The interaction between competition and unethical behaviour

Nick Feltovich*
Department of Economics
Monash University
Clayton VIC 3800, Australia
nicholas.feltovich@monash.edu

March 19, 2018

Abstract

We examine the interplay between unethical behaviour and competition with a lab experiment. Subjects play the role of firms in monopoly, weak competition (Bertrand-Edgeworth duopoly) or strong competition (Bertrand duopoly). Costs are determined either by a computer draw or a self-reported die roll, and pricing decisions are made with knowledge of one's own costs and – in duopoly – the rival firm's costs. Under self-reporting, lying is profitable and undetectable except statistically. We find that competition and lying are mutually reinforcing. We observe strong evidence that (behavioural) competition in both duopoly treatments is more intense when lying is possible: prices are significantly lower than when lying is impossible, even controlling for differences in costs. We also observe more lying under duopoly than monopoly – despite the greater monetary incentives to lie in the monopoly case – though these differences are not always significant.

Journal of Economic Literature classifications: K42, A13, D40, D03, C90.
Keywords: lying, dishonesty, truthfulness, experiment, monopoly, duopoly.

*Corresponding author. Financial support from Monash University is gratefully acknowledged. We thank Lana Friesen, Lata Gangadharan, Philip J. Grossman, Mehmet Y. Gürdal, Andreas Ortmann, Maroš Servátka, an associate editor, and two anonymous referees for helpful suggestions and comments.

1 Introduction

Within economics, the phenomenon of unethical behaviour is not well understood. As an example, consider lying. Only a few moments' observation or reflection is sufficient to conclude that few people are completely honest, but neither do we see the full extent of lying implied by a comparison of economic costs and benefits, and we doubt any economist would argue strongly for either of these polar cases.[1] Recent theories of unethical behaviour (and lying in particular), developed to account for evidence from the lab and the field, typically take a middle ground by treating individuals as having some kind of psychic distaste for such behaviour, but not an infinite one (see Section 2). Such a treatment often implies an interior solution for dishonesty – some, but not as much as monetary incentives alone would predict – and raises the question of what factors impact its extent.

One factor that has received some attention is the degree to which competitive pressures are present. This can matter for at least two reasons. First, competition can intensify monetary incentives for unethical behaviour. For example, for a monopolist, paying workers the legally mandated wage instead of a lower wage "off the books" might merely mean a smaller yacht, but if she faces competition from other businesses without such ethical principles, staying ethical herself might lose her all of her customers and drive her out of business. Second, even if monetary incentives are unaffected, competition can change perceptions of the morality of lying. As discussed in the next section, evidence from previous lab experiments suggests that the effect of dishonesty that explicitly affects others depends on the nature of the relationship between the protagonist and those who are affected. To greatly simplify, dishonesty appears to be viewed as more distasteful when it harms others who are perceived as friends, accomplices or innocent bystanders, but less distasteful when they are viewed as rivals, compared to dishonesty "in a vacuum" that either affects no-one at all, or is at the expense of the experimenter only.

A natural question that has received less attention in the literature is the converse: *how does unethical behaviour affect the severity of competition?*[2] There is a substantial experimental literature suggesting that individuals do not compete with one another as severely as standard game theory would imply (e.g., cooperation in the Prisoners' Dilemma and in voluntary contributions to public goods, Bertrand oligopoly).[3] If this reduced competition is due to positive other-regarding preferences (wanting others to have higher payoffs, other things held equal), and if these preferences are conditional on ethical behaviour by others, then evidence of dishonesty by others may lead to increased competition by weakening these "pro-social" preferences, possibly even turning them into "anti-social" preferences. Evidence from other settings (see Section 2) suggests this may indeed be the case. Even some other potential causes of reduced competition – such as beliefs that others will behave cooperatively – have a similar implication.

This study is an investigation of both sides of this interaction: (a) whether unethical behaviour is more prevalent when competition is more severe, and (b) whether competition is more severe when unethical behaviour is more prevalent. We use a lab experiment in which subjects play the role of price-setting firms. We have two treatment

[1] See Mazar et al. (2008) for some statistics concerning the extent of dishonest behaviour, and see Gibson et al. (2013) for examples of people behaving ethically at substantial personal costs.

[2] Note that we are using the word "competition" in two different ways: "structural" competition as determined by the numbers of buyers and sellers in the market, and "behavioural" competition as determined by how sellers' prices relate to equilibrium behaviour (in the previous and current paragraphs, respectively).

[3] Dufwenberg and Gneezy (2000) find that pricing is much less competitive in duopoly than the Bertrand prediction, though the prediction fares better as the number of oligopolists increases. See also Ledyard (1995) for examples in the public-good literature, and Holt (1995) for some from the oligopoly literature.

variables. First, we vary the severity of structural competition: monopoly, Bertrand-Edgeworth duopoly, or Bertrand duopoly. We also vary how firms' costs are determined: by die rolls made by the subject who inputs the result into the computer ("self-roll" treatment) or by computer-simulated die rolls ("computer-roll" treatment). In the computer-roll treatment, manipulation of costs is impossible. In the self-roll treatment, under-reporting the die rolls can be done with impunity (die rolls are unobservable by the experimenter and by other subjects). Importantly, the firm's costs are determined by the *reported* die roll, not the actual roll. Of course, there are few real-life settings where this literally happens: where a firm can mis-report its costs and have the lie become its true costs. But there are many examples where the effect of some form of unethical behaviour is to lower a firm's cost, and where this can be done largely out of sight: e.g., mislabelling (free range, organic, dolphin-safe, country-of-origin, etc.); rigging emissions-reporting software; adulteration of food, medicines, or cosmetics; use of sub-standard construction materials; tax avoidance/evasion; underpaying workers; and so on. So, while we will typically use "lying" as a label for the kind of unethical behaviour possible in our experimental setting, our study speaks to a broader class of misbehaviour than just lying.

Like most previous studies, we find little evidence in our experiment for either complete honesty or completely opportunistic behaviour. However, our results suggest that some apparent honesty may actually be imperfect understanding of the incentives to lie. We also find that within either of the duopoly settings, prices are lower in the self-roll treatment than in the computer-roll treatment. That is, *behavioural competition is indeed more severe when unethical behaviour is possible*. Importantly, this is not simply a consequence of lower reported costs in the self-roll treatment, but holds even when we control for costs econometrically. Finally, we find that in our self-roll treatment, reported costs are lower as structural competition increases from monopoly to Bertrand-Edgeworth duopoly and thence to Bertrand duopoly. Although these differences are not always significant, they do provide some evidence that *lying is more prevalent the more severe competitive pressures are*.

2 Literature review

Experimental studies of lying have varied greatly in their settings, such as the task used (and in particular whether the experimenter can observe lying directly, or only statistically by comparison between treatments or to known distributions); the effect of lying on others (e.g., whether lying helps or harms third parties, whether those parties are known to have behaved in good or bad ways previously); what information subjects are given about others' honesty (e.g., whether either honesty or dishonesty can be viewed as a norm); who is lied to (other subjects, the experimenter); how lying is framed (e.g., whether subjects are told explicitly they may do so); and so on.[4] Nearly all studies find substantial lying (an exception is Abeler et al., 2014), but there is often pronounced heterogeneity in its extent (Hurkens and Kartik, 2009; López-Pérez and Spiegelman, 2012; Gibson et al., 2013), and "maximal" lying (that maximises money payment) is rare.

When lying affects someone who can be viewed as an opponent or rival, behaviour typically appears to be consistent with *anti-social* preferences.[5] For example, lying occurs more in competitive settings than in comparable

[4]There are also experimental literatures examining other kinds of unethical behaviour, such as sabotage (Carpenter et al., 2010) and corruption (Serra and Wantchekon, 2012), as well as an old literature on deception in bargaining: unstructured bargaining with cheap talk (e.g., Roth and Murnighan, 1982), or structured bargaining with incomplete information, where deception involves opportunistic pooling instead of outright lying (e.g., Mitzkewitz and Nagel, 1993).

[5]By contrast, when lying affects someone who can be thought of as an innocent bystander (i.e., someone with whom there is no interaction

non-competitive settings (Schwieren and Weichselbaumer, 2010; Belot and Schröder, 2013; Rigdon and D'Esterre, 2015; Faravelli et al., 2015; Dutcher et al., 2016). However, the evidence these studies provide for anti-social preferences is weakened by a feature they share, with the exception of Rigdon and D'Esterre (2015): the own monetary benefit from a given extent of lying is on average larger under competition than under no competition. This feature arises naturally when the experiment holds constant the overall payment – due to the negative externality introduced by competition (lying has a similar effect on the total payoff to tournament competitors as it does under a piece rate, but also increases one's chance of winning the tournament at the expense of rivals) – and makes it difficult to distinguish between an intensified preference for lying due to competition per se and a straightforward response to greater monetary incentives to lie.[6] However, results from other experiments do suggest some element of anti-social preferences. For example, subjects seem to lie more when they were treated badly in a previous interaction, such as receiving a low offer in the dictator game (Houser et al., 2012). A similar result occurs in Alempaki et al.'s (2016) experiment, though they also report a positive effect: decreased lying when a high offer is received.

Similar to these last few papers – where lying depends on others' previous behaviour – is a strand of the literature looking at how behaviour is affected by others' previous honesty or dishonesty. Sánchez-Pagés and Vorsatz (2007) find that receivers expend resources to punish lying senders when this option is available, suggesting an aversion to being lied to. Peeters et al. (2008) show that some receivers *reward* truthful senders when the opportunity is available (see Anbarci et al., 2015, for a similar result). Brandts and Charness (2003) find evidence of both rewarding truth-telling and punishing lies in a 2x2 game preceded by cheap talk. Not only do individuals sometimes incur costs to punish bad behaviour directed at them, but some will do so even when the bad behaviour is directed at others; see Fehr and Fischbacher (2004), Ohtsubo et al. (2010) and Kriss et al. (2016) for examples of such "third-party" punishment. By the same token, some studies have found "third-party rewarding" of good behaviour (Seinen and Schram, 2006), and even both rewarding and punishment by third parties in the same setting (Nikiforakis and Mitchell, 2014). Researchers have also proposed other notions of disutility from bad treatment besides lying: e.g., "betrayal aversion" (Bohnet and Zeckhauser, 2004; Aimone and Houser, 2012; Aimone et al., 2013). While these studies do not closely resemble the setting we use in the current paper, their results suggest that subjects in our experiment may indeed behave more competitively when they have reason to believe they have been lied to.

A few studies are worthy of special mention due to close relevance to the current paper. The earliest studies using this procedure, to our knowledge, were those of Batson et al. (1997) and Batson et al. (1999). Their subjects were asked to use a coin toss to assign a "prize", either to oneself or to an anonymous other subject; nearly all subjects tossing the coin reported "winning the toss".[7] Fischbacher and Föllmi-Heusi (2013) seem to have been

other than the effect of the lie), an accomplice (lying is mutually beneficial, or requires coordination to be effective), or someone with a positive outside relationship (e.g., a friend or relative), behaviour is usually consistent with *pro-social* preferences. See Gneezy (2005) and Wiltermuth (2011) for examples of the first type, Weisel and Shalvi (2015), Barr and Michailidou (2016) for examples of the second type, and Houser et al. (2016) for an example of the third type.

[6] Rigdon and D'Esterre's (2015) design is similar to several of these other papers, and quite different from ours. Subjects perform a task individually for a piece rate in one treatment, and the same task in competition (against one or more other subjects) in another treatment. The major difference is that Rigdon and D'Esterre keep the same piece rate under both individualistic and competitive treatments, rather than multiplying it by the number of competitors. Thus, in contrast to other studies that maintain the expected piece rate and therefore raise the monetary incentives to lying, Rigdon and D'Esterre halve the expected piece rate, thus lowering (though not quite halving) the monetary incentives to lie.

[7] These studies used psychology experiment methods, including deception and post-session debriefing. Subjects were told that they would perform a real-effort task, with performance determining their chance of winning a raffle, and that the "prize" meant being given an intrinsically rewarding task instead of a tedious task. In fact, there was no "other subject" and no task, and all subjects were given equal

the first to use this procedure along with experimental-economics methodology. They find substantial heterogeneity in subjects' reported die rolls, with some subjects apparently completely honest, some completely opportunistic, and some in between. Like Ariely (2012), they report that changes to the monetary stakes had little effect, and importantly, implementing "double anonymity" also had little effect. Finally, a study by Cappelen et al. (2013) finds that lies are more common in market settings. Our experiment, with monopoly and duopoly treatments framed as such, therefore ought to encourage lying by subjects (though this may be less so in the monopoly cells, where subjects may view the setting as one of individual decisions rather than market interaction).

3 The experiment

Subjects in the experiment played 40 rounds, in two blocks of 20 rounds each. In the first block, all subjects played the role of monopolists who could produce 10 indivisible units in each round. The demand side of the market was automated; each simulated buyer demanded exactly one unit at any price of \$1.50 or less. In the first block, there were 10 such buyers, so market demand per round was given by

$$Q(p) = \begin{cases} 10, & p \leq 1.50 \\ 0, & p > 1.50. \end{cases}$$

In each round, a seller's per-unit cost $c \in \{0.20, 0.30, ..., 1.20\}$ was determined; after learning this cost, the subject chose price $p \in \{c, c + 0.10, ..., 1.50\}$, and profit was given by $\pi = Q(p)(p - c) = 10(p - c)$. Production was to order, and price was constrained to be no more than buyer valuation and no less than unit cost, so unsold units and negative profits were impossible.

One of our treatment variables was the way these unit costs were determined, and in particular, whether dishonesty could play a role. In our "computer-roll" treatment, c was determined by simulated die rolls: the computer would make two draws from $\{1, 2, ..., 6\}$, add them together, and multiply by \$0.10 to get c. In our "self-roll" treatment, the die rolls were made by the subject, and otherwise unobserved. The subject rolled two dice, entered the results of both rolls into the computer, and c was calculated from these self-reported rolls as in the computer-roll treatment. Thus lying about the die rolls was impossible in the computer-roll treatment, and both easy and undetectable (except perhaps statistically) in the self-roll treatment.

In the second block of 20 rounds, seller costs were determined in the same way (self- or computer-roll) as in the first block. But rather than all subjects being monopolists, we varied the severity of competition. In our "strong competition" treatment, subjects were Bertrand duopolists. In each round, they were paired, and after being informed of both own and opponent per-unit costs, they simultaneously chose prices. The 10 simulated buyers bought from the lower-priced seller, and each seller had enough capacity to serve the entire market. Our "weak competition" treatment was similar, except that we added 4 simulated buyers to make 14 in total; this Bertrand-Edgeworth duopoly is less sharply competitive than Bertrand duopoly since even the higher-priced seller is still assured of selling four units, rather than selling nothing as under strong competition. In both duopoly treatments, each seller sold to half of the buyers if their prices were equal.

Our last treatment, called "monopoly with feedback", is very similar to the setting in the first block, with subjects as monopolists facing 10 simulated buyers. The only difference from before is that each subject is paired in each

chances of winning the raffle.

round with another subject, and at the end of the round observes both their own results and those of the other subject, though the observed results have no direct payoff implications.[8]

Figure 1 shows a summary of the design, with the number of sessions and subjects in each cell. (Note that no

Figure 1: Experimental design and session information

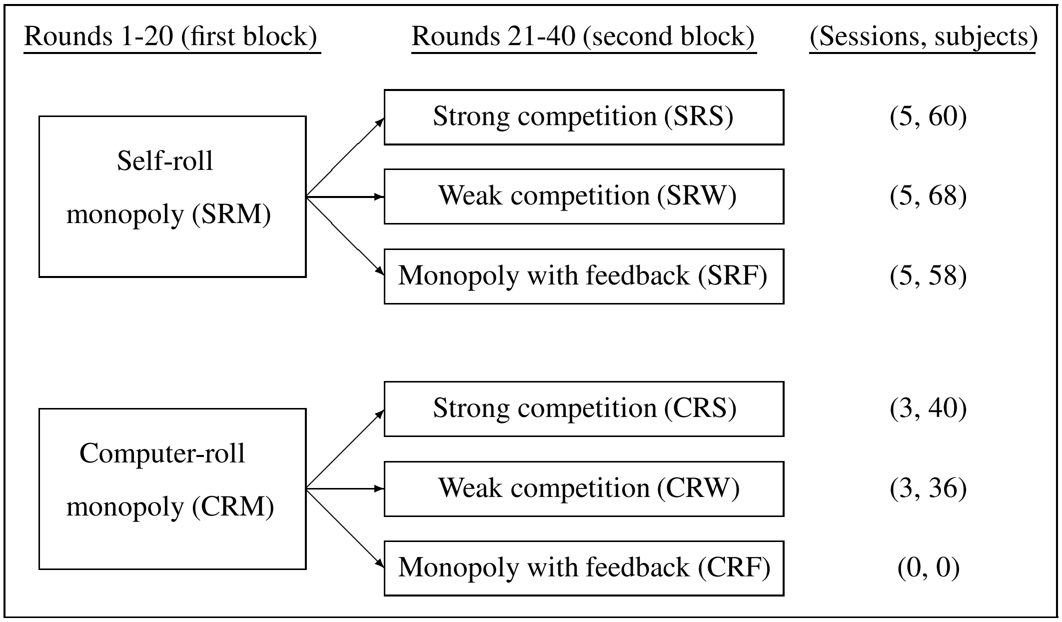

sessions were run in the "trivial cell" combining computer die rolls and monopoly with feedback, since little would be learned from this cell.) Our design allows examination of both of our main research questions. Comparison of block-2 cost-reporting decisions in the self-roll cells (SRS versus SRW versus SRF) will shed light on how lying varies with the severity of (structural) competition in the market. Comparison of second-block pricing decisions between self- and computer-roll treatments, holding the market constant (SRS versus CRS, and SRW versus CRW) will clarify how the level of (behavioural) competition changes when lying is possible relative to when it is not.

While neither of these comparisons directly involves the data from the first block, this part of the experiment is also important. The simpler setting in the first block, with no strategic interaction, provides subjects with an opportunity to gain experience in this environment before it becomes more complicated (with the introduction of strategic interaction in most cells), and in particular acquire some understanding of the monetary incentive to lie. The first block also allows us to measure subjects' intrinsic propensity to lie (the monetary incentive is transparent and no-one else's payoff is affected), as well as their understanding of the setting (due to the presence of a dominant price choice, again with no-one else's payoff affected).

[8]The reason for feedback in the second block is to allow a more direct comparison with our duopoly treatments, in which subjects not only compete, but also observe the reported cost of another subject in each round. Thus, our treatments in the second block hold constant the information subjects receive about *others'* reported costs, and hence the inferences subjects can make about others' lying. There is mixed evidence that observation of others' lying can affect honesty; see our literature review for details. See Lopez-Perez (2012) for a theoretical treatment of honesty norms, where the cost of lying is increasing in the fraction of other people observing the honesty norm, so that observing others lying can increase one's own propensity to lie via changes in beliefs regarding how much the norm is followed.

3.1 Experimental procedures

The experiment was conducted in MonLEE, the experimental lab at Monash University; subjects were recruited by ORSEE (Greiner, 2015) from the lab's subject pool, with no additional exclusion criteria. The experiment was computerised, and programmed using the z-Tree experiment software package (Fischbacher, 2007). All interaction took place anonymously via the computer program; subjects were visually isolated and received no identifying information about other subjects. The number of subjects in a session was kept low (10, 12 or 14, in a lab capable of accommodating 26) to provide an added measure of visual isolation; this was especially important in our design in order to ensure that subjects in the self-roll treatment could roll their dice in such a way that neither the experimenter nor any of the other subjects could see the result.

Sample instructions and screen-shots from the experiment are provided in the appendix. Subjects received instructions for the first block (and in the self-roll treatment, two dice) at the beginning of the session. Instructions for the second block were given after the first block had ended, so that no specific information about the second block was known during the first block. Both sets of instructions were given in writing and also orally (in an attempt to make the environment common knowledge), and subjects were prompted for questions before each block began. (They were also told they could ask questions after the block had begun, but this never happened.) There was no instructions quiz, nor were there any practice rounds.

Neutral phrasing was used in the instructions to the extent possible. Regarding die rolling in the self-roll treatment, the wording was "A round begins with you finding out your cost of production. You will roll two dice and enter the results into the computer. *The sum of these dice, multiplied by $0.10, will be your per-unit cost of production.*" [Emphasis in original.] Thus subjects were not told explicitly either that lying was allowed or that it was discouraged. Also, in duopoly cells, the rival firm was always referred to with wording like "the other seller" rather than "rival", "competitor" or "opponent" (or, for that matter, "partner"). This use of neutral phrasing, along with the relative complexity of the game being played (compared to other studies of lying discussed in Section 2), were intended to reduce demand effects, which might be extremely salient in a study of dishonesty.

Subjects were randomly matched in the second block using a one-population protocol, with no partitioning of individual sessions into smaller matching groups, so that in a given round, a subject was equally likely to be paired with any of the other subjects in the session. This was done in order to maximise the number of potential opponents a subject could face, and thus minimise incentives for repeated-game behaviour. For the same reason, no persistent information about the opponent (e.g., subject ID numbers) was given.

A round began with subjects being prompted to roll their dice and enter the results into the computer (in the self-roll treatment) or with a reminder that the computer would be determining their cost (in the computer-roll treatment). In the former, the die rolls were entered via two sets of radio buttons, so reports other than 1, 2, 3, 4, 5 or 6 were not possible. After costs were determined, subjects were prompted to choose their price. In the duopoly treatments, they were informed not only of their own cost, but also their opponent's cost, prior to this price choice. After all subjects had chosen prices, the round ended and feedback was given. In the first block, feedback consisted only of the subject's own results: cost, price, quantity and profit. In the second block, feedback consisted of this information and also the corresponding results of the paired subject.

After the last round, subjects completed a short questionnaire with some demographic questions.[9] Then, subjects

[9] See the appendix for sample instructions and screen-shots, including those for the questionnaire. Other experimental materials including the raw data are available from the corresponding author upon request.

were paid (exactly) the sum of their profits from four randomly-chosen rounds in Australian dollars (AUD), plus AUD 10 for completing the questionnaire.[10] Total earnings averaged about $36 for a session that typically lasted 45-60 minutes.

3.2 Monetary incentives

In this section, we present the monetary incentives in our various experimental cells, as well as the predictions arising from standard preferences (own-monetary-payment maximisation). The analysis here will also be useful in formulating our predictions under alternative preferences (see Section 3.3).

Monopoly cells

In the CRM, SRM and SRF cells, for any given cost, choosing a price of $1.50 is dominant, with each 10-cent decrease in price reducing one's own profit by one dollar. In SRM and SRF, money-payment maximisation also implies maximally under-reporting cost (c = $0.20) with each additional 10 cents in cost (i.e., one pip of the dice) reducing profit by one dollar. Hence, a subject whose utility depends only on own monetary payment will set a cost of $0.20 and a price of $1.50, yielding a profit of $13 compared to $8 (on average) from honest reporting of fair dice.

Computer-roll duopoly cells

In the CRW and CRS cells, each pair of realised costs determines a simultaneous-move pricing continuation game played between the two duopolists. Under weak competition, the best response to a high rival price is to under-cut that price by ten cents to sell all 10 units, while the best response to a low rival price is the monopoly price (selling only 4 units but at a high mark-up). As a result, for most cost pairs there is no pure-strategy equilibrium in prices; rather, sellers mix over a subset of the interval whose endpoints are the monopoly price and the highest price for which monopoly pricing is the best response (this is the support of the *Edgeworth (1925) cycle*). We will not go into the details of the equilibria of the various price-setting continuation games (these are available from the author upon request), but the resulting payoffs for the game with weak competition can be found in Figure 2. Since the game is symmetric, only row player payoffs are shown.[11]

Under strong competition, as is typical in Bertrand duopoly, the best response to a given rival price in the price-setting continuation game is usually to minimally under-cut that price. Equilibrium therefore requires pricing at or near the higher-cost seller's marginal cost. Again, we leave out the details of the equilibrium computations, and present the resulting symmetric game in Figure 3.

In the computer-roll treatment, the only subject decision is the price choice. Under standard preferences, we would expect to see equilibrium behaviour in the price-setting games. Other patterns of behaviour may be reasonable here, however, since mutual price increases benefit both sellers. Purely strategic incentives for collusive

[10] At the time of the experiment, one AUD was worth roughly 0.70 USD.

[11] There are a few cases (when sellers' costs are equal, at 0.20, 0.30, 0.80, 0.90, 1.00, 1.10 or 1.20) where multiple equilibria exist. In these cases, the payoffs in Figure 2 are calculated by taking the simple average of each equilibrium payoff; however, the equilibria are "nearly payoff equivalent" in the sense that order relationships across outcomes within any column are unaffected by choice of any particular one of the multiple equilibria. For example, if Firm 2's cost is 0.30, then Firm 1's payoff given a cost of 0.30 ($5.18), and based on the figure, is less than that given a cost of 0.20 but higher than that given any higher cost. The two equilibria when both costs are 0.30 give Firm 1 payoffs of $5.14 and $5.22 respectively, so using either individual equilibrium payoff instead of their average does not affect any of those order relationships. The same is true for Figure 3 below.

Figure 2: Firm 1 payoffs under *weak competition*, assuming equilibrium in continuation

		Firm 2 per-unit cost										
		0.20	0.30	0.40	0.50	0.60	0.70	0.80	0.90	1.00	1.10	1.20
	0.20	5.40	5.84	6.39	7.00	7.73	8.08	8.84	9.17	10.00	10.20	11.00
	0.30	4.80	5.18	5.42	6.00	6.72	7.10	7.83	8.18	9.00	9.17	10.00
	0.40	4.40	4.40	4.57	5.00	5.70	6.12	6.81	7.20	8.00	8.18	9.00
	0.50	4.00	4.00	4.00	4.09	4.68	5.14	5.79	6.22	7.00	7.20	8.00
Firm 1	0.60	3.60	3.60	3.60	3.60	3.91	4.18	4.77	5.25	6.00	6.22	7.00
per-unit	0.70	3.20	3.20	3.20	3.20	3.20	3.52	3.75	4.29	5.00	5.25	6.00
cost	0.80	2.80	2.80	2.80	2.80	2.80	2.80	2.95	3.33	4.00	4.29	5.00
	0.90	2.40	2.40	2.40	2.40	2.40	2.40	2.40	2.61	3.00	3.33	4.00
	1.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.27	2.47	3.00
	1.10	1.60	1.60	1.60	1.60	1.60	1.60	1.60	1.60	1.68	1.86	2.10
	1.20	1.20	1.20	1.20	1.20	1.20	1.20	1.20	1.20	1.20	1.40	1.77

Figure 3: Firm 1 payoffs under *strong competition*, assuming equilibrium in continuation

		Firm 2 per-unit cost										
		0.20	0.30	0.40	0.50	0.60	0.70	0.80	0.90	1.00	1.10	1.20
	0.20	0.50	0.67	1.17	2.50	3.50	4.50	5.50	6.50	7.50	8.50	9.50
	0.30	0.17	0.50	0.67	1.17	2.50	3.50	4.50	5.50	6.50	7.50	8.50
	0.40	0.17	0.17	0.50	0.67	1.17	2.50	3.50	4.50	5.50	6.50	7.50
	0.50	0.00	0.17	0.17	0.50	0.67	1.17	2.50	3.50	4.50	5.50	6.50
Firm 1	0.60	0.00	0.00	0.17	0.17	0.50	0.67	1.17	2.50	3.50	4.50	5.50
per-unit	0.70	0.00	0.00	0.00	0.17	0.17	0.50	0.67	1.17	2.50	3.50	4.50
cost	0.80	0.00	0.00	0.00	0.00	0.17	0.17	0.50	0.67	1.17	2.50	3.50
	0.90	0.00	0.00	0.00	0.00	0.00	0.17	0.17	0.50	0.67	1.17	2.50
	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.17	0.17	0.50	0.67	1.17
	1.10	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.17	0.17	0.50	0.67
	1.20	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.17	0.17	0.50

behaviour are weak, since our experiment uses random pairings and we publicly announce the number of rounds at the beginning of each block (making the setting one with finite rather than indefinite repetition). However, these incentives may still exist, and would be reinforced by even small pro-social preferences (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000) or efficiency seeking (Engelmann and Strobel, 2004). In this case, any systematic deviations from equilibrium price-setting would be in the direction implied by collusion: *higher* prices. On the other hand, recent work (Zizzo and Oswald, 2001; Abbink and Sadrieh, 2009) suggests that in some circumstances, people have anti-social preferences; these would imply *lower* prices than in equilibrium.

Self-roll duopoly cells

Predictions for the SRW and SRS cells must specify both price-setting and cost-reporting behaviour. Under standard preferences, the former is identical (for given costs) to the corresponding computer-roll cells. So, Figures 2 and 3 apply to the self-roll treatment as they did for the computer-roll treatment, but now they serve as the strategic forms for the cost-reporting decision in the SRW and SRS cells given standard preferences and equilibrium continuation behaviour. Clearly, reporting a die roll of 2 is strictly dominant in both SRW and SRS cells, just as in the SRM and SRF cells.

Under alternative preferences, the actual monetary impacts of lying on own and opponent payoffs may be important, so we provide a brief description of these impacts in Table 1, which combines the information in Figures 2 and 3 with some specific assumptions about how players report their costs. To keep the table a manageable size, we focus on a few simple own and rival cost-reporting strategies: (1) complete honesty (always reporting the actual die rolls), (2) under-stating by one pip, (3) under-stating by two pips and (4) maximal lying (always reporting a roll of two).

Table 1: Expected payoffs ($) of selected cost reporting strategies in self-roll-weak (SRW) and -strong (SRS) cells

Own cost	Rival cost reporting (SRW cell)				Rival cost reporting (SRS cell)			
reporting	Truthful	$<$ by 10¢	$<$ by 20¢	Minimum	Truthful	$<$ by 10¢	$<$ by 20¢	Minimum
Report truthfully	4.02	3.77	3.57	3.21	1.23	0.84	0.55	0.04
Under-state by 0.10	4.71	4.40	4.15	3.61	1.69	1.20	0.82	0.07
Under-state by 0.20	5.42	5.05	4.74	3.99	2.20	1.61	1.14	0.12
Report the minimum	8.16	7.61	7.10	5.40	4.51	3.59	2.75	0.50

In the table, own-payoff effects from changing one's cost-reporting strategy can be seen by moving up and down, while – due to symmetry – opponent-payoff effects can be seen by moving left and right. For example, if both players report truthfully under weak competition, each gets an expected payoff of $4.02 (assuming equilibrium continuation play). If one of the players decides instead to under-state cost by ten cents, his payoff rises to $4.71 and his opponent's falls to $3.77; that is, he gets 69 cents more and the opponent gets 35 cents less.

Recall that in the monopoly cells, under-stating cost by ten cents increases one's own profit by one dollar, with no effect on others' profits. In *absolute* terms (i.e., dollar gains and losses), under-stating costs helps oneself less and harms the rival more in either type of duopoly than in monopoly, and less in strong duopoly than in weak duopoly. However, in *relative* terms (percent gains and losses), the opposite is true: the effect on both own and opponent payoff is higher under strong than weak duopoly, and higher under weak duopoly than monopoly.

3.3 Behavioural hypotheses

In this section, we continue working from the analysis above, but we relax the assumption of standard preferences in two ways. First, lying is not costless: there is a cost of mis-reporting the die roll that increases (at an increasing rate) in the size of the mis-report. (We will typically refer to this as an "aversion" rather than a cost, to avoid confusion with the firms' monetary costs of production.) Second, preferences over money payments can be pro- or anti-social; that is, while individuals prefer more money to less for themselves, they might prefer more money to less for the opponent (in the case of pro-social preferences) or the reverse (for anti-social preferences), other things equal.

While we will not attempt a rigorous analysis here, for concreteness we display a sample utility function in (1).

$$U_1 = \pi_1 - \alpha(\pi_1 - \pi_0) - \beta\phi(c_1 - \hat{c}_1)^2, \tag{1}$$

Here, U_1 is utility, π_1 is the individual's own monetary payoff, π_0 is the rival's monetary payoff, c_1 is the individual's reported cost, and $\hat{c}_1$ is the "true" cost (according to the die rolls). The parameter α is an individual-specific measure of other-regarding preference: it is zero under standard selfish preferences, positive under pro-social preferences, or negative under anti-social preferences. The parameter $\beta \geq 0$ measures the individual's aversion to lying. The parameter ϕ is a game-specific scaling parameter meant to make the last term commensurable with the others; for example, it could be set to the game's expected monetary payment under truth-telling, or equivalently the difference between this expectation and the minimum possible monetary payment (see Table 1).

Monopoly behaviour

In the monopoly case, there is no rival, so other-regarding preferences are irrelevant, and subjects should still choose the dominant price of \$1.50. Reported costs (in SRM and SRF) are given by a simple marginal analysis exercise. The marginal benefit of under-reporting begins at just under one dollar for the first pip and decreases from there, while the marginal cost is increasing.[12] So, there are three potential solutions: honest reporting (for sufficiently high aversion to lying), maximal lying (for sufficiently low aversion), or an interior solution (for intermediate levels).

The effect of competition on lying

As in the monopoly case, the optimal extent of lying emerges from marginal analysis. It is here that the effects of lying on own and opponent money payoff, discussed in Section 3.2, become important. Recall that a strategy of under-stating die rolls by one pip increases own payoff by (nearly) one dollar in monopoly, a smaller amount in weak competition, and an even smaller amount in strong competition, and it decreases the opponent payoff by zero in monopoly, a small amount in weak competition, and a larger amount in strong competition – and that the total of the increase in own payoff and decrease in opponent payoff is typically a dollar or less. So, under the simplest assumptions of a cost of lying that depends only on the absolute size of the lie (i.e., increasing in the difference between reported and actual die rolls), and a benefit of lying equal to the monetary gain, we would expect to see less lying as competition increases under most kinds of preferences (i.e., all but anti-social preferences so extreme that one would sacrifice more than a dollar of own payoff to reduce the opponent's payoff by a dollar).

However, there are reasons to expect results in the opposite direction. First, competition itself could increase the benefits or lower the costs of lying. Intuitively, this seems plausible (consider, for example, the old saying "all is fair in love and war", with the implication that behaviour that is usually unacceptable becomes acceptable when competition is heightened). As noted in Section 2, some evidence from prior studies supports this possibility, with more lying in competitive settings compared to similar non-competitive settings. While (as also noted in that section) most of these studies also increased monetary incentives in the competitive setting, at least one did not, yielding a clean result for the effect of competition. Even for those studies that did increase incentives under competition, the strength of some of their results suggests that they are not due solely to the change in incentives. (Cappelen et al.'s (2013) finding of more lying in market settings is also relevant here.)

[12]The increasing marginal cost is by assumption. The decreasing marginal benefit is due to the increasing likelihood of the minimum-cost constraint of 20 cents becoming binding as under-stating becomes more severe.

A second potential reason for more lying in our more competitive settings is that either the benefit or cost of lying may involve a comparison to some measure of the overall stakes in the setting (as in the sample utility function (1) above). For example, a lie that increases money earnings by one dollar may be viewed by the decision maker as innocuous when thousands of dollars can be gained or lost, but as severe when playing for a one-dollar prize. We are not aware of any tests of relative-versus-absolute considerations in this particular setting, but it is similar in spirit to findings from other parts of the literature. A well-known example is Thaler's (1980) suggestion that search behaviour depends on relative as well as absolute price savings (when asked, people state that they would be more willing to spend time and effort travelling to pay \$20 rather than \$25 on a clock-radio than to pay \$495 rather than \$500 on a television). As noted at the end of Section 3.2, the benefit of lying in *relative* terms is highest under strong competition and lowest under monopoly.

Based on these reasons, we formulate our first hypothesis.

Hypothesis 1 *In the self-roll treatment, reported costs will be lower under strong duopoly than under weak duopoly, and lower under weak duopoly than under monopoly.*

The effect of lying on competition

As noted in Section 3.2, in duopoly the pair of costs – either randomly drawn in the computer-roll treatment or reported in the self-roll treatment – determines a pricing continuation subgame. Under selfish preferences, the resulting payoffs are as shown in Figures 2 and 3. As also noted there, pro-social preferences result in higher equilibrium prices, while anti-social preferences result in lower equilibrium prices.

To this last fact, we add a plausible conjecture about lying:

Conjecture 1 *For given reported costs, believing others are lying will make preferences less pro-social or more anti-social.*[13]

In a general sense, observed lying leading to a decline in other-regardingness is probably uncontroversial. However, there are several mechanisms by which this can happen, and we are implicitly ruling some of these out. One of these is simply direct reciprocity: inferring that one's current opponent has lied could make one less pro-social toward that person, without necessarily affecting one's attitudes toward others (Sánchez-Pagés and Vorsatz, 2007; Dufwenberg et al., 2017). Relatedly, inferring that the opponent was truthful could make one more pro-social (Peeters et al., 2008), or a combination of the two could occur (Brandts and Charness, 2003). A mechanism along these lines would imply that the effect on preferences should be correlated with the current opponent's reported cost (which is presumably more likely to be a lie the lower it is), though the strength of this correlation and the sign of the overall effect would depend on beliefs and on the pro- and anti-social effects of honesty and lying respectively.

Instead, the mechanism we propose is broader in both the source and the outcome. We conjecture that it is simply the knowledge that lying is likely to happen – rather than inferring it has happened in a particular case – that affects preferences. We also conjecture that the decrease in other-regardingness is not directed narrowly at those opponents believed to have lied, but rather at the group as a whole. That is, belief that the norm of honesty is likely to be violated weakens another norm: that of cooperation or generosity. The idea that weakening one norm can lead to the weakening of others was popularised as the "Broken Windows Theory" of civic disorder (Kelling and Wilson,

[13] In (1), this would be represented by α decreasing.

1982), and while its empirical validity in the field has been questioned, several studies in the lab have found support for it (e.g., Keizer et al., 2008, 2013; Keuschnigg and Wolbring, 2015; Berger and Hevenstone, 2016; Gächter and Schulz, 2016). A similar flavour is found in theoretical models, such as those of Levine (1998) and Ellingsen and Johannessen (2008), that propose individuals' pro-sociality is influenced by their beliefs about how pro-social their current opponent is. While these beliefs might come from past experience with that particular opponent, they could come from limited knowledge such as how the population on average behaves, or even just what kinds of behaviour are possible.

Since lying cannot occur in the computer-roll treatment but can (and almost certainly will) in the self-roll treatment, Conjecture 1 and the properties of the payoff matrices in Figures 2 and 3 yield our second hypothesis.

Hypothesis 2 *For given market type (strong or weak duopoly) and given reported costs, price choices will be lower in the self-roll treatment than in the computer-roll treatment.*

Note that the hypothesised effect on prices in Hypothesis 2 is *conditional on the costs*. Since under-reporting of costs is very likely in the self-roll treatment and impossible in the computer-roll treatment, costs are likely to be systematically lower in the self-roll treatment, which on its own would imply lower prices in that treatment. In Hypothesis 2, we are making a stronger prediction: prices are lower in the self-roll treatment even after controlling for the difference in reported costs.[14]

Before continuing, it is worth noting that while we consider our Conjecture 1 to be reasonable, there are other alternatives for which arguments could be made. One was mentioned above: the effect on preferences depending on the (perceived) honesty of the current opponent. As noted there, the overall effect on preferences is indeterminate in this case, but an alternative testable prediction would be a positive *interaction* between the self-roll treatment and the opponent reported cost. That is, prices ought to decrease by more, or increase by less, moving from the computer- to the self-roll treatment the lower is the opponent reported cost. A second alternative conjecture is that rather than preferences becoming less pro-social in the self-roll treatment due to the possibility that others are lying, they might become *more* pro-social due to the possibility that at least some others are telling the truth even when they are not constrained to do so (whereas in the computer-roll treatment, they are trivially honest because they have no opportunity to lie). We stay with Conjecture 1, and hence Hypothesis 2, for two reasons. First, the random-matching protocol used in our experiment makes it difficult for subjects to identify individual liars in the self-roll treatment, while over time they are likely to become aware that *some* others are lying. Second, evolutionary arguments suggest that even if both rewarding truth-telling and punishing lying exist in the population, the latter is likely to be stronger. However, we will revisit these alternative conjectures in our discussion of the experimental results (see Section 4.3).

4 Results

We conducted 21 experimental sessions with a total of 262 subjects (see Figure 1).[15]

[14] An alternative experimental procedure, which would have eliminated the need for such controls, would have the self-roll treatment conducted first, then used the observed distribution of *reported* die rolls for the computer-roll treatment instead of fair die rolls.

[15] Besides the sessions shown in the figure, there were two pilot sessions, conducted before the first "real" session, to test the program and ensure the instructions were understandable. Also, a real session had to be cancelled while in progress due to a hardware fault. The data from the pilots and the cancelled session are not included in any of our analysis, and we do not discuss them further.

4.1 Block 1 behaviour

Figure 4 shows time series for subjects' per-unit costs and price choices in block 1. Average block-1 costs were \$0.698 and \$0.547 in the computer-roll and self-roll treatments respectively. The former is almost exactly equal to the expected value associated with two fair dice, while the probability of \$0.547 or lower arising from truth-telling is vanishingly small.[16] A non-parametric robust rank-order test, using the 20-round average cost for each individual subject as the unit of observation, indicates that reported costs in the self-roll treatment are significantly lower than costs in the computer-roll treatment ($\grave{U} = 13.45$, two-tailed $q < 0.001$).[17] Hence we conclude that at least some subjects at some times are lying about their die rolls.[18] On the other hand, the right panel shows similar average price

Figure 4: Block 1 time series, cost and price

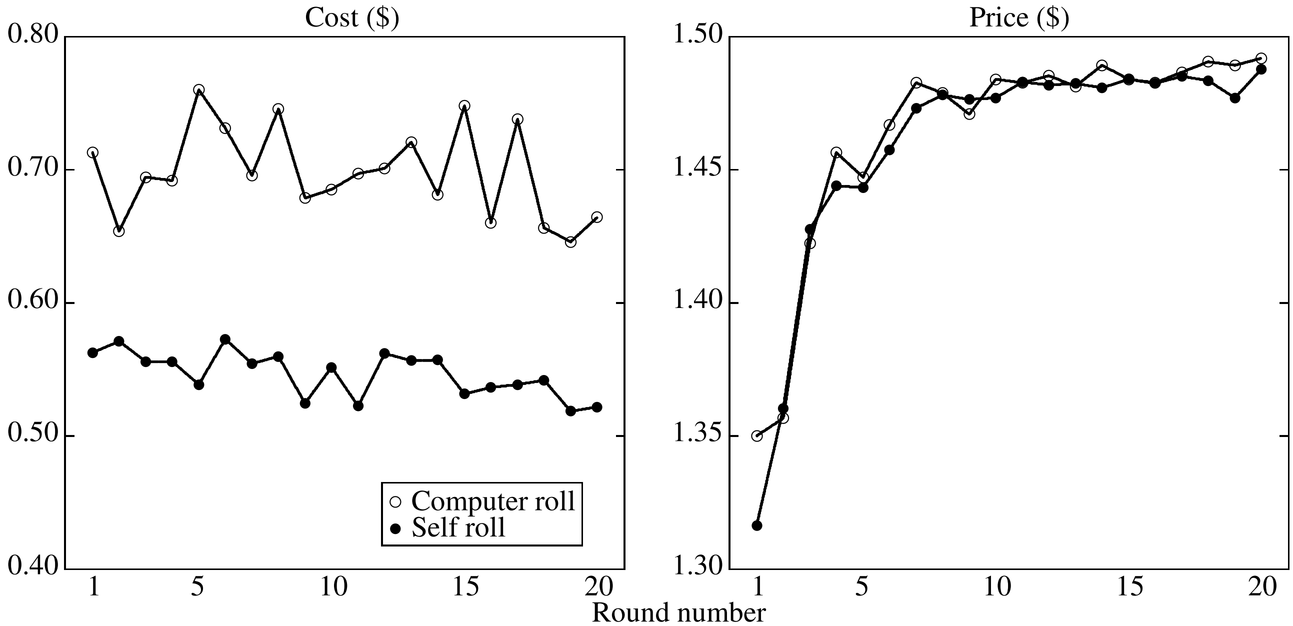

choices in the two treatments, and a robust rank-order test confirms there is no significant difference ($\grave{U} = 0.38$, two-tailed $q \approx 0.70$). In both, there is convergence toward the dominant action of $p = \$1.50$.

Figure 5 shows a disaggregation of the block-1 data, as a scatter-plot between individual subjects' average costs and their average price choices over the 20 rounds of this block. There is substantial heterogeneity in both costs and prices in both treatments, but importantly, the relationship between costs and prices differs between the two treatments. In the computer-roll treatment, there appears to be a positive correlation, though an OLS regression indicates that it is not significant.[19] In the self-roll treatment, there is a negative correlation and it is significant, with a slope coefficient of roughly –0.1. That is, higher costs are associated with lower prices in this treatment.

[16] A normal approximation yields a z-score of roughly –38.5.

[17] Since subjects did not interact, nor receive any feedback about others, in block 1, it is appropriate to treat each individual subject as independent. See Siegel and Castellan (1988) for descriptions of the non-parametric tests used in this paper, and Feltovich (2005) for critical values of the robust rank-order test used here and later. We refer to probabilities in our statistical tests with q (e.g., "q-value") to avoid confusing probabilities with prices, which continue to be denoted by "p".

[18] Indeed, nearly half (45 percent) of subjects in the self-roll treatment report average costs over rounds 1-20 that are lower than the lowest average cost of any subject in the computer-roll treatment. However, only 6 out of the 184 subjects in the self-roll treatment lie maximally (reporting the minimum possible cost of \$0.20 in every round), and only 14 (7.6 percent) do so in the last five rounds.

[19] This regression, and the corresponding one for the self-roll treatment, have average price choice over rounds 1-20 as the dependent

Figure 5: Block 1 scatter-plot of individual subjects' costs and price choices (round 1-20 averages)

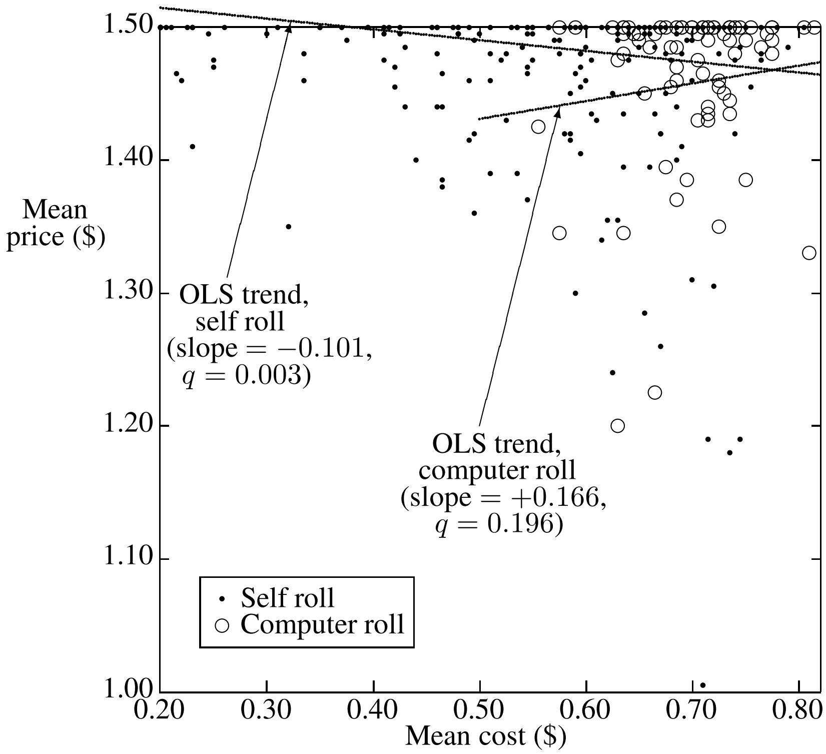

Keeping in mind that setting a price of $1.50 is strictly dominant, there is thus a negative relationship between lying about cost and decision errors in price choices. This suggests that one potential explanation for the observed level of truth-telling in our experiment may be an incomplete understanding of the experimental environment.[20]

4.2 Block 2 behaviour

Figure 6 shows the block-2 time series of subjects' costs in the three cells of the self-roll treatment. In all three cells, average costs start between $0.50 and $0.55 – roughly where they were at the end of block 1 – and continue decreasing over time, reaching about $0.45 by the last round of block 2. Differences across the cells are partially obscured by idiosyncratic noise over time within each treatment, but even so, we can see that reported costs are typically higher in the SRF cell than in SRW or SRS. Overall, average costs in the self-roll treatment are $0.50 under monopoly with feedback, $0.47 under weak competition and $0.48 under strong competition. These aggregates are not significantly different based on either a three-sample Jonckheere test or pairwise robust rank-order tests

variable, and a constant and the average cost over rounds 1-20 on the right-hand side. Other specifications, including Tobits replacing OLS, dropping one or more outliers, using only late rounds instead of all rounds from block 1, and using the fraction of optimal $1.50 choices as the dependent variable, yield similar results. We note that a positive relationship in the computer-roll treatment could be explained by a uniform random component in price choices (keeping in mind that price is constrained to be between the cost and $1.50) or by a "cost-plus" pricing heuristic.

[20] Another potential explanation is "moral cleansing" (Sachdeva et al., 2009): after unethical behaviour, people subsequently behave more morally or altruistically to atone. In our setting, subjects could be atoning for lying by reducing their prices. This is arguably not altruistic since our use of automated buyers means that no-one benefits from these lower prices, though one cannot rule out the possibility subjects have altruistic feelings toward the experimenter and incorrectly believe lower prices benefit him. In any case, harming oneself to atone for past sins - irrespective of any benefit to others - is certainly consistent with traditional religious forms of moral cleansing.

Figure 6: Block 2 time series, cost in $ (self-roll treatment)

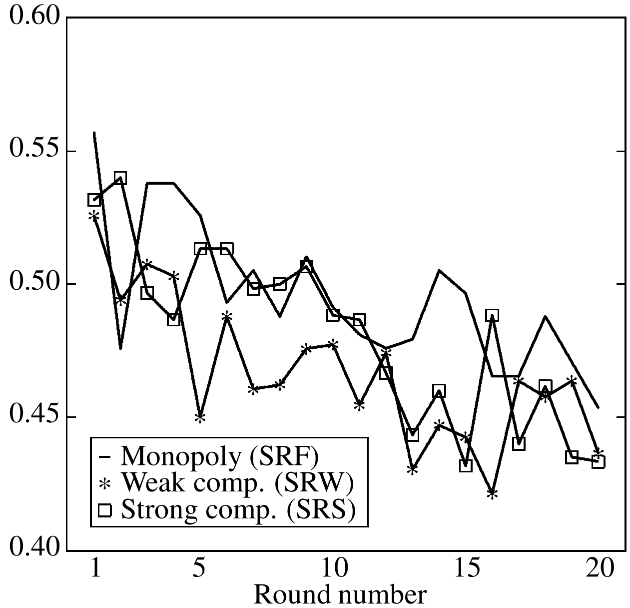

(two-tailed, session-level data, $q > 0.20$ in all cases). (However, as we will see shortly, some of these differences will become significant once we control for block-1 behaviour.) Also, there continue to be significant differences between the self-roll and computer-roll cells, with lower 20-round average costs in the SRS cell than in the CRS cell (robust rank-order test, two-sided, session-level data, $q \approx 0.036$) and in the SRW cell compared to the CRW cell ($q \approx 0.036$).[21] Thus, again we conclude that lying is taking place in the self-roll treatment, though again, the level of lying is far less than monetary payment maximisation would require.

Figure 7 shows the block-2 time series of price choices in the four duopoly cells. (Not shown is the SRF cell, where average prices are close to the dominant choice of $1.50: at least $1.45 in all rounds, and at least $1.48 in 17 out of the 20 rounds.) Within either weak or strong competition, price choices are noticeably higher in the computer-roll treatment than in the self-roll treatment. Non-parametric testing of 20-round averages finds significantly higher prices in the CRS cell than in the SRS cell, and in the CRW cell versus the SRW cell ($q \approx 0.036$ for both comparisons). However, comparison between the unconditional means is not very useful here, as we saw above that costs are lower in the self-roll cells than in the corresponding computer-roll cells, and we would expect lower costs to lead to lower prices.[22]

[21] Due to the interaction among subjects within sessions in block 2, we use session-level averages for non-parametric tests involving these data. Because of this, and because of the larger number of experimental cells than in block 1, significance results will generally be weaker here than in the block-1 results.

[22] Similarly, it should be unsurprising that within both the self-roll and the computer-roll treatment, prices are significantly lower under strong competition than under weak competition (robust rank-order test, two-sided, session-level data, $q \approx 0.008$ for self-roll and $q = 0.10$ for computer-roll; note for the latter that no lower two-sided q-value is possible for a 3-session-to-3-session comparison), and in the self-roll treatment, they are significantly lower under either than under monopoly with feedback ($q \approx 0.008$).

Figure 7: Block 2 time series, price in $ (duopoly cells)

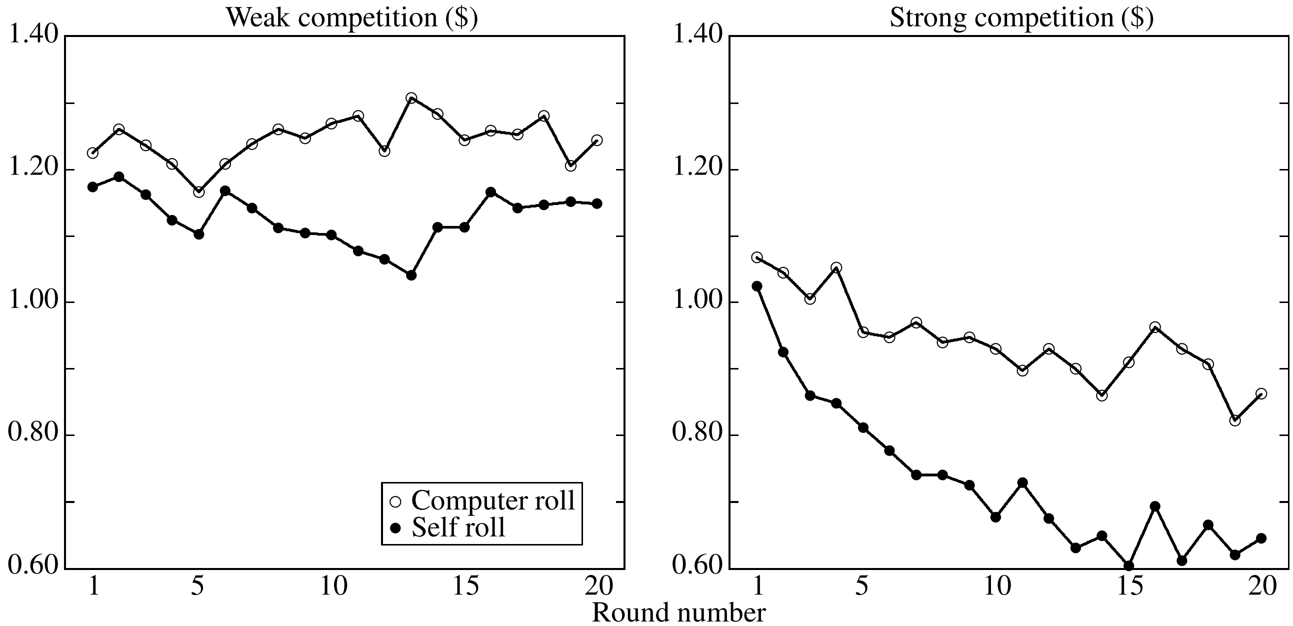

4.3 Regression analysis

For both the comparison of costs in Figure 6 and the comparisons of prices in Figure 7, relying on the unconditional means is unsatisfying. In the latter case, the differences in costs between self- and computer-roll treatments means that the differences in unconditional means are likely to over-state the differences in competition intensity due to the treatment. In the former case, uncontrolled variation in individual characteristics is likely to add noise to the unconditional averages, making significance more difficult to achieve and possibly obscuring true treatment effects.

To address both issues, we present estimation results from several panel Tobit models. We look first at costs, in Table 2. The dependent variable in each model shown there is the reported cost, which is restricted to be between 0.2 and 1.2 (hence the Tobit specification). The main explanatory variables in these cost regressions are indicators for the weak- and strong-competition treatments, so that the baseline is the monopoly-with-feedback treatment. Since these treatments were only present in the second block, we leave the block-1 data out of our sample (except for constructing one of our explanatory variables, as noted below). Also, since we are interested in reported costs rather than the computerised die rolls, we leave the computer-roll data out of our sample. Thus, our sample is block 2 of the self-roll treatment.

Additional explanatory variables for these Tobits are indicators for female and native-born students, the round number, the products of each of these with our treatment indicators (to allow for treatment effects that vary by demographic characteristics and over time), and subject random effects.[23] Models 1 and 2 are identical except for the presence of one additional variable in the latter: the individual subject's mean reported cost over block 1. This

[23] Additional Tobits, not reported here, included the complete set of demographic variables collected in the experiment. None of the additional variables were significant, nor did their inclusion affect the signs or significance levels of the variables shown in Table 2. Our results are also not substantially affected by using linear panel regressions (instead of Tobits) with clustering by individual subjects, clustering by session or bootstrapped standard errors. The same is true for the results reported in Table 3 below, and for still other regressions based on those models: including an equilibrium-price-squared variable to allow for a non-linear effect of the equilibrium price, or including the previous opponent's reported cost (as with Model 3 in Table 2).

variable serves as a proxy for the subject's intrinsic propensity toward lying, and thus helps to control for unobserved heterogeneity in our subject pool. Similarly, Models 2 and 3 are identical except for one variable added to the latter: the cost reported by the subject's previous opponent. (Even though subjects are re-matched in each round, they still may react to the previous-round result if lying is affected by perceptions of the strength of the honesty norm in the population.)

All three models in Table 2, as well as those presented in Table 3 below, were estimated using Stata (version 12). Both tables show marginal effects and standard errors for each variable, estimated using Stata's "margins"

Table 2: Tobit marginal effects (average unless specified), dependent variable = cost, block 2 of self-roll cells (standard errors in parentheses)

Variable	[1]	[2]	[3]
Strong competition indicator...avg. effect	–0.015	-0.025^{**}	-0.023^{*}
	(0.025)	(0.013)	(0.012)
...females	–0.031	-0.032^{*}	–0.029
	(0.038)	(0.018)	(0.018)
...males	–0.002	–0.019	–0.019
	(0.034)	(0.017)	(0.017)
...native born	–0.009	-0.055^{**}	-0.053^{*}
	(0.058)	(0.028)	(0.027)
...born elsewhere	–0.017	–0.017	–0.016
	(0.027)	(0.014)	(0.014)
Weak competition indicator...avg. effect	–0.025	–0.006	–0.004
	(0.025)	(0.012)	(0.012)
...females	–0.026	-0.035^{**}	-0.032^{*}
	(0.036)	(0.018)	(0.018)
...males	–0.024	0.020	0.021
	(0.033)	(0.017)	(0.017)
...native born	–0.086	-0.052^{*}	-0.048^{*}
	(0.060)	(0.029)	(0.028)
...born elsewhere	–0.010	0.006	0.008
	(0.027)	(0.014)	(0.014)
Female	0.033	0.003	0.005
	(0.020)	(0.010)	(0.010)
Native born	0.080^{***}	0.027^{**}	0.022^{*}
	(0.026)	(0.013)	(0.013)
Average cost in block 1		0.812^{***}	0.789^{***}
		(0.032)	(0.032)
Cost reported by prev. opponent			0.086^{***}
			(0.012)
Round number	-0.0034^{***}	-0.0034^{***}	-0.0027^{***}
	(0.0005)	(0.0005)	(0.0005)
N	3680	3680	3496
$\|ln(L)\|$	531.09	658.67	649.32

Notes: All marginal effects estimated over entire sample.
*,**,***: *Significantly different from zero at the 10%, 5%, 1% level.*

command. These are average marginal effects unless otherwise stated, in which case they are estimated conditional on particular values for one or more right-hand-side variables (including interactions if applicable) over the entire sample. For example, the "Strong competition indicator...females" marginal effect in the second row of Table 2 reports the marginal effect when the female indicator is set to 1 (and averaged over all values of the other variables).

The results from Model 1 show that when we do not control for subject heterogeneity via the "average cost in block 1" variable, the estimated coefficients of the competition-treatment indicators are negative (as predicted) but not significantly different from zero. By contrast, the results from Models 2 and 3 show that when we include the control for subject heterogeneity, both coefficients continue to have a negative sign, and that of the strong-competition treatment is significantly different from zero (though only weakly in Model 3: $q \approx 0.054$). Since the extra variables in Models 2 and 3 are highly significant, Model 3 outperforms Model 2, which in turn outperforms Model 1. Recalling that lying implies lower costs, we can thus conclude that lying increases under strong competition compared with monopoly.[24]

Examination of the remaining variables' average marginal effects, along with those conditional marginal effects of our main treatment variables shown in the table, yields a few interesting, though more tangentially relevant, conclusions. The significantly negative effect of the round number suggests increased lying over time, consistent with Figure 6. The significantly positive – and close to 1 – value for the effect of our "average cost in block 1" variable implies that individual heterogeneity in propensity to lie (whether from heterogeneous attitudes toward dishonesty or in understanding of the benefits to lying in this setting) persists from the first block to the second block. The positive and significant effect of the previous opponent reported cost suggests that subjects are more likely to under-report cost if their previous opponent had done so (or conversely, more likely to report truthfully if the previous opponent had done that). The conditional marginal effects of our weak- and strong-competition indicators are negative and sometimes significant for females and native-born subjects, while insignificant for males and subjects born elsewhere. We should be careful in drawing conclusions based on subgroup-level results, since we had no ex-ante hypotheses regarding any of them. However, these results suggest that lying by females and Australian-born subjects might be more affected by competition than lying by males and subjects born elsewhere.[25]

Result 1 *As we move from monopoly to weak duopoly and thence to strong duopoly, under-reporting of die rolls increases, though not always significantly.*

This result supports our Hypothesis 1, though we acknowledge that this support is somewhat weak due to the mixed significance results.

We next move to prices, in Table 3. The dependent variable is the price choice, which is constrained to be between the per-unit cost and \$1.50, so again we use a Tobit model. Since we are interested in whether competition intensifies in the self-roll treatment compared to the computer-roll treatment, we use the data from both of these treatments, but leave out the data from block 1 and from the SRF cell (where there was no competition, and also no computer-roll cell for comparison).

[24] The strong- and weak-competition indicators' marginal effects just miss being significantly different from each other ($q \approx 0.101$ in Model 3).

[25] Note that this statement is about different responses to the treatment, not about overall propensity to lie. The significant positive effect of the "native born" variable in the Model 1 results suggests that overall, lying by native-born subjects is significantly less prevalent than by subjects born elsewhere, while the corresponding effect for females also suggests less lying than by males, though not significantly so. A coherent, though speculative, explanation for these results would be that males are more intrinsically prone to unethical behaviour than females, leaving less room for more intense competition to further increase dishonesty (thus showing less of a treatment effect than females), and similarly for foreign- versus native-born subjects.

The main explanatory variable in these regressions is an indicator for the self-roll treatment. Because lying in this treatment means that costs are lower on average than in the computer-roll treatment, and because lower costs entail lower equilibrium prices, we control for the difference in costs across treatments. We do so in two ways. First, each model includes an "equilibrium price" variable, equal to the equilibrium price given the own and opponent costs and the market. Model 4 has no additional cost variables. Model 5 includes the own cost and the opponent cost as separate variables along with their interactions with the self-roll indicator, to allow for the possibility that costs affect pricing behaviour in ways other than via equilibrium pricing (for example, if subjects attempt to collude, prices may differ systematically from their equilibrium levels). Model 6 further includes the squares of these cost variables and their interactions with the self-roll indicator; thus, Models 5 and 6 allow linear own- and opponent-cost effects and quadratic effects, respectively. Additional explanatory variables are the female and native-born indicators and the round number, their interactions with the self-roll treatment indicator, and the average block-1 price chosen by the subject (as an indication of preference for high or low prices generally).

In all three models, the average marginal effect of the self-roll indicator is negative and significant, indicating that even controlling for differences in costs between self- and computer-roll treatments, prices are lower overall – and therefore competition is more intense – under self-roll. This is consistent with our Hypothesis 2, and thus seems to validate our Conjecture 1.[26] Further evidence for Conjecture 1 comes from the differences across treatments in the effect of the round number. Though prices decline over time in both self- and computer-roll treatments, the effect is significantly stronger ($p < 0.001$ in all three models) in the former. This is consistent with cooperation deteriorating more when lying is possible compared to when it is not, and is in the spirit of other results from the literature, such as the decay of above-equilibrium contribution levels in repeated public-good games (Isaac et al., 1985), as cooperative subjects learn that others are free-riding. Additionally, examination of the conditional marginal effects of the self-roll dummy suggests that its effect is present for both males and females, but limited to foreign-born subjects (the effect on natives is insignificant and close to zero). As with the cost results, care should be taken in drawing conclusions from these subgroup-level results, given the lack of ex-ante hypotheses.

Result 2 *When under–reporting of die rolls is possible, prices are lower (controlling for reported costs).*

This result supports our Hypothesis 2.

Again with this caveat about ex-ante hypotheses in mind, we close the analysis with some interesting though tangential results from the other variables. The equilibrium-price variable has the expected positive (and significant) marginal effect. Females choose lower prices than males, while overall there is no significant difference in pricing between native-born subjects and those born elsewhere (this last result suggesting that native-born subjects price higher on average in the self-roll treatment, and lower in the computer-roll treatment, than subjects born elsewhere). The positive and significant sign of the average price from block 1 suggests a positive correlation between prices in block 1 and prices in block 2, which is somewhat surprising given that one is under monopoly (where high prices

[26] Accordingly, we do not see evidence for the alternative conjecture discussed in Section 3.3 that subjects behave more cooperatively to reward anticipated truth telling across the board (which would imply a *positive* marginal effect). We also do not observe evidence for a positive interaction between the self-roll dummy and the rival cost, as would be implied by a strategy of conditional cooperation. Specifically, in [5], the estimated coefficient of the interaction term between self-roll and rival cost is not significantly different from zero (and indeed is *negative* rather than positive), and in both [5] and [6], the marginal effect of the rival cost is not significantly different between the self-roll and computer-roll treatments. We have found alternative specifications where this difference becomes significant, but we view these as inappropriate (e.g., [6] but with the equilibrium price removed as an explanatory variable).

Table 3: Tobit marginal effects (average unless specified), dependent variable = price, duopoly cells (standard errors in parentheses)

Variable	[4]	[5]	[6]
Self-roll treatment...avg. effect	-0.044^{***}	-0.031^{***}	-0.043^{***}
	(0.011)	(0.012)	(0.012)
...females	-0.037^{**}	-0.028^{*}	-0.037^{**}
	(0.016)	(0.016)	(0.017)
...males	-0.051^{***}	-0.034^{**}	-0.047^{***}
	(0.015)	(0.016)	(0.017)
...native born	0.001	0.012	0.001
	(0.021)	(0.022)	(0.022)
...born elsewhere	-0.059^{***}	-0.044^{***}	-0.057^{***}
	(0.013)	(0.014)	(0.014)
Equilibrium price	0.642^{***}	0.508^{***}	0.481^{***}
	(0.019)	(0.033)	(0.035)
Own cost		0.159^{***}	0.138^{***}
		(0.021)	(0.023)
Rival cost		-0.025^{*}	–0.010
		(0.014)	(0.015)
Weak duopoly treatment	-0.028^{*}	0.051^{**}	0.065^{***}
	(0.014)	(0.022)	(0.023)
Average price in block 1	0.210^{**}	0.234^{**}	0.211^{***}
	(0.102)	(0.103)	(0.103)
Female	-0.044^{***}	-0.044^{***}	-0.044^{***}
	(0.011)	(0.010)	(0.010)
Native born	–0.002	–0.003	–0.001
	(0.012)	(0.012)	(0.012)
Round number...avg. effect	-0.004^{***}	-0.005^{***}	-0.005^{***}
	(0.0005)	(0.0005)	(0.0005)
...self-roll	-0.006^{**}	-0.006^{*}	-0.006^{***}
	(0.0006)	(0.0006)	(0.0006)
...computer-roll	-0.002^{***}	-0.002^{**}	-0.002^{***}
	(0.0008)	(0.0008)	(0.0008)
Functional form for own and rival cost	None	Linear	Quadratic
N	4040	4040	4040
$\|ln(L)\|$	285.63	234.99	226.36

Notes: All marginal effects estimated over entire sample.

*,**,***: *Significantly different from zero at the 10%, 5%, 1% level.*

mainly suggest understanding of the setting) and the other is under duopoly (where high prices carry information about attitudes toward, or beliefs about, the rival firm). The positive and significant effect of the weak-duopoly indicator suggests that subjects are more competitive across-the-board under strong competition than under weak competition. One potential explanation for this is loss aversion: subjects may view selling zero as a severe loss under strong competition, whereas selling four units instead of ten under weak competition may be a loss, but a milder one. Finally, even though we also control for costs through the equilibrium price, the marginal effect of the own cost is positive and significant in [5] and [6], implying that subjects behave more competitively when they have low costs than when they have high costs. On the face of it, this suggests that truthfulness (associated with higher reported costs) and cooperativeness (associated with higher prices) might be positively correlated. However, we find this positive correlation in both self-roll and computer-roll treatment, and the marginal effects are not significantly different between the two treatments (indeed, in [5] it is actually higher in the computer-roll treatment). It is unclear why pricing would be more competitive – even in the computer-roll treatment – when costs are lower.

5 Discussion

We have investigated the interplay between competition and unethical behaviour. We find strong evidence of an effect in one direction: competition becomes more intense when facing someone who might be behaving unethically. This is observed not only in the descriptive statistics from our experiment – where lower prices in our self-roll cells than the corresponding computer-roll cells are a natural consequence of lower costs in the self-roll cells (due to lying about die rolls) – but also in our regression results, where we control for the direct effects of costs in several ways. To our knowledge, this direction of the relationship between competition and unethical behaviour has not been studied before, although it fits in well with findings from other settings in which individuals willingly bear a cost in order to punish bad behaviour (see Section 2).

By contrast, evidence of an effect in the other direction is weaker. Qualitatively, our regression results are consistent with more lying about die rolls under strong competition than under weak competition, and more lying under weak competition than under monopoly. However, only the difference between strong competition and monopoly is significant at conventional levels. On the surface, this would seem inconsistent with previous experimental results like those mentioned at the end of Section 2, where lying becomes more likely when competing against a rival who also can lie. However, as noted in Section 3.2, our setting has the feature that as competition increases (as we move from monopoly to weak duopoly to strong duopoly), the positive effect of a given lie on one's own money payoff becomes smaller in absolute terms (though it increases the negative effect on the rival's money payoff). By contrast, nearly all of the earlier studies had (on average) competition *increasing* the monetary benefit to oneself from a given lie. Thus, a very simple model of dishonesty – with cost of lying based on its extent and benefit based on the resulting increase in money payment, with no other-regarding component – would imply *less* lying as competition increases in our setting, but *more* lying in these other settings. The fact that we actually observe even weakly more lying suggests that such a simple model is lacking. Perhaps subjects view the negative externality of lying on rivals as a positive benefit, rather than neutrally or as a cost. Perhaps the cost of lying is reduced in market or competitive settings. Alternatively, subjects may view the benefits or costs of lying in a relative instead of absolute sense (e.g., taking account of the overall stakes involved). There may well be other potential explanations.

One other result from our experiment is worth re-iterating: in block 1, when all subjects are monopolists, we saw a significant negative correlation across subjects in the self-roll treatment between average reported costs and average price choices. (By contrast, there was a positive but insignificant correlation in the computer-roll treatment.) Since lying about the die rolls should reduce reported costs, while the dominant price choice for a monopolist is the maximum, this means that more truthful reporting of costs is associated with failure to choose the optimal price, and suggests that one of the factors behind truth-telling in our experiment may be an incomplete understanding of the setting. (As noted previously, our use of computerised buyers in the experiment probably rules out the obvious alternative explanation of pro-social preferences toward buyers.) This may be true in other experimental studies as well: unless the decision-making environment is quite transparent, the extent of honesty in the population may be over-estimated due to some potentially opportunistic subjects not figuring out the incentive to lie.

Our attention in this study has been focussed on behaviour, but it is worthwhile to take a moment to consider the welfare consequences of our main results, and specifically their implication that unethical behaviour and intensified competition should occur in tandem in real markets. If our market setting is taken literally, both lying and competitive pricing will benefit consumers via lower prices. Taken more figuratively, we would still expect increased competition in pricing to be social-welfare-increasing, though its benefit may be limited if its effect in the longer term is to drive sellers out of the market. The effect of lying is harder to discern, and will depend sensitively on what assumptions are made about consumers' preferences and the nature of the good. As discussed in the introduction, many of the real-life analogues of this market involve goods whose quality may not be discerned until long after they have been bought and used (e.g., finding out that a fast-food outlet uses adulterated ingredients or under-pays its workers may cause disutility even to someone who bought from it years ago). For such goods, a buyer's current willingness-to-pay may over-state her true realised value – possibly by even more than the decrease in price resulting from unethical seller behaviour. Then, the net effect on consumer surplus and overall social welfare could be ambiguous, even after taking into account the intensified competition.

We close by mentioning some extensions of this study that may be suitable for future research. Some of these involve additional treatments added to the current experiment. For example, we did not conduct a CRF (computer-roll, monopoly with feedback) cell because we expected behaviour to be similar to that in our SRF cell, but this assumption could be tested empirically. A new treatment where die-roll subjects were paired with computer-roll subjects could help understand whether subjects lie when they know their rival cannot lie, and distinguish between pricing competitively due to own lying versus others' lying. A treatment in which subjects rolled dice but the results were entered by a monitor (to prevent lying) could shed light on whether lower prices in our self-roll treatment were indeed due to lying, or rather to the die-rolling itself. (Perhaps generating their own costs makes subjects choose lower prices.)

Other extensions involve larger changes to our design. In our setting, lying that harmed others seems to have made preferences less pro-social (or more anti-social). This raises the question of what would have happened if the environment were such that lying *benefited* others. If the change in preferences we observed was due to the negative consequence of the lying on others, then lying in this new environment ought to *strengthen* pro-social preferences. However, if the weakening was due to lying representing a general decline in ethical behaviour, then even lying that benefits others may *weaken* pro-social preferences. An experimental test would serve to distinguish between these two mechanisms. A second experimental extension, using triopolies instead of duopolies, could examine the effect on competition in situations where one, but not the other, rival seller is likely to have under-stated costs. There is also scope for future theoretical work. We justified our hypotheses in Section 3.3 by appealing to intuition and previous results from the literature, but it ought to be possible to construct a model that entails our hypotheses. Such a model might contain an aversion to lying and (positive or negative) other-regarding preferences, with the extent of these determined endogenously.

References

Abbink, Klaus and Abdulkarim Sadrieh (2009), "The pleasure of being nasty", *Economics Letters* 105, pp. 306-308.

Abeler, Johannes, Anke Becker and Armin Falk (2014), "Representative evidence on lying costs", *Journal of Public Economics* 113, pp. 96-104.

Abeler, Johannes, Daniele Nosenzo and Collin Raymond (2016), "Preferences for truth-telling", CESifo Working Paper Series No. 6087.

Aimone, Jason A., Sheryl Ball and Brooks King-Casas (2013), "Social and economic risk: results from an individual level betrayal aversion measure", working paper, Virginia Tech.

Aimone, Jason A. and Daniel Houser (2012), "What you don't know won't hurt you: a laboratory analysis of betrayal aversion", *Experimental Economics* 15, pp. 571-588.

Alempaki, Despoina, Gönül Doğan, and Silvia Saccardo (2016), "Deception and reciprocity", working paper.

Anbarci, Nejat, Nick Feltovich and Mehmet Y. Gürdal (2015), "Lying about the price? Ultimatum bargaining with messages and imperfectly observed offers", *Journal of Economic Behavior & Organization* 116, pp. 346-360.

Ariely, Daniel (2012), *The Honest Truth About Dishonesty: How We Lie to Everyone – Especially Ourselves*, New York: HarperCollins.

Banerjee, Ritwik (2016), "Corruption, norm violation and decay in social capital", IZA Discussion Paper No. 9859.

Barr, Abigail and Georgia Michailidou (2016), "Complicity without connection or communication", CeDEx discussion paper 2016-14, University of Nottingham.

Batson, C. Daniel, Diane Kobrynowicz, Jessica L. Dinnerstein, Hannah C. Kampf, and Angela D. Wilson (1997), "In a very different voice: unmasking moral hypocrisy", *Journal of Personality and Social Psychology* 72, pp. 1335-1348.

Batson, C. Daniel, Elizabeth R. Thompson, Greg Seuferling, Heather Whitney and Jon A. Strongman (1999), "Moral hypocrisy: appearing moral to oneself without being moral", *Journal of Personality and Social Psychology* 77, pp. 525-537.

Belot, Michèle and Marina Schröder (2013), "Sloppy work, lies and theft: a novel experimental design to study counterproductive behaviour", *Journal of Economic Behavior & Organization* 93, pp. 233-238.

Berger, Joël and Debra Hevenstone (2016), "Norm enforcement in the city revisited: an international field experiment of altruistic punishment, norm maintenance, and broken windows", *Rationality and Society* 28, pp. 299-319.

Bohnet, Iris and Richard Zeckhauser (2004), "Trust, risk and betrayal", *Journal of Economic Behavior & Organization* 55, pp. 467-484.

Bolton, Gary E. and Axel Ockenfels (2000), “ERC: a theory of equity, reciprocity, and competition”, *American Economic Review* 90, pp. 166-193.

Brandts, Jordi and Gary Charness (2003), “Truth or consequences: an experiment”, *Management Science* 49, pp. 116-130.

Cappelen, Alexander W., Erik Ø. Sørensen and Bertil Tungodden (2013), “When do we lie?” *Journal of Economic Behavior & Organization* 93, pp. 258-265.

Carpenter, Jeffrey, Peter Hans Matthews and John Schirm (2010), “Tournaments and office politics: evidence from a real effort experiment”, *American Economic Review* 100, pp. 1-18.

Diekmann, Andreas, Wojtek Przepiorka and Heiko Rauhut (2015), “Lifting the veil of ignorance: an experiment on the contagiousness of norm violations”, *Rationality and Society* 27, pp. 309-333.

Dufwenberg, Martin and Uri Gneezy (2000), “Price competition and market concentration: an experimental study”, *International Journal of Industrial Organization* 18, pp. 7-22.

Dufwenberg, Martin, Maroš Servátka and Radovan Vadovič (2017), “Honesty and informal agreements”, *Games and Economic Behavior* 102, pp. 269-285.

Dutcher, Glenn, Daniela Glätzle-Rützler and Dmitry Ryvkin (2016), “Don’t hate the player, hate the game: uncovering the foundations of cheating in contests”, Working Papers in Economics and Statistics 2016-29, University of Innsbruck.

Edgeworth, Francis Y. (1925), “The pure theory of monopoly”, in Francis Y. Edgeworth, ed., *Papers Relating to Political Economy*, Volume 1, New York: Burt Franklin, pp. 111-142.

Ellingsen, Tore and Magnus Johannesson (2008), “Pride and prejudice: the human side of incentive theory”, *American Economic Review* 98, pp. 990–1008.

Engelmann, Dirk and Martin Strobel (2004), “Inequality aversion, efficiency, and maximin preferences in simple distribution experiments”, *American Economic Review* 94, pp. 857-869.

Faravelli, Marco, Lana Friesen and Lata Gangadharan (2015), “Selection, tournaments, and dishonesty”, *Journal of Economic Behavior & Organization* 110, pp. 160-175.

Fehr, Ernst and Urs Fischbacher (2004), “Third-party punishment and social norms”, *Evolution and Human Behavior* 25, pp. 63-87.

Fehr, Ernst, and Klaus M. Schmidt (1999), “A theory of fairness, competition and cooperation”, *Quarterly Journal of Economics* 114, pp. 817-868.

Feltovich, Nick (2005), “Critical values for the robust rank-order test”, *Communications in Statistics – Simulation and Computation* 34, pp. 525-547.

Fischbacher, Urs (2007), “z-Tree: Zurich toolbox for ready-made economic experiments”, *Experimental Economics* 10, pp. 171-178.

Fischbacher, Urs and Franziska Föllmi-Heusi (2013), "Lies in disguise: an experimental study on cheating", *Journal of the European Economic Association* 11, pp. 525-547.

Gächter, Simon and Jonathan F. Schulz (2016), "Intrinsic honesty and the prevalence of rule violations across societies", *Nature* 531, pp. 469-499.

Gibson, Rajna, Carmen Tanner and Alexander F. Wagner (2013), "Preferences for truthfulness: heterogeneity among and within individuals", *American Economic Review* 103, pp. 532-548.

Gneezy, Uri (2005), "Deception: the role of consequences", *American Economic Review* 95, pp. 384-394.

Greiner, Ben (2015), "Subject pool recruitment procedures: organizing experiments with ORSEE", *Journal of the Economic Science Association* 1, pp. 114-125.

Holt, Charles A. (1995), "Industrial organization: a survey of laboratory research", in *Handbook of Experimental Economics*, J.H. Kagel and A.E. Roth, eds., Princeton University Press, Princeton, NJ., pp. 349-443.

Houser, Daniel, John A. List, Marco Piovesan, Anya Samek and Joachim Winter (2016), "Dishonesty: from parents to children", *European Economic Review* 82, pp. 242-254.

Houser, Daniel, Stefan Vetter and Joachim Winter (2012), "Fairness and cheating", *European Economic Review* 56, pp. 1645-1655.

Hurkens, Sjaak and Navin Kartik (2009), "Would I lie to you? On social preferences and lying aversion", *Experimental Economics* 12, pp. 180-192.

Innes, Robert and Arnab Mitra (2013), "Is dishonesty contagious?" *Economic Inquiry* 51, pp. 722-734.

Isaac, Mark, Kenneth F. McCue and Charles Plott (1985), "Public goods provision in an experimental environment", *Journal of Public Economics* 26, pp. 51-74.

Keizer, Kees, Siegwart Lindenberg and Linda Steg (2008), "The spreading of disorder", *Science* 322, pp. 1681-1685.

Keizer, Kees, Siegwart Lindenberg and Linda Steg (2013), "The importance of demonstratively restoring order", *PLoS ONE* 8, e65137.

Kelling, George L. and James Q. Wilson (1982) "Broken windows: the police and neighbourhood safety", *The Atlantic Monthly* (March), pp. 29-39.

Keuschnigg, Marc and Tobias Wolbring (2015), "Disorder, social capital, and norm violation: three field experiments on the broken windows thesis", *Rationality and Society* 27, pp. 96-126.

Kocher, Martin G., Simeon Schudy and Lisa Spantig (2017), "I lie? We lie! Why? Experimental evidence on a dishonesty shift in groups", forthcoming, *Management Science*.

Kriss, Peter H., Rosemarie Nagel and Roberto A. Weber (2013), "Implicit vs. explicit deception in ultimatum games with incomplete information", *Journal of Economic Behavior & Organization* 93, pp. 337-346.

Kriss, Peter H., Roberto Weber and Erte Xiao (2016), “Turning a blind eye: on the robustness of costly third party punishment in humans”, *Journal of Economic Behavior and Organization* 128, pp. 159-177.

Ledyard, John O. (1995), “Public goods: a survey of experimental research”, in *Handbook of Experimental Economics*, J.H. Kagel and A.E. Roth, eds., Princeton University Press, Princeton, NJ., pp. 111-194.

Levine, David K. (1998), “Modeling altruism and spitefulness in experiments”, *Review of Economic Dynamics* 1, pp. 593-622.

López-Pérez, Raul (2012), “The power of words: a model of honesty and fairness”, *Journal of Economic Psychology* 33, pp. 642-658.

López-Pérez, Raul and Eli Spiegelman (2013), “Why do people tell the truth? Experimental evidence for pure lie aversion”, *Experimental Economics* 16, pp. 233-247.

Lundquist, Tobias, Tore Ellingsen, Erik Gribbe and Magnus Johannesson (2009), “The aversion to lying”, *Journal of Economic Behavior & Organization* 70, pp. 81-92.

Mazar, Nina, On Amir and Dan Ariely (2008), “The dishonesty of honest people: a theory of self-concept maintenance”, *Journal of Marketing Research* 45, pp. 633-644.

Mitzkewitz, Michael and Rosemarie Nagel (1993), “Experimental results on ultimatum games with incomplete information”, *International Journal of Game Theory* 22, pp. 171-198.

Nikiforakis, Nikos and Helen Mitchell (2014), “Mixing the carrots with the sticks: third party punishment and reward”, *Experimental Economics* 17, pp. 1-23.

Ohtsubo, Yohsuke, Fumiko Masuda, Esuka Watanabe and Ayumi Masuchi (2010), “Dishonesty invites costly third-party punishment”, *Evolution and Human Behavior* 31, pp. 259-264.

Peeters, Ronald, Marc Vorsatz and Markus Walzl (2008), “Rewards in an experimental sender-receiver game”, *Economics Letters* 101, pp. 148-150.

Rigdon, Mary L. and Alexander P. D’Esterre (2015), “The effects of competition on the nature of cheating behavior”, *Southern Economic Journal* 81, pp. 1012-1024.

Roth, Alvin E. and J. Keith Murnighan (1982), “The role of information in bargaining: an experimental study”, *Econometrica* 50, pp. 1123-1142.

Sachdeva, Sonya, Rumen Iliev and Douglas L. Medin (2009), “Sinning saints and saintly sinners: the paradox of moral self-regulation”, *Psychological Science* 20, pp. 523-528.

Sánchez-Pagés, Santiago and Marc Vorsatz (2007), “An experimental study of truth-telling in a sender-receiver game”, *Games and Economic Behavior* 61, pp. 86-112.

Schurr, Amos and Ilana Ritov (2016), “Winning a competition predicts dishonest behavior”, *Proceedings of the National Academy of Sciences* 113, pp. 1754-1759.

Schwieren, Christiane and Doris Weichselbaumer (2010), “Does competition enhance performance or cheating? A laboratory experiment”, *Journal of Economic Psychology* 31, pp. 241-253.

Seinen, Ingrid and Arthur Schram (2006), “Social status and group norms: indirect reciprocity in a repeated helping experiment”, *European Economic Review* 50, pp. 581-602.

Serra, Danila and Leonard Wantchekon, eds. (2012), *Research in Experimental Economics, v. 15: New Advances in Experimental Research on Corruption*, Emerald Group Publishing Limited, Bingley, UK.

Siegel, Sidney and N. James Castellan, Jr. (1988), *Nonparametric Statistics for the Behavioral Sciences*, McGraw-Hill, New York.

Thaler, Richard (1980), “Toward a positive theory of consumer choice”, *Journal of Economic Behavior & Organization* 1, pp. 39-60.

Weisel, Ori and Shaul Shalvi (2015), “The collaborative roots of corruption”, *Proceedings of the National Academy of Sciences* 112, pp. 10651-10656.

Wiltermuth, Scott S. (2011), “Cheating more when the spoils are split”, *Organizational Behavior and Human Decision Processes* 115, pp. 157-168.

Zizzo, Daniel J. and Andrew J. Oswald (2001), “Are people willing to pay to reduce others’ incomes?” *Annales d’Économie et de Statistique* 63-64, pp. 39-62.

A Sample instructions from the experiment

Below are the instructions from one of our cells: self-roll and (in the second block) weak competition. The instructions for other cells are similar, and available from the author upon request. Note that in all cells, the instruction for the first and second blocks (labelled “Part 1” and “Part 2” respectively) were distributed before the first block began and between blocks 1 and 2 respectively.

Instructions (Part 1)

You are about to participate in a decision making experiment. Please read these instructions carefully. If you have a question at any time, please feel free to ask the experimenter. We ask that you not talk with the other participants during the experiment, and that you put away your mobile phones and other devices at this time.

This experiment has two parts. These instructions are for Part 1; you will receive new instructions after this part has finished. Part 1 is made up of ***20 rounds***, each consisting of a simple computerised market game. You are a ***seller***; you will remain in this role throughout the experiment.

The market game: In each round, you can produce up to ***10 units*** of a hypothetical good. There are 10 automated buyers in your market. Each will buy a single unit of your good as long as the price is ***$1.50 or less***. No buyer can buy more than one unit in a round.

A round begins with you finding out your cost of production. You will roll two dice and enter the results into the computer. ***The sum of these dice, multiplied by $0.10, will be your per-unit cost of production***. So your per-unit cost will be at least $0.20, and at most $1.20. After you learn your cost, you choose your price, which is entered as a ***multiple of $0.10, between your per-unit cost and $1.50*** inclusive. (Don’t type the dollar sign).

Profits: Your profit for the round is given by the formula

Profit = (quantity sold) * (price – cost)

Payments: Your payment for this experiment depends on the results. At the end of the experiment, ***two*** rounds from Part 1 will be chosen randomly for each of you. ***You will be paid the total of your profits from those selected rounds***, plus whatever you earn in the remainder of the experiment. Payments are made privately and in cash at the end of the session.

Instructions (Part 2)

In Part 2, you will continue in a similar setting as in Part 1, for an additional ***20 rounds***. Feel free to refer to your instructions from Part 1 for details.

The main difference from the setting of Part 1 is that now, you and another seller will be producing the ***same good*** and serving the ***same market***. The other seller in your market faces a situation ***identical*** to yours. His/her per-unit cost is determined in the same way as yours is, and he/she chooses a price as you do. Of course, the actual values of the other seller's cost and price may differ from yours.

A second difference from Part 1 is that now, there are ***14 buyers*** in your market instead of 10. As before, each buyer will buy one unit of the good as long as the price is $1.50 or less, and no buyer will buy a second unit. But buyers will buy from the ***lower-priced seller*** if they can. So, if your price is ***lower*** than the other seller's price, you will sell ***10 units*** and the other seller will sell 4. If your price is ***higher*** than the other seller's, you will sell ***4 units*** and the other seller will sell 10. If your prices are equal, you will each sell ***7 units***.

As in Part 1, your profit is given by the formula

Profit = (quantity sold) * (price – cost)

The other seller in your market is chosen randomly in each round, so may differ from round to round. He/she will see no identifying information about you, nor will you see any such information about him/her.

After the last round of Part 2, you will be shown another screen containing a questionnaire. Once all participants have completed the questionnaire, ***two*** rounds from Part 2 will be chosen randomly for each of you. ***The total of your profits from those selected rounds*** will be added to what you earned from Part 1, and you will receive an additional $10 for completing the questionnaire.

B Sample screenshots from the experiment

Below are sample screenshots adapted from the experiment. They are modified from the original screenshots to improve readability, by removing information from the top of the screen (round number, history table, time remaining in the stage), and grey space from the left, right, and bottom of the screen. Additional screenshots are available from the author upon request.

Cost reporting screen (block 1, self-roll treatment):

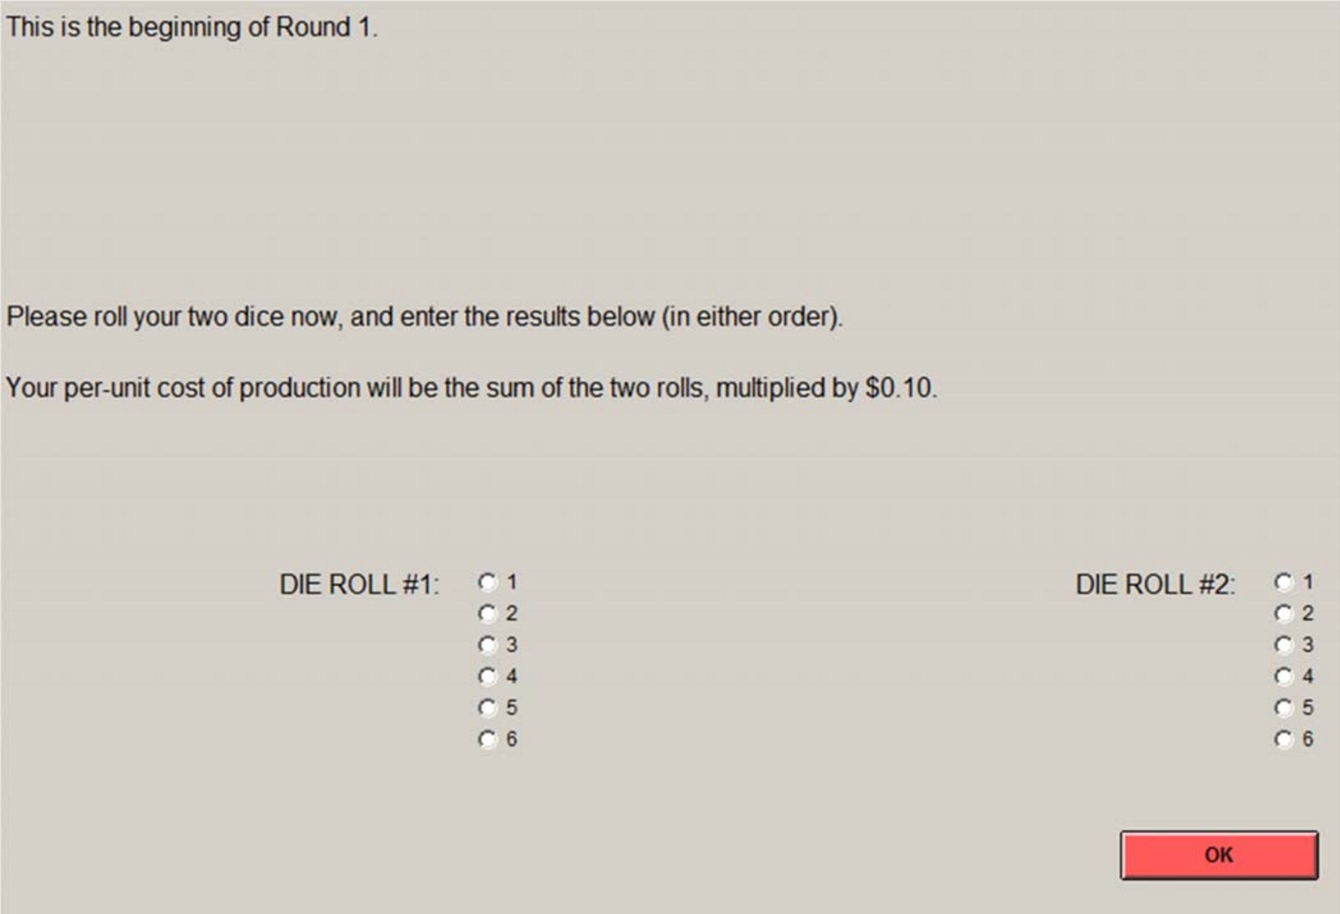

Price choice screen (block 1, self-roll treatment):

Your two die rolls were **3** and **4**, so your cost of production is **0.70**.

Please enter the price (in $) you will charge. It can be any multiple of 0.10 between 0.70 and 1.50 inclusive.

MY PRICE $:

OK

Price choice screen (block 2, strong-competition treatment):

Your two die rolls were **4** and **3**, so your cost of production is **0.70**.

The two die rolls of the seller matched to you were **1** and **1**, so his/her cost of production is **0.20**.

Please enter the price (in $) you will charge. It can be any multiple of 0.10 between 0.70 and 1.50 inclusive.

MY PRICE $:

OK

Remember that if your price is **less than** the other seller's price, you will sell **10** units.

If your price is **greater than** the other seller's price, you will sell **0** units.

And if your prices are **equal**, you will each sell **5** units.

Feedback screen (block 2, strong-competition treatment):

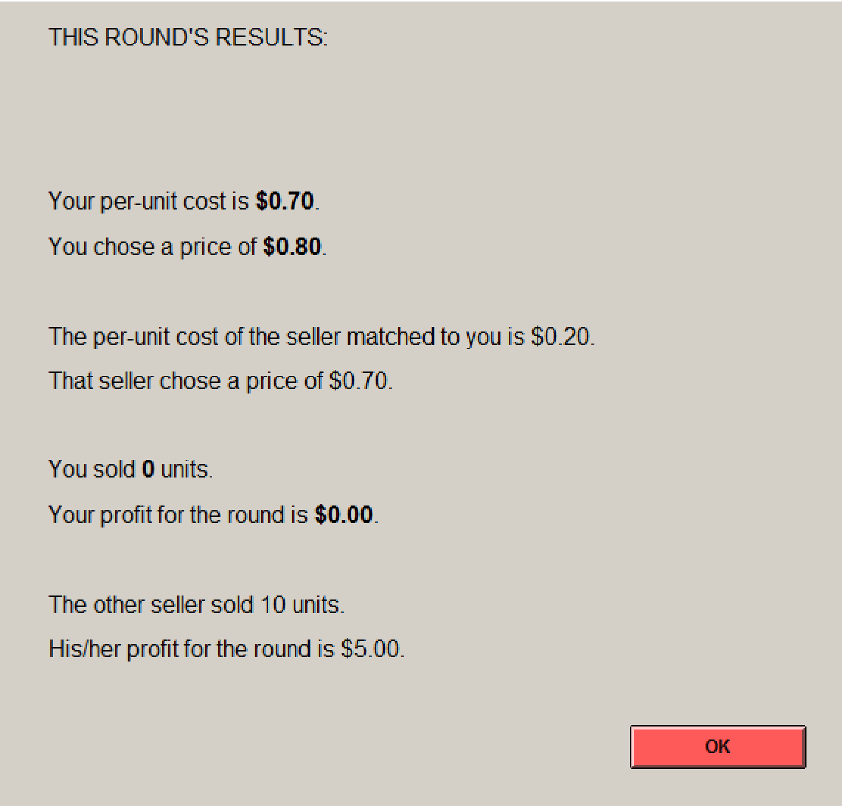
THIS ROUND'S RESULTS:

Your per-unit cost is **$0.70**.

You chose a price of **$0.80**.

The per-unit cost of the seller matched to you is $0.20.

That seller chose a price of $0.70.

You sold **0** units.

Your profit for the round is **$0.00**.

The other seller sold 10 units.

His/her profit for the round is $5.00.

OK

Questionnaire screen (after block 2 ends):


Please answer the following questions.

What is your age, to the nearest year?

What is your gender?
- Female
- Male

Were you born in Australia?
- Yes
- No

How many years have you lived in Australia (to the nearest year)?

What is your area of study?
- Economics
- Other business
- Other

Since you began studying at university, how many economics units have you completed?
- 0
- 1
- 2
- 3
- 4 or more

Finished