Inducing good behavior

van der Veen, A.

Publication date
2012

1. Introduction

Using the experimental method, we analyze mechanisms to induce "good" behavior in four cases. Good is seen from the perspective of a superior in a hierarchical relation, where command and control is not an economically feasible option. The four cases are:

1. A government wants to induce good behavior with the help of subsidies. We analyze whether it is more effective to introduce the subsidy gradually or introduce it in one big step.

2. A government wants to induce good behavior with the help of automatic bonuses and automatic fines. We analyze which of the two instruments is the more effective one.

3. An employer wants to induce good behavior from a worker by rewarding desired behavior and punishing undesired behavior. This time the application of instruments is not automatic, but a discretionary power of the employer. Again, we analyze which of the instruments is more effective.

4. A government wants to induce good behavior from limitedly liable bidders in an auction. The government doesn't want the bidders to overbid, in a situation where post auction bankruptcy is undesirable. We compare the English auction to the first-price sealed-bid auction, with respect to the likelihood of post auction bankruptcy.

We use laboratory experiments, while we also could have used mechanism design to construct a theoretically optimal mechanism.[1] However, most models used in mechanism design, assume agents to be rational, selfish, and making decisions not affected by emotions. Experimental evidence, both from the lab and from the field, shows that said assumptions often do not hold.[2] As we do not have at our disposal a unifying theory of human behavior, we rely on laboratory experiments to study the four cases. In each case, we will confront subjects with two commonly used mechanisms, and study which of the two performs best.

In Chapter 2, we investigate how to introduce subsidies aimed at steering behavior. In 2009, the Japanese government introduced a 10% subsidy on solar power panels. As the subsidy turned out to be less effective than planned, it is expected to be raised in the future (Leader, 2009). In the same year, the Chinese government announced a 50% subsidy on these panels, the highest

[1] For a discussion of mechanism design, see Myerson (1981).
[2] For an overview of this problem, see e.g., Tirole (2002).

such subsidy in the world (Ideas, 2009). As subsidies are important instruments for governments, we test whether an introduction in one step or a gradual introduction is more effective.

In our experiment we use a public good game, where participants decide every round how much to contribute. The total contribution is raised by some fraction (20% in our case), free of cost for the participants. The pot is equally split and paid out to all participants, independent of their contributions. These rules make contributing 0 the dominant strategy for each participant. We augment this game with a subsidy. The subsidy we use is a reduction in the cost of contributing. If the subsidy is 0.45, contributing 10 to the common pot cost the participant only $(1-0.45)\times 10 = 5.5$.

We compare two treatments, the quick treatment and the gradual treatment. In both treatments the subsidy begins at a level of 0 and after a certain amount of rounds the subsidy starts increasing. In the quick treatment the subsidy switches to the target-level in one step. In the gradual treatment the subsidy is slowly raised each round until the target-level is reached. When it is reached, the subsidy stays at the target-level until the end of the experiment.[3]

From experiments without subsidies, we already know that we can expect some participants to contribute. Furthermore, the subsidy makes contributions more effective and we know for example from Isaac and Walker (1988) and Isaac, Walker, and Williams (1994) that participants tend to contribute more when contributions are more effective. In the literature, two explanations are offered. One is the existence of material altruists, who care about the payoff of other people, and give more because their contribution is made more effective (Goeree, Holt, and Laury, 2002). The other one is the existence of conditional cooperators in public good games (Offerman, Sonnemans, and Schram, 1996; Fischbacher, Gächter, and Fehr, 2001; Brandts and Schram, 2001). Conditional cooperators choose their contribution conditional on their expectations concerning what others are going to contribute. If contributing is made more effective, they could become more optimistic about others contributing and therefore contribute more themselves

In contrast to this literature we do not focus on the reason why people react to a subsidy, but focus on how they react to the implementation of the subsidy, either quick or gradual. Interestingly, the idea of conditional cooperators could still play a role. If conditional cooperators expect the other players to react more to an introduction in one step and less to an introduction in small steps, they could also be inclined to react stronger to an introduction in one step. Another option is that it is not so much expectations that drive the results, but anchoring (Tversky and Kahneman, 1974). The initial subsidy serves as a reference point: participants only change their behavior if there is a noticeable change in the subsidy.

The experiment shows a difference in the change of contributions between the two treatments, but only if the target-level is high enough. We compare target-levels of 0.45 and 0.75. In treatments where the target-level is 0.45 subjects do not respond differently to a quick or gradual increase of the subsidy: contributions to the public good are hardly raised during the experiment

[3] To check whether a possible treatment effect could be explained by distraction (as faced in real live), we ran treatments with and treatments without a second task to be performed by the participants simultaneously with the public good task. This addition of an extra task does not produce a difference in contribution.

anyway. When the target-level is 0.75, again subjects hardly respond to a gradual increase of the subsidy, but when the subsidy is introduced in one step they very significantly raise their contribution to the public good. From the experiment we can conclude that to influence behavior, it is better to introduce a substantial subsidy at once than in small steps.

While in Chapter 2 we focus on authorities using subsidies to influence behavior, in Chapter 3 the focus is on authorities using either punishment or reward to encourage good behavior. In 2009, the Dutch tax authorities increased the fine for not reporting savings from 100% to 300% and announced further increases (Tweede Kamer, 2009). In 2003, the South Korean tax authorities started rewarding taxpayers having high compliance levels (NTS, 2004). Punishment of bad behavior and reward of good behavior are instruments often used by authorities. In an experiment, we test which instrument works better.

We investigate the question with the help of an inspection game, with two players, one called the inspector and the other one called the inspectee. In each round, both inspector and inspectee independently and simultaneously make a decision. The inspector decides whether to do a costly inspection and the inspectee decides whether to work, which is costly for the inspectee. The inspector has to pay the inspectee a wage (higher than the cost of working), except when the inspectee decided not to work and the inspector decided to inspect. When the inspectee works, the payoff of the inspector is enlarged more than an inspection costs.

To the baseline game we add either an automatic fine or an automatic bonus, but only if the inspector chose inspection. Fines are paid by the inspectee and received by the inspector; bonuses are paid by the inspector and received by the inspectee. After each round of the game, players are randomly rematched in new pairs of one inspector and one inspectee, but during the whole experiment a participant only plays one of the two roles.

We observe that the inspectee performs better under automatic fines than under automatic bonuses. This result is in line with predictions of the mixed strategy Nash equilibrium where players make their decisions dependent on the payoffs of the other player. If an inspectee knows that an automatic fine is introduced that adds to the payoff of the inspector, the inspectee will expect the inspector to inspect more often in order to collect the fine. To avoid the fine the inspectee will decide to work more often and this is what we see happen. However, this can not be the whole story. In line with the previous reasoning adding automatic bonuses should lead to less work, and this is not what we observe. There is only an insignificant difference in the decision to work for both treatments. These results can be fairly well explained by recent behavioral models based on respectively impulse balance equilibrium (Selten and Chmura, 2008) and quantal response equilibrium (McKelvey and Palfrey, 1995). We can conclude that automatic fines work better than automatic bonuses, but in contrast to the standard game theoretical prediction, automatic bonuses are not detrimental to the decision to work.

In Chapter 4 we focus again on punishment and reward, but this time in the context of employers and workers in a fairly standard labor relationship. With this context in mind we changed the set-up of the experiment on various points, although the basis is still the inspection game.

In contrast to the previous experiment, both punishing and rewarding are now at the discretion of the inspector (from now on called employer) and costly for the employer, while just as in the previous experiment, punishing reduces the payoff of the inspectee (from now on called worker) and rewarding increases the payoff of the worker. In each treatment, we use a cost/effect ratio of either 1:1 or a ratio of 1:3. A cost/effect ratio of $1 : x$ means that a punishment [reward] that costs the employer 1, costs [contributes to] the worker x. Another difference is that employers and workers stay matched in the same pair for all rounds during the experiment. Finally, if the employer decides to inspect, an extra stage is added, in which the employer can choose either to punish the worker, reward the worker, or do nothing.

The literature gives us some indications for what we could expect to happen, but the literature is not conclusive. In the psychological literature, Skinner (1965) concludes from experiments on animals that unlike rewarding, punishing has no lasting effect. Furthermore, psychologists find that supervisors rewarding good behavior perform better in inducing hard work than supervisors punishing bad behavior (Sims, 1980; Podsakoff, Bommer, Podsakoff, and MacKenzie, 2006; George, 1995). However, this research is based on questionnaires which makes identifying cause and effect difficult.

In experimental economics, studies have investigated the strength of negative and positive reciprocity (Abbink, Irlenbusch, and Renner, 2000; Brandts and Sola, 2001; Charness and Rabin, 2002; Offerman, 2002; Brandts and Charness, 2004; Falk, Fehr, and Fischbacher, 2003; Charness, 2004; Al-Ubaydli and Lee, 2009). These studies found no or weak evidence for positive reciprocity, which would undermine the idea that workers would react positively to rewards. Although they found stronger evidence for negative reciprocity, it is difficult to draw conclusions from this evidence. On the one hand workers are perhaps more eager to avoid punishment, but on the other hand we could perhaps expect a negative spiral of punishment, less working, more punishment etc.

In our experiment, we see in general stronger results for treatments with a cost/effect ratio of 1:3 compared to those with a cost/effect ratio of 1:1, and we will further focus on the treatments with a cost/effect ratio of 1:3. We compare single instrument treatments in which employers have just one instrument, either punishment or reward, and the baseline treatment where they have no instrument at all. For the single instrument treatments, we find that the workers work more compared to the baseline treatment. Moreover, it does not matter whether the instrument is punishment or reward. With respect to inspection, we see less costly inspections in the punishment-only treatment compared to the baseline and the reward-only treatments. Therefore, whith just one instrument available, the punishment-only treatment increases the payoff of the employer most.

We could expect that the payoff of employers in a treatment with both instruments would be at least as high as it is in a treatment with only punishment: employers could just ignore the reward instrument. This however, is not the case. In the two instrument treatment, reward is used more often than punishment and in a questionnaire, participants in both roles (employer and worker) state that rewarding good behavior is more appropriate than punishing bad behavior. In the two instruments treatment, workers work as much as in the single instrument treatments, but what makes the punishment-only treatment more profitable for the employer than other treatments, is the fact that the employer needs fewer inspections. We conclude that only adding the possibility to punish to the baseline is most profitable for the employer, but when the possibility to reward is also added, the positive effect seems to decrease.

In Chapter 5, we deal with the question how an auctioneer could prevent winners in an auction from going bankrupt afterwards. The context is one in which winners have to file for bankruptcy, if it turns out that the value of the object is less than the price paid for it. Bankruptcy may be very undesirable in the case of license auctions where a government sells the right to exploit radio frequencies and where bankruptcy of the operator would interrupt communication via those frequencies or decrease competition. Another situation where post auction bankruptcy may be undesirable is when a government selects a (critical) supplier, using a procurement auction.

The problem of post auction bankruptcy is widespread in practice. An extreme example is the 1996 C-Block auction by the Federal Communications Commission in the US: all major bidders (winning bids $10.2 billion in total) went bankrupt (Zheng, 2001). Governments have used various methods to overcome the bankruptcy risk. The literature mentions for example: surety bonds, a kind of third party guarantee (Calveras, Ganuza, and Hauk, 2004), multi-sourcing, where bidders can only win part of the contract (Engel and Wambach, 2006) and finally the average bid auction, where the winner is the one with a bid closest to the average (Decarolis, 2010). We analyze whether a simple choice of auction type could mediate the problem. In a laboratory experiment, we compare the English auction[4] and the first-price sealed-bid auction[5], two auction types that are used frequently to sell licenses and to procure goods and services.

Our experimental design is a straightforward implementation of the problem. Half of the participants take part in English auctions and the other half in first-price sealed-bid auctions. For each auction, the common value of the object is the sum of three numbers (signals) randomly generated. Each auction has three participants and each participant receives one of the signals, but is not informed about the value of the other signals. For each of the treatments, in half of the auctions if participants make a loss, they go bankrupt, and only incur a minimal cost. In the other half of the auctions participants have to cover their full losses.

[4] In the English auction, the auctioneer increases a counter indicating the price of an object. Each bidder can step out of the auction by stopping the counter. The other bidders are informed about the price where this bidders steps out,and the counter restarts from that point. The last bidder who remains in the auction wins the object and pays the price where the penultimate bidder steps out.

[5] In the first-price sealed-bid auction, all bidders simultaneously submit a bid. The highest bidder wins and pays a price equal to her own bid.

The literature gives us some intuition about what to expect. Klemperer (2002) states for example that bidders that can go bankrupt, will bid more aggressive as the downside risk is capped by the bankruptcy option. However, the literature is inconclusive with respect to the question which auction-type will perform better. In case of auctions with a common value studied here, we can expect higher winning bids and therefore more bankruptcy in the English auction than in the first-price auction (Milgrom and Weber, 1982). However, in English auctions bidders know when the other bidders step out of the bidding and they could use this information to make a more informed guess about the true value of the object and therefore overcome the bankruptcy risk. We find that when bankruptcy is a possibility, in auctions of both types more bidders make losses than in the unlimited liability case. This increase is not significantly different between both types of auction formats. The result contradicts the predictions of a Nash equilibrium analysis. Eyster and Rabin's (2005) "cursed equilibrium" model explains our findings quite well. We conclude that a choice of either the English or the First-Price auction does not overcome the bankruptcy problem and that the cursed equilibrium model helps to explain this.