number between 1 and 25 inclusive. The participant you are randomly matched to will be doing the same thing. In the actual experiment you will be free to change your mind and enter a different price up until the stated deadline.

After entering your price, please wait for further instructions (do not use the computer while you wait).

Example round 2 (Pricing round)

Round 1 is now closed. Click on "Refresh" in Internet Explorer.

On the screen you can see the price you chose in round 1, the price the other seller chose in round 1, and your earnings from round 1 in lab dollars. Check you can calculate your earnings from the two prices set by comparing with the relevant cell in the payoff table we have provided you.

Now, as instructed, enter the price you would like to set in round 2 in the given box. After entering your price, please wait until further instructions.

Example round 3 (Strategy round)

Round 2 is now closed. Click on "Refresh" in Internet Explorer.

Sometimes a round may differ in that instead of entering a single price, you enter a pricing strategy. We will call this a strategy round. This is just a different way for you to specify your price for the round. It does not differ in any other aspect.

Let me read from the screen which should appear in front of you.

This is a strategy round, which means you are not told what price the other seller set in the previous round.

The table below has 25 boxes. Above each box is a price the other seller might have set in the previous round (*i.e. in this case round 2*). In each case, enter the price you would like to set for this round. This, together with the price the other seller actually set in the previous round, will determine your actual price for this round.

For example, if the other seller had set a price of L$1 in the previous round, your actual price for this round will be the price you enter in the box corresponding to L$1. Likewise, if the other seller had set a price of L$25 in the previous round, your actual price for this round will be the price you enter in the box corresponding to L$25.

Please enter your strategy for this round now. You must fill in each of the 25 boxes with a whole number between 1 and 25. Use the tab key to shift between boxes. After entering your pricing strategy, please wait until further instructions.

Round 4 (Pricing round)

Round 3 is now closed. Click on "Refresh" in Internet Explorer.

In this example, round 4 is a pricing round, so you will observe the prices chosen by the other seller in the previous round. This is actually the price implied by his pricing strategy in round 3. For example, his price for round 3 is determined by applying his strategy to your round 2 price. Likewise, your price for round 3 is determined by applying your strategy to his round 2 price. Please check you can work out where your price in round 3 comes from.

Further rounds

We will stop here (at round 4) for this practice session. In the actual experiment, further rounds would continue in the same way. They will either be pricing rounds or strategy rounds. When you enter the website for a new round you will see whether it is a pricing or a strategy round. It should now be clear that asking for your strategy is just another way of having you specify which price you would like to set for a round.

Remember also that after 12 rounds, we will introduce a 5% chance each round that the experiment will end for you (as well as the other seller you are matched with).

We will also email you to remind you when a round is near to ending if you have not already entered the required information. Remember, failure to complete a round by the stated deadline can result in your removal from the experiment with no payment. Do not leave it to the last minute. You should email us if you ever encounter any problems so we can rectify the problem quickly.

The deadlines for logging in to enter the required information will be kept standard. Unless otherwise specified, they will be every Tuesday, Thursday and Sunday at 8pm. However, you cannot login to a new round until the last round is finalized and a new round started. Typically this will be by no latter than the next morning following the closure of each round. We will email you as soon as each round is started. (We will also give you an option to unsubscribe from this service if you wish).

Communication

Unless otherwise specified, you are not to communicate with others about the experiment. Some selected pairs of participants will be allowed to send messages to each other in each of the rounds. If you are a selected participant, then from the first round onwards, the following extra information will appear on your screen.

"Clicking on the highlighted link will allow you to post a message to the other seller. This is optional. If you do so, your message will be observed by the other seller from the start of the subsequent round. Note any message that could be used by the other seller to infer your identify (either directly or indirectly) will be banned, as will any defamatory, obscene or abusive message. We will check all messages for such information before a subsequent round is started. Otherwise, you are free to send any text based message."

As with entering prices or strategies, when sending a message you can change your entry up until the deadline for each round. In that case, only your final entry will be used for the subsequent round (and so seen by the other seller).

4 Payoff table

The other seller's price

Your price	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
1	23.02	33.77	44.51	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45
2	23.53	45.02	66.51	88	88	88	88	88	88	88	88	88	88	88	88	88	88	88	88	88	88	88	88	88	88
3	1.53	33.77	66	98.23	129	129	129	129	129	129	129	129	129	129	129	129	129	129	129	129	129	129	129	129	129
4	0	0	42.98	85.95	128.93	168	168	168	168	168	168	168	168	168	168	168	168	168	168	168	168	168	168	168	168
5	0	0	0	51.16	104.88	158.6	205	205	205	205	205	205	205	205	205	205	205	205	205	205	205	205	205	205	205
6	0	0	0	0	58.33	122.79	187.26	240	240	240	240	240	240	240	240	240	240	240	240	240	240	240	240	240	240
7	0	0	0	0	0	64.47	139.67	214.88	273	273	273	273	273	273	273	273	273	273	273	273	273	273	273	273	273
8	0	0	0	0	0	0	69.58	155.53	241.49	304	304	304	304	304	304	304	304	304	304	304	304	304	304	304	304
9	0	0	0	0	0	0	0	73.67	170.37	267.07	333	333	333	333	333	333	333	333	333	333	333	333	333	333	333
10	0	0	0	0	0	0	0	0	76.74	184.19	291.63	360	360	360	360	360	360	360	360	360	360	360	360	360	360
11	0	0	0	0	0	0	0	0	0	78.79	196.98	315.16	385	385	385	385	385	385	385	385	385	385	385	385	385
12	0	0	0	0	0	0	0	0	0	0	79.81	208.74	337.67	408	408	408	408	408	408	408	408	408	408	408	408
13	0	0	0	0	0	0	0	0	0	0	0	79.81	219.49	359.16	429	429	429	429	429	429	429	429	429	429	429
14	0	0	0	0	0	0	0	0	0	0	0	0	78.79	229.21	379.63	448	448	448	448	448	448	448	448	448	448
15	0	0	0	0	0	0	0	0	0	0	0	0	0	76.74	237.91	399.07	465	465	465	465	465	465	465	465	465
16	0	0	0	0	0	0	0	0	0	0	0	0	0	0	73.67	245.58	417.49	480	480	480	480	480	480	480	480
17	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	69.58	252.23	434.88	493	493	493	493	493	493	493
18	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	64.47	257.86	451.26	504	504	504	504	504	504
19	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	58.33	262.47	466.6	513	513	513	513	513
20	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	51.16	266.05	480.93	520	520	520	520
21	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	42.98	268.6	494.23	525	525	525
22	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	33.77	270.14	506.51	528	528
23	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	23.53	270.65	517.77	529
24	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	12.28	270.14	528
25	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	268.6

5 Screen shots from the experiment

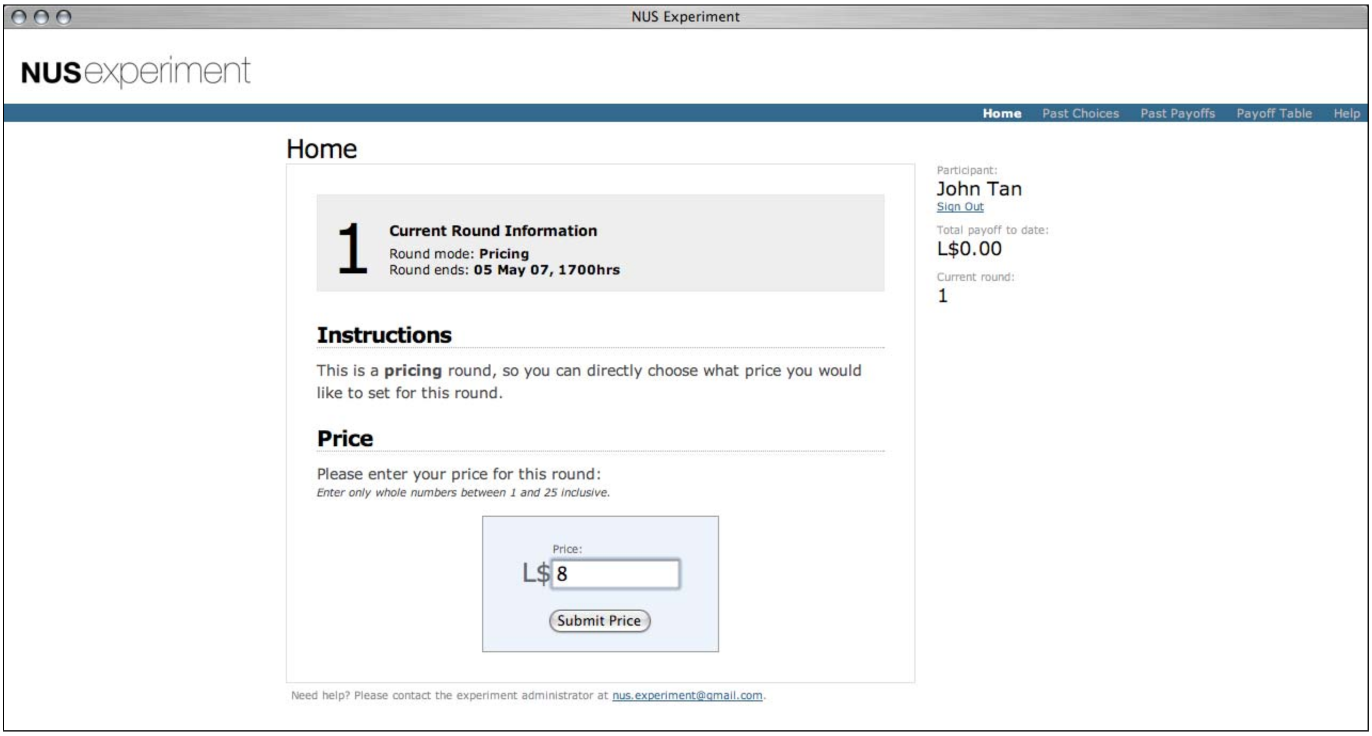

Figure: Screen shot of hypothetical price round

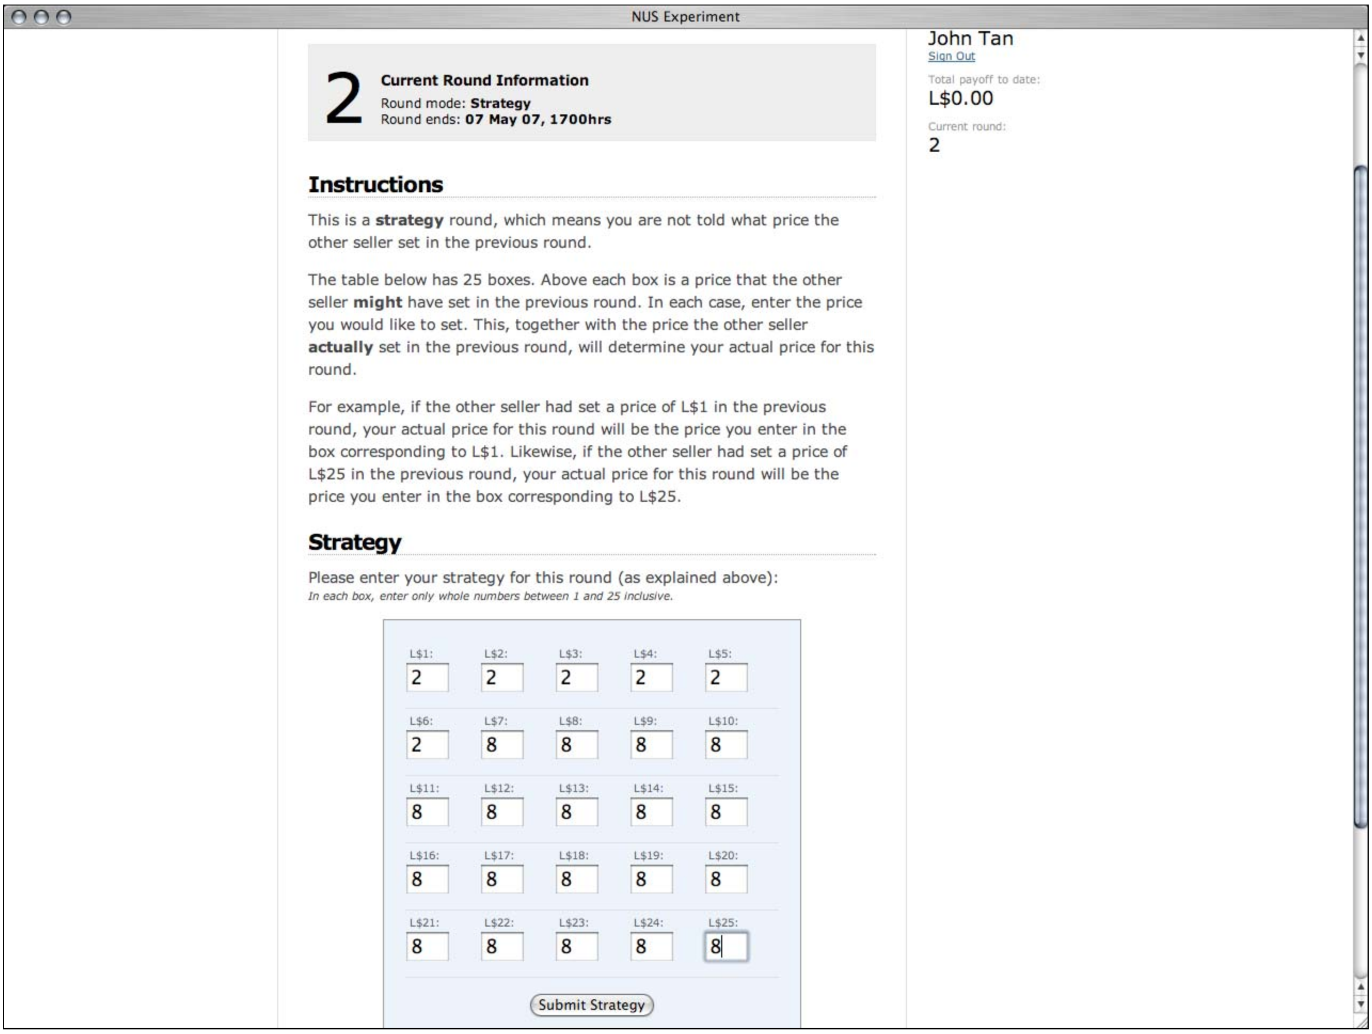

Figure: Screen shot of hypothetical strategy round

6 Additional results and specifications

In this section, I give some additional results beyond those noted in the main paper. I also give the results for alternative specifications, including those mentioned in the main paper.

6.1 Additional descriptive statistics

Subjects were well compensated for their efforts. The average earnings are S$3.30 per round. For the amount of time required to enter a price (or strategy) for a round, this compares favorably with the S$8.74 hourly wage for undergraduate research assistants at the time. There is also considerable variation in earnings across subjects. The standard deviation of earnings is S$2.82 per round. In 20.3% of rounds, subjects obtained no payoff, while in 23.5% of rounds, subjects obtained S$5.41 per round (corresponding to the monopoly payoff when both subjects set a price of L$23). In 137 cases, subjects obtained more than S$10 in a round. In total, including the S$10 for attending the lab briefing, earnings varied from S$16.10 to S$483.82 excluding the 12 subjects who chose to drop out (and who did not obtain anything beyond the initial S$10).

The average price across all 5,362 observations is L$12.83. Taking the average price for each subject first and then taking the average of this over the 200 subjects implies a mean price of L$12.50. The difference between the two measures arises because the subjects that belong to long-lasting markets contribute more observations to the total sample, leading to a higher average price given prices increase slightly over time. The average price in round 1 is L$11.815 versus L$23 from round 69 onwards, although the latter average only reflects the six subjects remaining at round 69. Calculating the change in price for each of the 176 subjects (i.e. removing markets with voluntary dropouts) in each round, the average increase in price is L$0.0591 per round when averaged over all observations. When averaged for each subject first and then averaged across the 176 subjects, weighting each subject equally, the average increase is L$0.0479 per round.

Defining stable prices as three consecutive rounds in which both subjects in the market set constant prices, 71.1% of the 1,736 stable price observations involve the monopoly price of L$23, 19.2% are less than or equal to the one-shot Nash price of L$2, and only 9.7% involve prices between the one-shot Nash price and monopoly. (Not surprisingly, there are no stable prices above the monopoly price.) In no cases did subjects choose different stable prices. Indeed, subjects in the experiment never set constant but unequal cooperative prices.

6.2 No strategy-round effect

I consider two different tests of the hypothesis that there is no strategy-round effect: (i) that individual subjects set the same prices in their strategy rounds as in their other rounds; and (ii) that subjects who are in a strategy round set the same prices as other subjects who are not in a strategy round. Since prices tend to increase in later rounds, both tests control for the number of rounds at which strategies are elicited.

To test (i), I estimate the following fixed effects model:

$$p_{it} = \alpha_i + \lambda_t + \beta I_{it} + \varepsilon_{it}, \tag{1}$$

where p_{it} is the price of subject i in round t and I_{it} is an indicator variable which equals 1 if round t is a strategy round for subject i and 0 otherwise (α_i captures subject-specific effects and λ_t captures round-specific effects). The estimate of β is -0.274 with a standard error of 0.445, so I cannot reject (for any reasonable significance level) that subjects set the same price in strategy rounds as in other rounds.[4]

To test (ii), I conduct a Wilcoxon rank-sum test (also known as Mann Whitney test) of the null hypothesis that subjects who are asked for strategies set the same prices as other subjects who are not asked for strategies in the same round. Each observation is the average price of a market in a particular round. I do this test one strategy round at a time (i.e. rounds 10, 15, 20 ...) to ensure observations used in the test are independent. I cannot reject (at the 5% significance level) that the price set in strategy rounds are the same as the price set in price rounds for any rounds for which strategies are elicited.[5] Thus, I do not find any evidence of a systematic strategy-round effect.

6.3 No last-round effects

I avoid having a known last round or implicit time limit (reflecting that subjects know the experiment cannot last beyond a certain hour) by using a fixed 5% chance of ending (after round 12) in every round for each market. This suggests there should not be any last-round effect.

Paralleling the tests of strategy-round effects, I consider two different tests of the hypothesis that there is no last-round effect: (i) that individual subjects set the same prices in their last-rounds as in their other rounds; and (ii) that subjects who are in their last round set the same prices as other subjects who are not in their last round. Since

[4]Standard errors are clustered at the subject level to account for the fact sometimes more than one strategy is elicited from the same subject.

[5]In two rounds in which strategies were elicited, towards the end of the experiment, the average price in each market was the same. Since all the observations are ties, the Wilcoxon rank-sum test cannot be formally constructed, although clearly we cannot reject that prices are equal in this case.

prices tend to increase in later rounds, both tests control for the round number in which the last round occurs.

To test (i), I estimate the following fixed effects model:

$$p_{it} = \alpha_i + \lambda_t + \beta I_{it} + \varepsilon_{it}, \tag{2}$$

where p_{it} is the price of subject i in round t and I_{it} is an indicator variable which equals 1 if round t is the last-round for a subject i and 0 otherwise (α_i captures subject-specific effects and λ_t captures round-specific effects). Excluding the 12 markets where subjects dropped out, the estimate of β is -0.079 with a standard error of 0.389, so I cannot reject (for any reasonable significance level) that subjects set the same price in their last round as in other rounds.[6]

Subjects that dropped out may have done so due to low earnings. This would be the case for subjects (or those they are matched to) that had particularly low prices just before dropping out. Including these 12 markets implies an estimate of β in (2) of -0.388 with a standard error of 0.362, so does not change the result. Estimating (2) for the 12 dropped markets separately gives an estimate of β of -1.870 with a standard error of 0.988, which is marginally significant (p-value of 0.071), indicating there may indeed be a last-round effect for these markets but with the usual causation reversed (i.e. for these markets, low prices may have caused the game to end by leading one of the subjects to drop out).

To test (ii), I conduct a Wilcoxon rank-sum test of the null hypothesis that the average price set in markets that are in their last round are the same as the average price set in markets that are not in their last round, at the same point in time. Since the samples are not matched, a rank-sum test is appropriate. Each observation is the average price of a market. I do this test one round at a time to ensure observations used in the test are independent. I cannot reject (at the 5% significance level) that the price set in markets that are about to end are the same as the price set in markets that are not about to end for 33 out of 34 rounds for which there are sufficient observations to conduct the test (after excluding the 12 markets where subjects dropped out). Note for this test I can only consider rounds which are the last round for some markets. There are two rounds where the average price in each market (those in their last round and those not in their last round) are all the same. Clearly for these two rounds, we also cannot reject the null hypothesis. Thus, I do not find any evidence of a systematic last-round effect.

[6]Standard errors are clustered at the subject level to account for the fact sometimes more than one strategy is elicited from the same subject.

6.4 Test of faculty or major effect

Although subjects are never matched to other subjects in the same major, it is possible that subjects discuss their strategy choices with classmates from the same major. From a theoretical point of view, this discussion of strategies with non-competitors should not have any effect on outcomes. Nevertheless, to see if this could be a factor in the results, I tested if there is indeed a "major effect" or "faculty effect" on pricing strategies.

Table A1: Test of faculty or major effect on H_{it}, Δ_{it}, and p_{it}

	Subjects in major/faculty with H_{it} defined	Test on (H_{it} - predicted H_{it} under Grim)	Test on (Δ_{it} - predicted Δ_{it} under Grim)	Subjects in major/faculty in round 20	Test on Price in round 20
Faculty					
Arts & Social Science	35	0.7302	0.1219	33	0.9686
Business	15	0.6509	0.5525	15	0.3271
Computing	8	0.9883	0.1402	9	0.5489
Engineering	28	0.5498	0.2554	29	0.1794
Science	43	0.9064	0.4876	32	0.8664
Major					
Bio Engineering	3	0.6165	0.8435	1	0.5139
Chemical Engineering	6	0.5565	0.9235	7	0.8223
Chemistry	8	0.2525	0.8979	5	0.7819
Civil Engineering	1	0.1967	0.2169	0	
Communications & Media	1	0.1540	0.1543	2	0.3761
Computer Engineering	4	0.2823	0.2296	5	0.9731
Computer Science	4	0.9079	0.1513	4	0.1555
Economics	19	0.3942	0.6845	14	0.2258
Electrical Engineering	5	0.6957	0.3270	4	0.6794
English	1	0.2261	0.1064	1	0.1027
Environmental Engineering	2	0.3302	0.1540	3	0.4368
Finance	9	0.2911	0.4354	7	0.6869
Food, Science & Technology	2	0.8114	0.1098	1	0.2073
Geography	1	0.8087	0.5237	2	0.1346
History	1	0.5273	0.5237	1	0.1917
Industrial & Systems	2	0.4799	0.6717	1	0.1917
Life Science	6	0.9598	0.6472	3	0.1749
Management	3	0.8942	0.1336	2	0.4481
Marketing	5	0.3789	0.7862	5	0.2770
Math	9	0.1552	0.6371	8	0.7537
Mechanical Engineering	1	0.5631	0.1064	3	0.0632
Pharmacy	7	0.8228	0.5250	7	0.4933
Physics	5	0.5575	0.5416	1	0.2073
Political Science	1	0.4843	0.5237	1	0.8704
Psychology	8	0.5839	0.1213	10	0.6960
Real Estate	1	0.3329	0.3852	1	0.6350
Sociology	1	0.1875	0.8600	0	
Statistics	1	0.1787	0.9459	2	0.1519

Notes: Reported test results are p-values from two-tailed Wilcoxon rank-sum test to test whether subjects in each particular faculty or major has the same distribution for the constructed variable for all other subjects. When a subject has multiple strategies elicited, the median value of the constructed variable is taken across the multiple strategies.

To do this I first calculate the difference between the two measures of whether strategies are harsh or disproportionate (H_{it} and Δ_{it}) and their predicted values under Grim for each case, as is done in Section 4.3 of the main paper. For each subject I take the median value of this difference. Then I use a Wilcoxon rank-sum test to test whether each particular major has the same distribution for this variable as all other majors. Note the rank-sum test is used since the subjects in a major are not matched to subjects outside of the major. This test is done for all 28 majors for which there are subjects for which H_{it} and Δ_{it} can be calculated. In the few cases where a subject put down two majors (i.e. a double major student), I assign them to both majors since they could be communicating with subjects from either. I do the same test at the faculty level for all five faculties represented in the sample (Arts and Social Science, Business, Computing, Engineering, Science). As shown in Table A1, I cannot reject that students in any particular faculty or major have the same distribution in either of the punishment measures as all other subjects in the experiment, even at the 10 percent level. Thus, I do not find any evidence that the choice of pricing strategy is linked to a subject's faculty or major.

For completeness, I also checked whether prices were significantly different for any particular faculty or major from the others. I consider this at round 20, which allows some time for the possibility of sharing information on pricing strategies, but without losing too many observations due to markets ending. Using the same non-parametric test as above, of whether subjects in one major (or faculty) have the same distribution of prices as all other subjects setting prices in the same round, I cannot reject that the distributions are the same in 25 of the 26 majors with available price data (at the 10 percent level) and in all 26 majors (at the 5 percent level). Similarly, for faculties, I cannot reject that the distribution of subjects' prices in any particular faculty is the same as all other faculties (the lowest p-value is 0.1794 which is for Engineering). The results are shown in Table A1.

6.5 Some further notes on the definition of harshness

Note that the identifying assumption in the definition of harshness H_{it} is that i's maximum intended price $\overline{s}_{it}$ is taken as the cooperative price that i intends to set. The results using H_{it} are almost identical if in the calculation of H_{it}, $\overline{s}_{it}$ is constrained to be less than or equal to 23 to take into account that the *intended* cooperative price is unlikely to ever be above the monopoly price of 23. (Recall, there are no stable prices above the monopoly price.)

Note the extent of any punishment for deviations involving j setting prices above $\overline{p}_{it}$ is not measured using our definition of H_{it}, consistent with the definition (ii) of Disproportionate punishment strategies in Section 3 of the main paper. It is not clear whether these should be considered deviations from cooperation once one moves away from a very strict interpretation of Grim and Optimal.

Recall from the main paper that for Best-response, $H_{it} \approx 0.555$. This moderate characterization of harshness makes sense even though sometimes the myopic best-response to j's deviation price leaves j with zero profit. Anticipating such a punishment, j can either match or further lower his price to ensure a relatively high payoff. This is in contrast to the case with Disproportionate.

Note that if i adopts Lenient, then $\overline{p}_{it} = 1$ and $H_{it} = 0$, reflecting that there is no immediate punishment. More generally, $\overline{p}_{it} = 1$ arises when s_{it} never increases in j's previous price p, in which case $H_{it} = 0$.

6.6 The three subjects selected out of Section 4.3 analysis

Out of the 132 subjects that have strategies elicited, H_{it}, Δ_{it} and Σ_{it} could not be defined for three subjects (subjects 56, 174 and 178), since the highest intended price of these subjects was not above the one-shot Nash price of L$2 in any of their elicited strategies. These are cases where there is no possibility of cooperation and so no sense in which a subject's punishment for deviating from cooperation can be measured. In Table A2 I give the price paths for the two markets corresponding to these three subjects.

The time series of prices in both markets appears to be consistent with subjects having difficulty coordinating on cooperative prices during the experiment. One of the markets involved a subject dropping out. None of the price paths is consistent with the idea that subjects had been cooperating until cooperation broke down due to a defection, resulting in an immediate harsh punishment with low prices.

Table A2: Price paths for three subjects selected out of Section 4.3 analysis

	Market=35		Market=44	
Round	Subject=56	Subject=174	Subject=178	Other subject
1	3	3	5	12
2	1	3	5	12
3	2	3	10	3
4	2	1	4	5
5	2	21	4	12
6	3	1	5	3
7	2	2	5	2
8	3	1	2	5
9	2	2	5	1
10	2	1	1	4
11	3	1	1	4
12	1	1	2	25
13	1	2	15	4
14	1	1	2	3
15	1	Drop out	1	23
16			5	1
17			1	2

6.7 Plots of possible disproportionate strategies

Figure A1 plots the 13 out of 204 elicited strategies that are classified as Disproportionate based on $\tau = 1$ in Section 4.4 of the main paper. Those which are also classified as Disproportionate based on $\tau = 0$ (i.e. exact fits) appear in the first row of Figure A1. These are textbook-like adoptions of trigger strategies by four different subjects. In none of these cases did the other subject in the same market also adopt a similar disproportionate punishment strategy, suggesting these are not part of an equilibrium strategy profile in which both subjects adopt trigger strategies. Only two of these subjects gave these same strategies in their first strategy elicitation. The remaining eight assigned strategies in Figure A1 can be separated into those that still exhibit features of disproportionate punishment strategies (the second row) and those that do not (the third row). The strategies corresponding to subjects 12, 45, 50 and 124 have relatively high values of Σ_{it} (number of upward steps) compared to Δ_{it} (greatest upward step), suggesting these do not actually involve disproportionate punishments. A visual inspection of these four strategies, shown in the third row of Figure A1, confirms this view.

Figure A1: Disproportional punishment strategies (assigned based on $\tau = 1$)

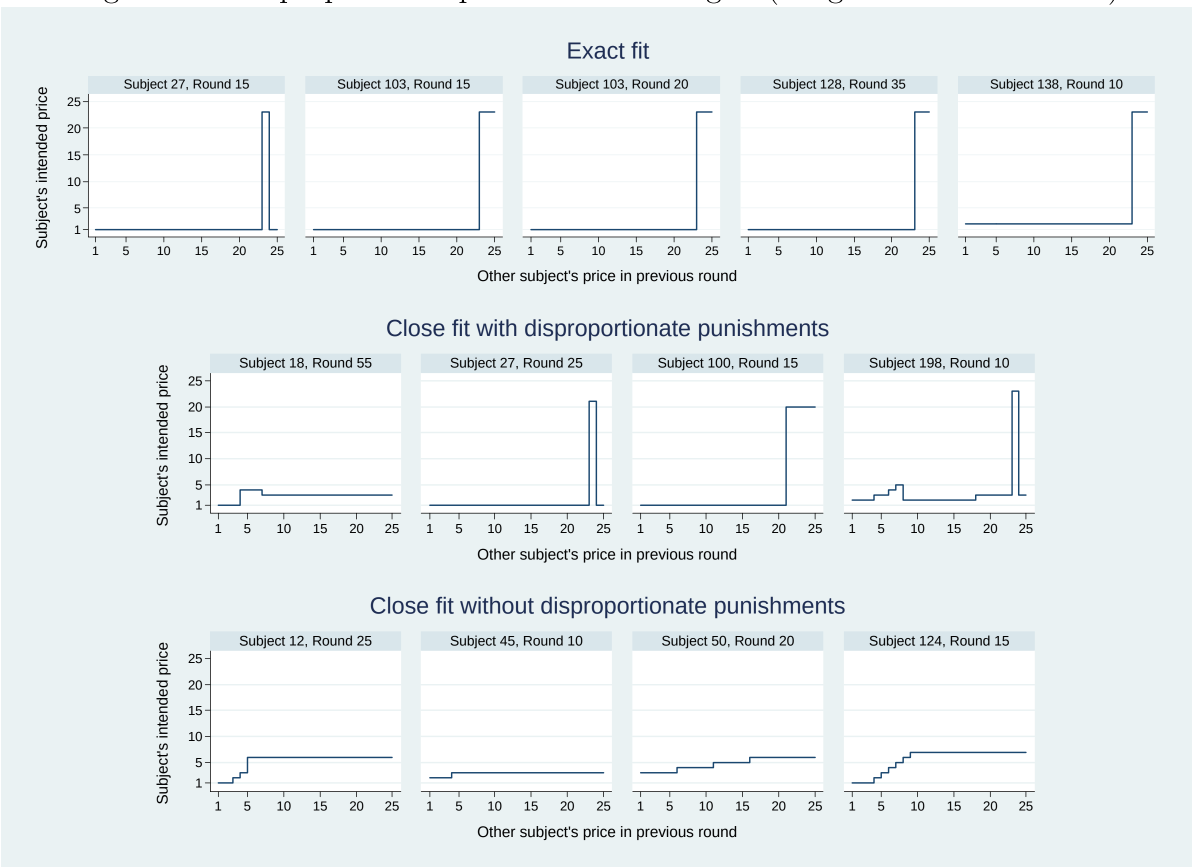

6.8 Alternatives to Table 1

In the main text, Table 1 shows the number and properties of different elicited strategies. It shows their assignment to each type of strategy k according to the two values of τ, along with the corresponding average of $RMSE_{it}^{k}$, H_{it}, Δ_{it}, and Σ_{it} across subjects for each k. It is based on the first elicited strategies for each subject. In this section, I show a similar table using the last elicited strategy for each subject (see Table A3) and using the weighted proportion assigned (see Table A4).[7] Comparing to the results in Table 1, the results are similar regardless of whether the first elicited strategies, the last elicited strategies, or all elicited strategies (appropriately weighted) are considered.

Table A3: Assignment of last elicited strategies and properties for each assignment

		Number assigned	RMSE	Harshness index H_{it}	Greatest upward step Δ_{it}	Number of upward steps Σ_{it}
Based on exact fit	($\tau = 0$)					
Disproportionate	$p^p = 1$	2	0	1	22	1
	$p^p = 2$	1	0	0.95	21	1
Tit-for-tat		8	0	0.52	1	24
Best-response		18	0	0.55	1	21
Non-cooperative		4	0	-	-	-
Lenient		18	0	0	0	0
Unassigned		81	1.95	0.54	4.78	12.51
Based on close fit	($\tau = 1$)					
Disproportionate	$p^p = 1$	4	0.21	1	20.75	1
	$p^p = 2$	5	0.59	0.73	9.20	2.60
Tit-for-tat		8	0	0.52	1	24
Best-response		45	0.34	0.56	1.49	20.96
Non-cooperative		5	0.12	0.50	1	1
Lenient		20	0.07	0.02	0.20	0.45
Unassigned		45	3.05	0.51	5.98	9.42

Notes: Based on last strategy elicited for each of the 132 subjects that had strategies elicited. Aside from column 1, all other entries are averages of the respective measure across elicited strategies of the corresponding type. The RMSE for unassigned strategies corresponds to the average of the RMSE across these strategies, calculated based on the closest fit strategy in each case.

[7] To calculate the weighted proportion assigned, for each subject I put weight $1/n$ on each strategy observation to calculate the proportions, where n is the number of strategy observations for the subject. For the 78 subjects with just one strategy observation, $n = 1$, so their strategy observations get a weight of one. If a subject has three strategy observations, for example, this implies each strategy observation is assigned a weight of 1/3. The weighted proportion assigned to a particular strategy from all 132 subjects is then just the sum of all the weights for that strategy divided by 132, which is reported in percentage terms. These weights are used to calculate weighted averages for other variables like RMSE, Harshness Index etc.

Table A4: Assignment of 204 strategies, and properties for each assignment: Weighted Proportions

		Weighted proportion	RMSE	Harshness index H_{it}	Greatest upward step Δ_{it}	Number of upward steps Σ_{it}
Based on exact fit	($\tau = 0$)					
Disproportionate	$p^p = 1$	1.39%	0	1	22	1
	$p^p = 2$	0.76%	0	0.95	21	1
Tit-for-tat		5.18%	0	0.52	1	24
Best-response		13.38%	0	0.55	1	21
Non-cooperative		3.53%	0	-	-	-
Lenient		13.45%	0	0	0	0
Unassigned		62.31%	1.90	0.53	4.28	12.86
Based on close fit	($\tau = 1$)					
Disproportionate	$p^p = 1$	2.59%	0.21	1	19.54	1
	$p^p = 2$	3.79%	0.58	0.71	9	2.90
Tit-for-tat		5.68%	0.08	0.52	1.09	23.64
Best-response		35.10%	0.33	0.56	1.45	20.93
Non-cooperative		3.91%	0.06	0.39	0.77	0.77
Lenient		14.84%	0.06	0.01	0.20	0.30
Unassigned		34.09%	3	0.50	5.48	9.42

Notes: Based on weight of $1/n$ for each strategy observation from a subject that has n strategy observations. Aside from column 1, all other entries are weighted averages of the respective measure across strategies of the corresponding type. The RMSE for unassigned strategies corresponds to the weighted average of the RMSE across these strategies, calculated based on the closest fit strategy in each case.

6.9 Evidence from price paths

According to Table 1 in the main paper, there are 45 subjects who have a first elicited strategy that is unassigned when $\tau = 1$. It is possible these unassigned strategies could arise because previously subjects had been cooperating using disproportionate punishment strategies, but cooperation had broken down leading to some punishment phase that is characterized by these unassigned strategies. Here I show there is no evidence to support that this is the reason these strategies are unassigned.

Six of these subjects had other (later) elicited strategies that were assigned and so are dealt with in the alternative specifications of Table 1 above. The remaining 39 subjects had strategies elicited but never had any of these elicited strategy assigned to one of the five types of strategies specified in Section 3 of the main paper. That is, based on the assignment of strategies using $\tau = 1$, their strategy was only classified as unassigned. Figure A2 gives the price paths for the 32 markets corresponding to these 39 subjects.

None of these price paths points towards subjects that had been cooperating until cooperation broke down due to a defection, resulting in an immediate harsh punishment with the non-cooperative prices of the subjects dropping immediately to low levels. Rather, the pattern of prices in most markets seems to be more consistent with subjects having difficulty coordinating on cooperative prices in the first place. In a few cases, the

Figure A2: Price paths for 32 markets with unassigned strategies

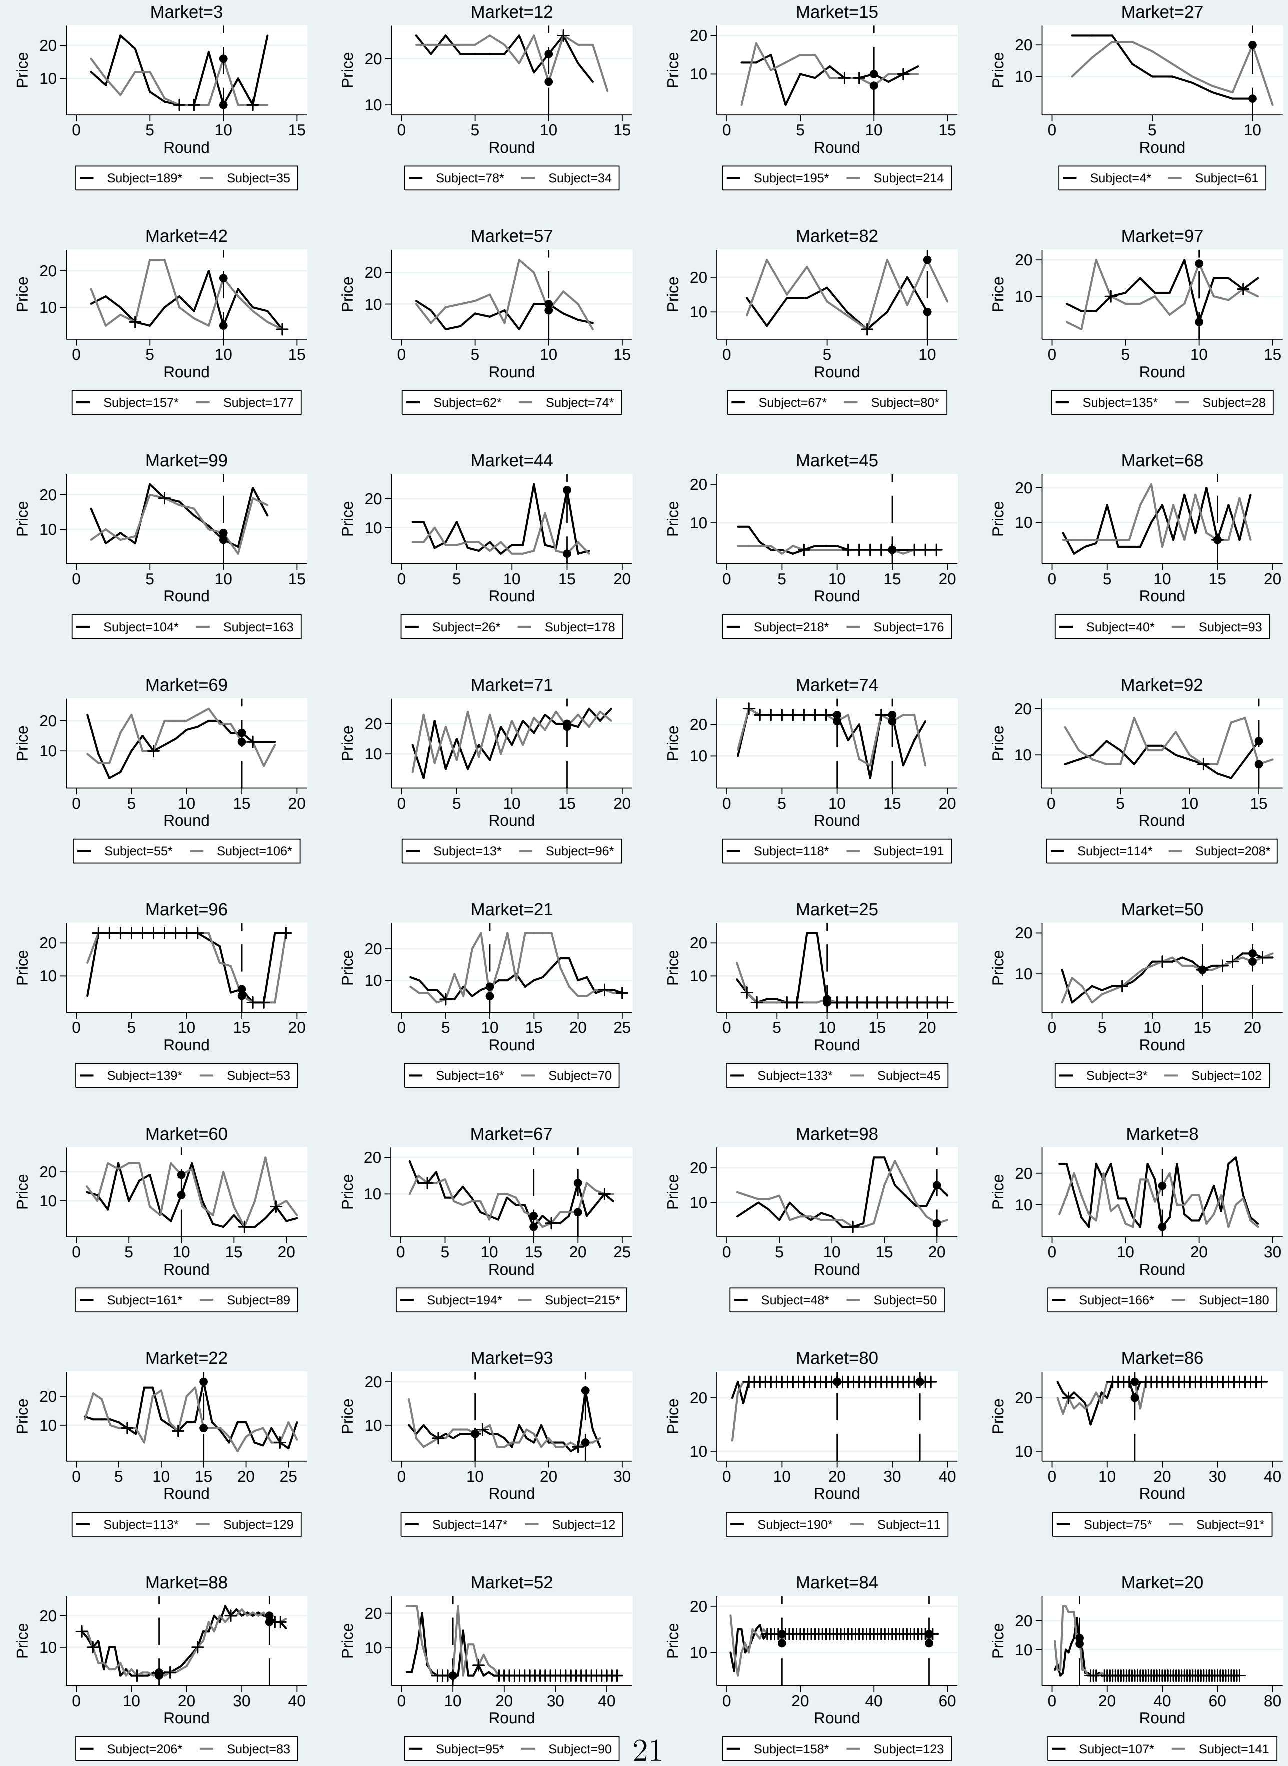

*Subjects with unassigned strategies
+Indicates common price
•Strategy rounds

pattern of prices is consistent with cooperation have been reached at the point strategies were elicited, notably markets 80, 84 and 86. In at least one case (market 96), there did seem to be cooperation in the earlier rounds but low prices at the time the strategies were elicited. In this case the breakdown of cooperation seems to have happened gradually in a tit-for-tat type fashion over a period of about five rounds. In a few cases, the prices seem to have settled at non-cooperative levels by the time strategies are elicited (e.g. markets 25 and 45) but in these cases there was no sudden drop in prices in previous periods that might indicate disproportionate punishment strategies had been used earlier.

Another possible concern arises for the 12 markets (or 24 subjects) that were affected by the 12 subjects that voluntarily dropped out. Subjects who leave voluntarily may be more likely to be those who experience low prices and profits, possibly due to the use of harsh punishments. If strategies are not elicited for these subjects, since they leave the experiment before their strategies have been elicited, then the results may understate the use of harsh punishments.

Table A5: Price paths for the four markets where one subject dropped out before strategies were elicited

	Market=24		Market=48		Market=56		Market=72	
Round	Subject dropping out	Other subject	Subject dropping out	Other subject	Subject dropping out	Other subject	Subject dropping out	Other subject
1	15	14	4	8	23	3	11	13
2	10	11	5	2	23	23	10	5
3	8	3	4	4	23	23	6	5
4	17	4	7	3	23	23	6	4
5	8	9	6	5	23	23	5	4
6	2	4	10	5	23	23	3	4
7	3	3	6	10	23	23	Drop out	2
8	3	3	5	4	23	23		
9	3	2	4	5	23	23		
10	21	2	5	3	23	23		
11	10	3	1	4	Drop out	23		
12	2	4	3	2				
13	Drop out	3	2	3				
14			2	2				
15			Drop out	5				

This potential bias is limited by the fact out of these 12 markets, seven did involve strategies being elicited at least once, so only five did not. One of these five markets would not have involved strategies being elicited anyway (based on the random draw of rounds in which strategies were scheduled to be elicited in the experiment). The four remaining markets ended (due to voluntary drop out) in rounds 7, 10, 12, and 14, so indeed relatively early in the experiment. Table A5 gives the price paths for subjects in these four markets (markets 24, 48, 56 and 72). After round 1, both subjects in market 56 priced at the monopoly price in every round suggesting the drop out was caused purely by the failure of one of the subjects to observe the deadline rather than low prices due to possibly harsh punishment strategies. Therefore, once I equally weight subjects in the strategy analysis, the potential bias in the strategy analysis caused by voluntary dropouts can at most affect 3% of the total sample. Even if subjects in all three markets actually used disproportionate punishment strategies it would not change the overall conclusions of the paper.

As can be seen from table A5, the pattern of prices does not suggest that subjects had been cooperating until cooperation broke down due to a defection, resulting in harsh punishment. Rather, the pattern of prices seems to be more consistent with subjects having difficulty coordinating on cooperative prices in the first place. In the round corresponding to each subject dropping out, the other subject set prices of 2, 3 and 5, so prices are only actually equal to the non-cooperative price in one case.

6.10 Probits to test for adoption of strategies over time

Here I report the results from 24 different probit regressions of a dummy variable for a subject's adoption of a particular strategy (based on the assignment with $\tau = 0$ or $\tau = 1$) on a constant and the corresponding number of rounds at which the subject was asked for their strategy. See Table A6.

Table A6: Estimates of how strategy choices depend on number of rounds

	Using first elicited strategy ($n = 132$)		Using last elicited strategy ($n = 132$)	
	$\tau = 0$	$\tau = 1$	$\tau = 0$	$\tau = 1$
Disproportionate	-0.035 (0.078)	-0.015 (0.034)	-0.001 (0.015)	-0.008 (0.013)
Tit-for-tat	0.040* (0.021)	0.038* (0.020)	0.019** (0.008)	0.019** (0.008)
Best-response	0.024 (0.018)	0.021 (0.016)	0.013* (0.007)	0.003 (0.007)
Non-cooperative	-0.026 (0.046)	-0.026 (0.046)	-0.037 (0.038)	-0.010 (0.016)
Lenient	-0.024 (0.026)	0.002 (0.020)	0.007 (0.007)	0.009 (0.007)
Unassigned strategy	-0.017 (0.016)	-0.047* (0.024)	-0.020*** (0.007)	-0.017** (0.008)

Notes: Each entry in the table represents the coefficient (and standard error) on the number of rounds at which subjects are asked for their strategy from a probit regression in which the dependent variable is the assignment of a certain type of strategy.
*** Significant at 1%. ** Significant at 5%. * Significant at 10%.

The probits are run on each of the five different types of strategies identified as well as unassigned strategies. The dependent variable is a 0-1 variable to indicate whether the elicited strategy takes on a particular form. For example, for the regression involving Disproportionate, the dependent variable is 0 except when the elicited strategy is assigned to be Disproportionate based on a particular value of τ (either 0 or 1). The independent variables are a constant and the number of rounds at which the strategy was elicited. The idea is that if subjects are more likely to adopt a particular strategy (e.g. Disproportionate) after playing the game for many rounds, then those strategies elicited later in the experiment would be more likely to be this strategy. A probit regression is used to pick up any such relationship.

Two different specifications are used for each strategy, either using the first elicited strategy from each subject or the last elicited strategy from each subject. For both specifications I do everything twice, once with strategies classified based on $\tau = 0$ and again with strategies classified based on $\tau = 1$. Thus, for each possible strategy classification, I run four probits. Since there are six different possible strategy classifications, this meant I ran 24 probits altogether.

The results are similar across the different approaches. The only strategy adoption that has a consistent significant relationship with the number of rounds across all specifications is Tit-for-tat, which is more likely to be used later in the experiment as indicated by the positive coefficient on the number of rounds at which the strategy is elicited.

The postive results for Tit-for-tat should be viewed with caution since they are based on observing different subjects at different points of time. Focusing only on the 54 subjects that had strategies elicited multiple times does not reveal any relationship between the strategy choice and time. Indeed 44 out of these 54 subjects chose the same strategy (based on a classification with $\tau = 0$) the last time their strategy was elicited as the first time their strategy was elicited, meaning most subjects seem to be quite stable in their strategy choice. For the remaining subjects which changed their strategies, there was no clear pattern. One interpretation of these results is that subjects that had more experience setting prices before submitting their strategies were more likely to adopt Tit-for-tat, but once they fixed a strategy it was not likely to be revised.

6.11 Alternative to Table 2

In Table 2 of the main paper, I presented the estimated effects of strategy choice on the PDV of earnings. For the 54 subjects that had multiple strategy observations, I used the following rule: Whenever a subject's strategy classification remains the same over multiple strategy rounds, only the first of these observations is used in the regression. In this Section, I report the results of the analysis when instead I use the last of these observations in the regression (Table A7). As can be seen, the results are very similar to those in the main paper. After controlling for past earnings, the only strategy that leads to significantly different earnings is Tit-for-tat, which leads to higher earnings.

Table A7: Estimated effect of last strategy choice on PDV of earnings

	Model 1		Model 2	
	$\tau = 0$	$\tau = 1$	$\tau = 0$	$\tau = 1$
Disproportionate	84.844***	65.123***	13.703	10.424
	(22.492)	(13.952)	(16.206)	(12.707)
Tit-for-tat	112.692***	88.984***	46.049**	27.411***
	(13.863)	(8.415)	(20.845)	(9.602)
Best-response	44.605***	68.075***	0.331	13.573
	(9.216)	(7.403)	(10.189)	(8.781)
Non-cooperative	29.627***	34.888***	3.003	6.515
	(8.872)	(9.573)	(5.900)	(7.480)
Lenient	69.709***	60.889***	7.279	-1.389
	(10.657)	(10.267)	(11.248)	(11.792)
Unassigned	67.559***	60.923***	14.023	5.705
	(6.014)	(8.494)	(8.835)	(10.131)
past earnings per round			16.877***	17.444***
			(2.527)	(2.418)
p-values (Wald tests of equal coefficients)				
Ho: Disproportionate = Tit-for-tat	0.294	0.146	0.108	0.097*
Ho: Disproportionate = Best-response	0.100	0.850	0.373	0.778
Ho: Disproportionate = Non-cooperative	0.024**	0.076*	0.460	0.750
Ho: Disproportionate = Lenient	0.544	0.807	0.634	0.351
Ho: Disproportionate = Unassigned	0.447	0.785	0.980	0.676
Ho: Tit-for-tat = Best-response	0.000***	0.050*	0.017**	0.070*
Ho: Tit-for-tat = Non-cooperative	0.000***	0.000***	0.027**	0.021**
Ho: Tit-for-tat = Lenient	0.015**	0.036**	0.035**	0.003***
Ho: Tit-for-tat = Unassigned	0.003***	0.016**	0.074*	0.015**
Ho: Best-response = Non-cooperative	0.244	0.007***	0.783	0.440
Ho: Best-response = Lenient	0.078*	0.574	0.532	0.165
Ho: Best-response = Unassigned	0.034**	0.494	0.152	0.400
Ho: Non-cooperative = Lenient	0.006***	0.065*	0.685	0.498
Ho: Non-cooperative = Unassigned	0.001***	0.040**	0.147	0.937
Ho: Lenient = Unassigned	0.858	0.998	0.442	0.523

Notes: Dependent variable is PDV of earnings as measured in Singapore dollars (S\$1≈US\$0.70). Whenever a subject's strategy classification remains the same over multiple strategy rounds, only the last of these observations is used in the regression. There are 151 such observations when $\tau = 0$ and 153 observations when $\tau = 1$. Robust standard errors, clustered on subjects, are reported in parentheses and used in the Wald tests.

*** Significant at 1%. ** Significant at 5%. * Significant at 10%.

6.12 Evidence of higher prices under communication

As with many earlier studies (e.g. see the survey in Haan *et al.*, 2009), I do find subjects set significantly higher prices when allowed to communicate. To show this, I calculate robust confidence intervals for the Hodges - Lehmann median difference in prices between subjects where no communication is allowed and those where communication is allowed. These confidence intervals are robust to the possibility that the population distributions in the two groups are different in ways other than their medians (see Newson, 2002). Figure A3 shows a 95% confidence interval around the difference in the median prices of the two groups by round.

There is no significant difference in the median prices between the two groups of subjects in round 1 (when no communication is possible prior to prices being set). From round 2, prices are significantly higher for subjects in the group where communication is allowed in most rounds, based on a Wilcoxon rank-sum test between the two groups, with the exception of rounds 29, 31, 32, 34, 35, 38, and 41-46. Note from round 47 onwards there are no active markets left in which subjects can communicate.

Figure A3: Prices are significantly higher with communication

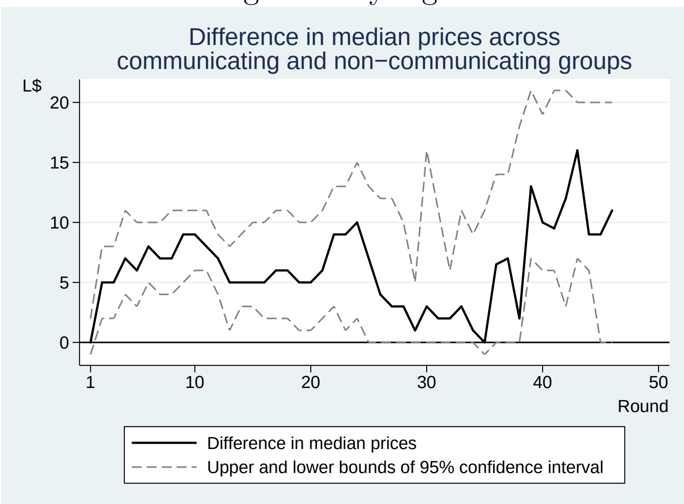

6.13 Effects of communication on strategies

Table A8 shows the coefficient estimates and standard errors from probit regressions of the adoption of each type of strategy (assigned based on $\tau = 0$ and $\tau = 1$) on a dummy variable which equals one in the communication treatment (subjects in 25 markets are allowed to send a message to each other each round) and equals zero otherwise. Each probit regression also includes a constant term. As can be seen from the results in Table A8, Lenient is significantly more likely under the communication treatment. Note that there are no observations for Non-cooperative under the communication treatment, so a probit was not conducted for Non-cooperative. This suggests that Non-cooperative is significantly less likely under the communication treatment. None of the other strategies are significantly more (or less) likely under the communication treatment.

Table A8: Estimates of how strategy choices depend on communication being allowed

	Using first elicited strategy ($n = 132$)		Using last elicited strategy ($n = 132$)	
	$\tau = 0$	$\tau = 1$	$\tau = 0$	$\tau = 1$
Disproportionate	0.500 (0.578)	0.153 (0.410)	0.833 (0.513)	0.283 (0.370)
Tit-for-tat	0.259 (0.419)	0.153 (0.410)	0.064 (0.404)	0.064 (0.404)
Best-response	-0.534 (0.382)	-0.251 (0.272)	-0.232 (0.347)	-0.403 (0.282)
Non-cooperative	-	-	-	-
Lenient	0.670** (0.299)	0.768*** (0.295)	0.510* (0.302)	0.714** (0.292)
Unassigned	-0.183 (0.263)	-0.272 (0.277)	-0.267 (0.262)	-0.147 (0.272)

Notes: Each entry in the table represents the coefficient (and standard error) on the dummy variable for whether messaging is allowed from a probit regression in which the dependent variable is the assignment of a certain type of strategy.
*** Significant at 1%. ** Significant at 5%. * Significant at 10%.

Corresponding to Table A8, Table A9 shows the coefficient estimates and standard errors from regressions of the three measures of punishment in Section 4.3 of the main paper on the dummy variable for the communication treatment. The main (robust) result in Table A9 have significantly less steps (the significant negative coefficient on Σ_{it}) consistent with the more significant use of Lenient in which there is no punishment at all.

Tables A10 and A11 repeat the above analysis but instead of the dummy variable for communication capturing whether subjects are allowed to communicate to each other, the dummy variable is defined based on whether either of the subjects in a market has sent a message to the other subject before the round in which the strategy is elicited. The results indicate actual communication has a somewhat more significant effect on strategy choice.[8] Disproportionate is significantly more likely in one out of the four specifications. Best response is significantly less likely in three out of the four specifications. Lenient is

[8]However, the results should be interpreted with caution. Using actual communication rather than the communication treatment raises possible endogeneity issues given the dummy variable on the right hand side of the regressions in Tables A10 and A11 captures the subjects' choice of whether they actually communicate (from those subjects allowed to communicate).

Table A9: Estimates of how punishment measures depend on communication being allowed

	Using first elicited strategy ($n = 132$)	Using last elicited strategy ($n = 132$)
H_{it}	-0.220	-0.023
	(0.250)	(0.254)
Δ_{it}	0.801	2.393**
	(0.954)	(1.137)
Σ_{it}	-3.452*	-4.907***
	(1.878)	(1.886)

Notes: Each entry in the table represents the coefficient (and standard error) on the dummy variable for whether messaging is allowed from a linear regression in which the dependent variable is H_{it} (the harshness index), Δ_{it} (the height of the greatest upward step of the strategy) or Σ_{it} (the number of upward steps of the strategy). For H_{it}, a generalized linear model with a logit link function and the random error from the binomial family is used to reflect that the dependent variable is a proportion between 0 and 1.
*** Significant at 1%. ** Significant at 5%. * Significant at 10%.

significantly more likely in all four specifications. The results from Table A11 are very similar to those presented in Table A9. Note the harshness measure is not significantly related to past communication, whichever way it is measured.

Table A10: Estimates of how strategy choices depend on communication occurring

	Using first elicited strategy ($n = 132$)		Using last elicited strategy ($n = 132$)	
	$\tau = 0$	$\tau = 1$	$\tau = 0$	$\tau = 1$
Disproportionate	0.538	0.199	0.876*	0.333
	(0.581)	(0.414)	(0.515)	(0.374)
Tit-for-tat	0.304	0.199	0.109	0.109
	(0.422)	(0.414)	(0.408)	(0.408)
Best-response	-0.861*	-0.431	-0.822*	-0.627**
	(0.469)	(0.286)	(0.470)	(0.304)
Non-cooperative	-	-	-	-
Lenient	0.738**	0.840***	0.575*	0.786***
	(0.303)	(0.299)	(0.306)	(0.296)
Unassigned	-0.164	-0.196	-0.139	-0.068
	(0.269)	(0.281)	(0.269)	(0.277)

Notes: Each entry in the table represents the coefficient (and standard error) on the dummy variable for whether there has been any communication in the market up till the round the strategy is elicited from a probit regression in which the dependent variable is the assignment of a certain type of strategy.
*** Significant at 1%. ** Significant at 5%. * Significant at 10%.

Table A11: Estimates of how punishment measures depend on communication occurring

	Using first elicited strategy ($n = 132$)	Using last elicited strategy ($n = 132$)
H_{it}	-0.275	-0.053
	(0.264)	(0.269)
Δ_{it}	0.937	2.762**
	(0.976)	(1.160)
Σ_{it}	-4.237**	-5.940***
	(1.912)	(1.912)

Notes: Each entry in the table represents the coefficient (and standard error) on the dummy variable for whether communication has occurred from a linear regression in which the dependent variable is H_{it} (the harshness index), Δ_{it} (the height of the greatest upward step of the strategy) or Σ_{it} (the number of upward steps of the strategy). For H_{it}, a generalized linear model with a logit link function and the random error from the binomial family is used to reflect that the dependent variable is a proportion between 0 and 1.
*** Significant at 1%. ** Significant at 5%. * Significant at 10%.

Given these findings, I explore whether the main analysis of the paper showing the distribution of key properties (the harshness of the strategy H_{it}, the height of the greatest upward step of the strategy Δ_{it}, and the number of upward steps of the strategy Σ_{it}) depends on whether communication is allowed or not.

As Figure A4 demonstrates, the nature of punishments in the subsample where communication is not allowed closely tracks the nature of punishments for the full sample used in the main paper. The subsample for which communication is allowed also has a broadly similar distribution for each of the measures of punishment as the subsample without communication. To test the distributions across the two treatments are not significantly different, I conduct a two-sample Kolmogorov-Smirnov test for the difference in the cumulative distributions. The p-values from this test are 0.634 for H_{it}, 0.512 for Δ_{it} and 0.076 for Σ_{it}, so the only possibly statistically significance difference in the distributions is for Σ_{it}. This likely reflects the greater proportion of cases with no upward steps (Lenient).

Despite the greater use of lenient strategies under the communication treatment, the overall results remain quite similar to before, in that the typical strategy is still very different from that implied by Grim. Recall that out of the 129 subjects in the full sample (a), 96.1% of them have punishment strategies that are less harsh (in their immediate response to a deviation) than implied by Grim for the same maximum intended price. In comparison, for the sample (b) where communication is not allowed, 97.0% of

Figure A4: Cumulative distributions of H_{it}, Δ_{it} and Σ_{it} for three different subsamples

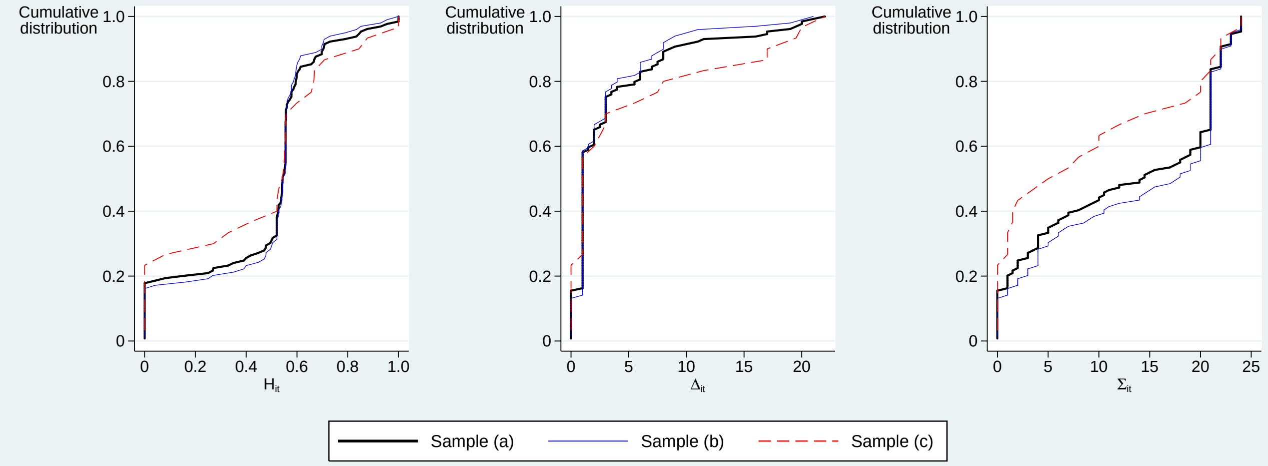

Notes: Median values of H_{it} (the harshness index), Δ_{it} (the height of the biggest step of the strategy) and Σ_{it} (the number of upward steps of the strategy) are constructed for each subject from the following subsamples:
(a) 129 subjects that had strategies elicited in which both subjects were willing to set price above p^n
(b) 99 of these subjects where communication was not allowed
(c) 30 of these subjects where communication was allowed

strategies are less harsh (in their immediate response) than implied by Grim for the same maximum intended price, while for the sample (c) where communication is allowed, 93.3% of strategies are less harsh (in their immediate response) than implied by Grim for the same maximum intended price.

To test formally whether the elicited strategies are more or less Harsh than implied by Grim, I calculate the difference between H_{it} and its predicted value under Grim. I take the median value of this by subject for the corresponding samples (a)-(c). Using a two-tailed Wilcoxon signed-rank test, I can reject the null hypothesis that this differenced variable is zero at the 1% significance level (p-value is zero to four decimal places) in all three samples (a)-(c). This is perhaps not surprising given the median difference between H_{it} and the predicted H_{it} under Grim across the three samples is -0.400 for sample (a), -0.400 for sample (b) and -0.403 for sample (c), so the typical subject has a one-period-ahead ahead strategy that is much less harsh than implied by Grim regardless of whether communication is allowed or not.

An alternative measure of whether punishment is disproportional to the deviation, is the height of the greatest upward step of the strategy Δ_{it}. To test formally whether the elicited strategies involve bigger or smaller values of Δ_{it} than implied by Grim, I calculate the difference between Δ_{it} and its predicted value under Grim. As before, I take the median value of this by subject for the corresponding samples (a)-(c). Using a two-tailed Wilcoxon signed-rank test, I can again reject the null hypothesis that this differenced variable is zero at the 1% significance level (p-value is zero to four decimal places) in all three samples (a)-(c). The median difference between Δ_{it} and the predicted

Δ_{it} under Grim is -L\$19 for subjects in sample (a), -L\$19 for subjects in sample (b) and -L\$18 for subjects in sample (c), reflecting that the typical subject has a one-period-ahead ahead strategy that is very proportional with no steps larger than L\$1 for L\$1 changes in the other subject's price regardless of whether communication is allowed or not.

The number of upward steps of the strategy of subject i in round t denoted Σ_{it} equals one for Disproportionate but is large for strategies with punishments that depend on the size of the deviation. For those strategies involving some punishment, Σ_{it} provides a measure of how independent is the punishment and so another measure of how disproportionate the punishment is. For the three samples considered in Figure A4, the proportion of subjects with a median value of Σ_{it} equal to one is 4.7% for the full sample (a), 3.0% for the subsample (b) where communication is not allowed and 10.0% for the subsample (c) where communication is allowed.

Overall the results remain robust whether or not we allow for the 25% of markets in which communication is allowed.

References

Camera, G., and M. Casari. 2009. "Cooperation among strangers under the shadow of the future," *American Economic Review*, 99(3), 979-1005.

Haan, M. A., L. Schoonbeek and B. M. Winkel (2009) "Experimental results on collusion. The role of information and communication," In J. Hinloopen and H. T. Normann (eds.), Experiments and Competition Policies, Cambridge University Press.

Newson, R. (2002) "Parameters behind "nonparametric" statistics: Kendalls tau, Somers' D and median differences," *Stata Journal* 2: 45-64.