

Will Green Taxes Undermine Moral Motivation?

Karine Nyborg¹

Department of Economics, University of Oslo

P.O.Box 1095 Blindern, NO-0317 Oslo, E-mail: karine.nyborg@econ.uio.no

Abstract

Both psychologists and economists have shown that under certain circumstances, economic incentives seem to undermine intrinsic or moral motivation. If such crowding-out effects are sufficiently strong, the effect of green taxes might even be counterproductive and increase, rather than decrease, environmental problems. However, moral motivation can also be reinforced by economic incentives, a fact which has received surprisingly little attention in the literature. In particular, reciprocal preferences and conditional cooperation may cause green taxes to support, rather than undermine, individuals' moral motivation to contribute to a better environment.

Green taxes and intrinsic motivation

For many years, environmental economists have recommended market-based instruments for environmental policy – mainly green taxes, subsidies, and tradable emission permits – claiming that such policy instruments are typically more efficient than command-and-control measures like prohibitions and mandatory standards (see, e.g., Cropper and Oates, 1992). More recently, however, some researchers have expressed doubt concerning the efficiency of green taxes, pointing out that monetary incentives may potentially undermine individuals' moral or intrinsic motivation to contribute to a better environment (see, in particular, Frey 1994, 1997, 1999). If such crowding-out effects are present, they might reduce, or potentially even reverse, the positive environmental impacts of green taxes.

In the present paper, I discuss arguments both in favor of and against the idea that green taxes undermine intrinsic motivation for environment-friendly behavior. Since arguments in both directions do exist, the question of which effects are strongest is ultimately an empirical one, which will not be addressed here. Nevertheless, since the debate may be rather confusing for policy makers trying to judge whether or not to use green taxes, my aim is to provide a reasonably systematic – although far from exhaustive – discussion of arguments I find relevant, as well as important, for this issue.

¹ This paper is partly based on my presentation “Do Green Taxes Undermine Moral Motivation?” (presented in Norwegian) at a conference organized by the Department of Economics, University of Oslo, for the Standing Committee on Finance and Economic Affairs and the Standing Committee on Energy and the Environment of the Norwegian Parliament (the Storting), Oslo, January 22, 2008. Thanks to the Research Council of Norway for funding through the RAMBU programme. Thanks also to the journal's reviewers for their clarifying comments.

As a background for the discussion, I will begin with a brief summary of certain quite robust findings from the recent literature in experimental economics. I will then discuss the predictions concerning motivational crowding-out effects of green taxes from, respectively, cognitive evaluation theory, the self-image/duty-orientation approach, and theories of cognitive dissonance and reciprocity/conditional cooperation.

Below, I will focus on internalized moral norms, that is, rules of behavior enforced *by the individual herself* through psychological mechanisms such as bad conscience, a warm glow of giving, or the desire for a decent self-image. Specifically, I will not discuss models of *social norms*, by which I mean rules of behavior enforced by others through social approval, social reproach and the like. In practice, social and moral norms are intertwined; it is the need to narrow my scope that prevents me from including both here. Furthermore, since the topic is that of moral motivation, my discussion will necessarily be limited to cases where individuals are aware that a moral issue – here, environmental protection – is at stake; firms' behavior, as well as consumer behavior in cases where, although taxes increase market prices, the consumer is unaware about the potential environmental effects of her actions, thus falls outside the scope of my discussion.

Moral motivation: Some relevant experimental results

The simplest textbook model of human behavior, the Homo Oeconomicus model, assumes that individuals care only about their own consumption of public and private goods. Many environmental goods, such as a stable global climate and the ozone layer, can be considered public goods; but although Homo Oeconomicus may well care for these goods, he will contribute voluntarily to their provision only to the extent that his private benefits of doing so exceeds his private costs. In large economies, this means that only a few who are very rich, or who care very strongly about the public good, will contribute, while the rest of the population are free-riders (Andreoni 1988)². Substantial evidence, however, documents that people do not conform to these predictions – that is, at least not all people, all of the time. The Homo Oeconomicus model can hardly explain, for example, why so many people bother to recycle their household waste even in the absence of economic incentives (see, for example, Bruvold et al. 2002), or why some people pay more for eco-labeled goods (Björner et. al, 2004). In particular, the recent literature on experimental economics has confirmed a substantial willingness to contribute to public goods, and to share with others, even strangers; however, it has also demonstrated that people punish each other much more, and react more strongly when disappointed, than Homo Oeconomicus would.³ In this section I will briefly review some relevant experimental findings illustrating that the Homo Oeconomicus model is not always sufficiently sophisticated for studying the question at hand here.

² Andreoni's result is based on what he calls "pure altruism", a model where the individual gets utility from own consumption and from a public good. Although Andreoni seems to have in mind that the preference for the public good involves a concern for *others'* access to this good, his model is formally equivalent to the Homo Oeconomicus model.

³ For a good, although not quite updated overview, see Kagel and Roth 1995; Colin Camerer's 2003 book, extensively cited below, can be highly recommended.

In an economic experiment, individuals make decisions within a controlled environment, and are paid in real money according to their choices. Usually, all decisions are made anonymously, excluding social sanctions or social rewards as an explanation for subjects' choices. One of the simplest games played in the laboratory is the *Dictator Game*: Subject *A*, the dictator, receives a sum of money, for example \$10. His task is to decide how to share this sum with another player *B*, whose role is purely passive. The typical finding is that dictators give an average of about 20 percent of their initial endowment to the recipient (Camerer, 2003); giving nothing and giving 40-50 percent are the most common strategies.

The *Ultimatum Game* is similar to the Dictator Game, except that player *B* can now accept or reject *A*'s offer; if *B* rejects, both get nothing. If *B* is like Homo Oeconomicus, he will never reject a strictly positive amount. Knowing this, *A* will, if he is also like Homo Oeconomicus, propose the lowest strictly positive amount possible; and *B* will accept.⁴ Nevertheless, a large number of experiments have established that the mean proposal is usually about 40 percent; that offers of 50 percent are very common; and, that if *A* offers less than 20 percent of his experimental endowment, *B* rejects in about half of the cases (Camerer 2003). These results are surprisingly robust to methodological variations in the experiments, such as the size of the stakes. Ultimatum game experiments have been conducted in a large number of societies with widely differing cultures (see, in particular, Heinrich et al., 2001); and although behavior does vary between cultures, no society seems to fit the Homo Oeconomicus predictions.⁵ The real puzzle here is not that proposers share so much (if responders are likely to reject small offers, substantial sharing can be rational even for Homo Oeconomicus); the puzzle is that responders reject strictly positive offers. One explanation which seems consistent with a substantial number of evidence is that of reciprocal preferences (Rabin 1993); I will return to this below.

In a *Public good game*, every subject is member of a small group (often 3-4 individuals), but without knowing the identity of the other group members. Each subject gets an initial endowment and is asked to share this money between herself and the group. All contributions to the group are multiplied (by the experimenters) by a factor larger than one, but smaller than the size of the group, and then shared equally between group members. A Homo Oeconomicus subject will thus prefer others in his group to contribute a lot, but will keep his own endowment for himself. The unique Nash equilibrium in a public good game with Homo Oeconomicus preferences is, accordingly, that nothing is contributed to the group. The typical finding from public good game experiments, however, is that in one-shot games, and in the first round of repeated games, subjects contribute on average 40-60 percent of their initial endowment (see, e.g., Ledyard, 1995; Fischbacher et al., 2001). Here, too, results appear to be robust with respect to stakes (Kocher et al., 2008). These numbers conceal a huge variation, though: Most studies find that a substantial proportion of subjects contribute nothing, quite a few contribute everything, while most contribute something in between

⁴ The following set of strategies is also a subgame perfect Nash equilibrium in this game: *A* offers zero, *B* accepts all offers.

⁵ In fact, some data indicates that the individuals who behave mostly like Homo Oeconomicus are small children, 3 – 4 years old (Fehr et al., 2008).

(e.g., Fischbacher et al., 2001). When a public good game is repeated several times, even if subjects are members of a new group in every round, average contributions tend to start off high, but then decline by rounds (Ledyard, 1995; Fehr and Gächter, 2002). After 10 rounds, average contributions may be as low as around 10 percent or even less (Fehr and Gächter 2000).

Substantial evidence is consistent with the hypothesis that this decline occurs, to a considerable extent, as a reaction from individuals to others' low contribution. If subjects are (without knowing so) matched into groups with high contributors only, contributions stay high (Gächter and Thöni 2005, Gunnthorsdottir et al. 2007). Typically, most studies find that a moderate share of subjects behave as free-riders; a very small share are unconditional cooperators, who contribute regardless of what others do; while a substantial share, often a majority, are conditional cooperators, meaning that their contribution is increasing in others' contributions (Fischbacher et al. 2001, Fischbacher and Gächter 2006).

A further interesting aspect of repeated public good games is the following. Some researchers have offered subjects the option to punish others: After one round has been played, subjects get to see an (anonymous) list of others' contributions in their group. Before the next round is played, each subject can punish whomever she wishes by giving up (for example) one unit of her own endowment; this will reduce the payoff of the punished person by (for example) three units. If two subjects will never be members of the same group twice, and this is known to all, Homo Oeconomicus would never punish anyone: Punishing is costly, and gives no pecuniary benefits to the punisher. Thus, the option to punish others should have no effect. Nevertheless, in experiments one finds that people punish low contributors to a considerable extent; and rather than decline by rounds, average contributions tend to stay high, or even increase up to one hundred percent (Fehr and Gächter 2000, 2002). Again, the puzzle is not really why people contribute, since the threat of being punished may make this perfectly rational for a self-interested subject; the puzzle is why anyone would bother to punish others.

The above might seem entirely unrelated to the issue of green taxes. However, I will soon return to the issue of reciprocity and conditional cooperation, explaining why I find this highly relevant for the current discussion.

Why economic incentives can undermine moral motivation

Although substantial experimental evidence does indicate that people contribute more to the common good than implied by the Homo Oeconomicus model, there is also evidence that such willingness may be crowded out by economic incentives. In a much quoted study, Gneezy and Rustichini (2000a) found that imposing a fine on parents arriving late to collect their children at day care *increased* the number of late-coming parents; Gneezy and Rustichini (2000b) found that among high school students collecting donations for a charitable cause, those who were paid by a percentage of their collections (financed by the researchers) collected *less* than those who were unpaid. Mellström and Johannesson (2008) found in a field experiment that women who were offered a monetary compensation for donating blood donated less than half of those who were not

offered monetary compensation.⁶ Results indicating crowding-out of moral motivation by economic incentives have also been reported by, e.g., Frey and Oberholzer-Gee (1997) and Brekke et al. (2003).

What are the reasons that economic incentives, such as environmental taxes, might undermine moral motivation? Here, I will focus on two related, but different strands of literature; cognitive evaluation theory, and theories relating to the concept of self-image.

Cognitive evaluation theory

Cognitive evaluation theory (Deci and Ryan, 1985) comes from the social psychology literature and is backed by a large amount of experimental evidence. After its introduction to economics, particularly pioneered by Bruno Frey (see Frey 1997; Frey and Jegen 2001), this theory has become influential in the economic literature on moral motivation.

Cognitive evaluation theory is concerned with intrinsically motivated behavior, that is, things the individual would have wanted to do even if no external incentives were present. Many of the psychological experiments underlying cognitive evaluation theory are based on maze games, puzzles and the like (see, e.g., Deci and Ryan 1985). A frequently observed phenomenon is that if subjects are given external incentives, for example in the form of monetary payment for good performance, they will improve their performance if the payment is substantial enough; however, they seem to lose interest in the task, and if the payment is subsequently abolished, performance does not return to its initial level.

This argument implies that the effect of economic incentives can be non-monotonic: While a sufficiently high economic incentive can improve performance, a low incentive may reduce performance, since the increased extrinsic motivation is insufficient to outweigh the decreased intrinsic motivation. For example, if your kids are happy to help out with washing the dishes, and you start paying them, they may actually refuse to help.

One aspect of cognitive evaluation theory which has been little discussed in economics, however, is its substantially context-specific predictions. As explained in Deci and Ryan (1985), an extrinsic incentive does not always undermine intrinsic motivation. Rather, it will do so if the incentive is perceived by the individual as *controlling*; if the incentive is, on the other hand, perceived as *acknowledging*, it will typically *reinforce* the intrinsic motivation. That is, if you start paying your child to wash the dishes and tell him that “I do this to improve your behavior, I’m fed up with you being so slow and reluctant” you should expect the incentive to work only if the payment is substantial enough to outweigh the negative effect you will create in the child’s intrinsic motivation. However, if you say that “this is to show that I appreciate your kind and efficient help”, you can expect the child’s intrinsic motivation to increase.⁷

⁶ Perhaps surprisingly, though, men’s blood donations did not differ significantly between treatments in this study.

⁷ Note that cognitive evaluation theory was initially developed for analysis of intrinsically interesting tasks, not *moral* motivation. To my knowledge, little has been done to explore whether intrinsic moral motivation and the motivation to perform intrinsically interesting tasks respond equivalently to extrinsic (external) motivation.

Applied in the environmental context, some researchers have claimed on theoretical grounds that a green tax will undermine people's moral motivation to be environment-friendly (Frey, 1999). Nevertheless, the above indicates that this will only occur if the tax is perceived as controlling by the public. If the tax is perceived as acknowledging or supportive of the public's moral motivation, the reverse should be expected to hold; that is, the tax may well reinforce the intrinsic moral motivation to behave environment-friendly. I will return to this below.

Self-image approaches

A slightly different reason that economic incentives may have counterintuitive effects on moral decisions has been formalized by Benabou and Tirole (2006).⁸ They assume that the individual does not know her own moral values perfectly. Although values are put to test whenever the individual makes actual decisions, actions are more easily remembered than values; thus, the individual learns how intrinsically "good" she is by observing her own behavior. Thus, one benefit of contributing to a better environment is that it increases the individual's confidence in her own moral goodness.

Now, if an economic incentive to contribute is imposed, and the individual chooses to contribute, she may be unsure of why she did so. Was it because she is intrinsically good, or simply because contributing was economically profitable for her? In brief, the economic incentive destroys the signaling effect to oneself of doing good deeds. A related argument is found in Ellingsen and Johannesson (2008).

Brekke et al. (2003) assume that people have a preference for a self-image as a socially responsible individual, and that a better self-image is achieved through striving towards the individual's idea of the morally ideal behavior. They assume, moreover, that the morally ideal behavior is *that behavior which would (hypothetically) have maximized social welfare, if chosen by everyone*. If individuals were identical this would correspond to using Kant's categorical imperative as the criterion for determining the morally ideal behavior. In their model, however, actual behavior is determined by weighing the concern for a better self-image against other interests of the individual, such as consumption and leisure. In a model like this, policy may influence individual contributions to public goods in three ways: through changing relative prices and through the opportunity set, like in any other economic model; but also through influencing individuals' perception of the morally ideal action.

The reason why policy may influence the morally ideal action is the following. When considering the question, "What would happen if everyone acted just like myself?", the individual will typically take public policy as exogenously given. If the public sector takes good care of the elderly, for example, the individual might think that the morally ideal way to spend her time is to look after her children, not her elders; if the government leaves the elderly to themselves, the individual may find it morally superior to care more for her elders and less for her kids (or her job). Brekke et al. (2003) show that this can, under certain conditions, lead to counterintuitive effects of economic incentives.

⁸ The model of Benabou and Tirole (2006) is broadly formulated, and can also be interpreted in terms of social reputation rather than self-image.

Let me illustrate by means of an example. Consider a person who purchases, each time she travels by air, climate tickets offsetting the climate emissions from her flight. Then, imagine that a global climate tax, substantial enough to keep global climate gas emissions at a socially optimal level, is introduced (that is, a tax which, according to the individual's own beliefs, is a Pigou tax). Will this person continue to purchase climate tickets after the tax has been implemented? Probably not. With the tax, her voluntary contributions have become superfluous, or at least less socially valuable; this will reduce, or remove, her motivation to contribute voluntarily. Hence, in this case the moral motivation to contribute is, indeed, crowded out by the environmental tax.

Now, is this an argument against the use of environmental taxes? I would say hardly. In the example of climate tickets, moral motivation for voluntary contribution decreases because voluntary contributions are not important anymore. There is nothing in this argument implying that relying on moral motivation is better for the environment than using an environmental tax. On the contrary, the self-image model of Brekke et al. (2003) implies that voluntary contributions will always be strictly lower than the first-best; with a Pigou tax, on the other hand, one could in fact reach the first-best. Thus, while moral motivation of this kind may be crucial for environmental protection in cases where taxes cannot be used, it is much less clear that relying on people's moral motivation is preferable (from the environment's point of view) when environmental taxes are feasible and efficient.

Why economic incentives can support moral motivation

The possible crowding-out effects of intrinsic motivation by economic incentives have been thoroughly discussed in the recent economics literature. Possible crowding-in effects, however, are less mentioned; perhaps because empirically, such effects may be hard to disentangle from the usual economic incentive effect of a tax. Nevertheless, there are reasons for an environmental tax to support, rather than crowd out, moral motivation.

Cognitive dissonance

One of the best documented phenomena in modern psychology is that of *cognitive dissonance*; the fact that we experience a negative feeling or drive when our behavior is at odds with our usually positive self-conception as good or smart persons (Aronson et al, 2005). The self-image approach by Brekke et al. (2003), discussed above, can be considered as an economic model of cognitive dissonance, but in which normative views are fixed and cognitive dissonance is reduced only through adapting one's actual behavior. In the psychology literature, however, much emphasis is placed on the other possible way to reduce the unpleasant experience of cognitive dissonance; namely changing one's values rather than behavior.

Based on this tradition, Östling (2009) assumes that the more costly it is to keep a moral ideal, the larger is the individual's tendency to subjectively reduce his perception of the moral value of this ideal. If it becomes more costly to drive an electric car, for example, the individual will be more prone to rethink his moral values, searching for reasons why driving an electric car is not so morally superior after all. This produces a positive reinforcement mechanism between moral motivation and economic incentives: By imposing an environmental tax, the government makes it relatively less

costly to be environment-friendly, thus making individuals more prone to keep, or reinforce, their environment-friendly moral values.

Reciprocity and conditional cooperation

The final argument I want to discuss, albeit not the least important one, is related to *reciprocal preferences* (Rabin 1993). This concept is used in slightly different ways by different writers, but a common definition is *a preference for repaying good intentions by good deeds, and for repaying bad intentions by bad deeds*. Note the word “preference”; reciprocity, defined this way, is very different from, for example, using a tit-for-tat strategy, which can be an equilibrium strategy in a dynamic game even for Homo Oeconomicus.

Roughly, one may say that a person with reciprocal preferences prefers to be good to the good and bad to the bad. If I cheat you, you may feel bad about it; you may be able to reduce your bad feelings, however, by cheating me back. Conversely, if I treat you kindly, and you have reciprocal preferences, you will feel bad if you treat me badly in return.

There is, by now, a vast body of experimental evidence indicating that reciprocal preferences are widespread. The high share of low offers being rejected in ultimatum games, discussed above, is one example; the prevalence of punishment in repeated public good games is another. The relevance of this phenomenon in the present context is that reciprocal individuals tend to be conditional cooperators: they prefer to contribute if others do so, but prefer not to contribute if others are not contributing. Fischbacher et al. (2001), for example, found that about 30 percent of their subjects were free-riders, behaving consistently with the Homo Oeconomicus model; more than 50 percent were conditional cooperators, while *none* were classified as unconditional cooperators. In a study by Fischbacher and Gächter (2006), 55 percent were classified as conditional cooperators, while only 20 percent were classified as “egoists”.

Conditional cooperation may be linked more explicitly to moral motivation via the self-image model by Brekke et al. (2003) outlined above. Assume again that the individual has a preference for a good self-image, and that self-image is better the closer the individual is to her morally ideal contribution. Assume, however, that the morally ideal contribution is not (or at least not only) based on the kantian-style reasoning discussed above, but is, instead, conditional: One’s moral obligation towards others is limited by what those others contribute (as proposed by Sugden 1984). If others contribute a lot, I feel a moral obligation to contribute too; but I have no obligation to let myself be exploited by others, and I may even have a moral obligation to punish others who do so (Fehr and Gächter, 2002). Then, the moral obligation to contribute is decreasing in others’ contribution.

Indeed, Hauge (2007) asked subjects in a repeated public good game what they found to be the morally ideal contribution for themselves in that game. A majority of her subjects responded with a conditional ideal; more precisely, the morally ideal contribution reported by most subjects equaled, roughly, the average contribution by others in their group. If not only actual behavior, but even one’s perceived moral obligation, depends on others’ behavior, then the *mere presence of free-riding behavior* will undermine moral motivation.

Let us return to the question we started out with in this section: Why might an environmental tax *support* moral motivation? The reason should, by now, be clear: If too many free-riders are out there, they will demoralize the conditional cooperators. That would not be important if the conditional cooperators were few; however, the experimental evidence indicates that they are many, possibly a majority.

In any population, individuals will be different. The free-riders, behaving (at least in the present context) consistently with the Homo Oeconomicus predictions, are not the only ones out there. They are, however, present. The only language Homo Oeconomicus understands and responds to is that of economic incentives, for example environmental taxes. Green taxes, thus, regulate the otherwise demoralizing behavior of free-riders; by forcing everyone to contribute at least a little, a green tax may keep up the moral motivation of the conditional cooperators. A green tax may thus well be interpreted as supporting, and acknowledging, conditional cooperators' moral motivation. A more systematic study of this idea, possibly in the context of a laboratory experiment, remains, however, to be done.

Conclusions

Will green taxes undermine moral motivation? Perhaps, but not necessarily; on the contrary, moral motivation may, in fact, be strengthened by a tax.

If the green tax is perceived as *controlling* by the public, cognitive evaluation theory predicts that intrinsic motivation *will* be undermined. Further, if a tax is perceived by the public as solving the environmental problem, the self-image approach also predicts that moral motivation will be undermined. In the latter case, however, one may ask whether such crowding-out is a problem; after all, if the problem *is* solved by the tax, there is not necessarily much to worry about (unless the tax is perceived as sufficient while in fact it is not).

An environmental tax applies in the same way for everyone, regardless of their motivation. A pollution tax, if sufficiently enforced, makes every polluter pay; complete free-riding is no longer possible. There is by now substantial experimental evidence that a large fraction of the population, perhaps a majority, are conditional cooperators. Conditional cooperators are, at the outset, willing to contribute to public goods, but withdraw their contributions if too much free-riding is taking place. By preventing free-riding, an environmental tax can constitute an important support to keep up conditional contributors' motivation.

Interpreted this way, a tax can well be seen as acknowledging and supportive of the population's moral motivation. If, moreover, this is the interpretation taken by the public, then cognitive evaluation theory, which has been used to argue that economic incentives *undermine* moral motivation, predicts that a green tax will *reinforce* intrinsic motivation.

Finally, it should be noted that while morally motivated individuals may contribute substantially for the sake of a better environment, even those with a strong moral motivation cannot be expected to contribute enough to bring the economy to its first-best situation (Brekke et al., 2003). That is, moral motivation is typically insufficient to solve environmental problems in an efficient way. When a green tax is feasible, this instrument *is* typically able to bring the economy to its first-best.

A caveat is in order. Most of the evidence cited above comes from small-scale experiments performed in rather artificial contexts. Although an impressively large number of experiments have achieved similar results, we do not yet have sufficient knowledge concerning how people respond to similar situations in economies with large-scale, even global, public goods. This reservation applies, however, both to much of the crowding-out and crowding-in arguments presented above.

References

- Andreoni, J. (1988): Privately Provided Public Goods in a Large Economy: The Limits of Altruism, *Journal of Public Economics* **35**, 57-73.
- Andreoni, J. (1990): Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving, *Economic Journal* **100**, 464-477.
- Aronson, E., T.D. Wilson, and R.M. Akert (2005): *Social Psychology*. Fifth Edition. New Jersey: Pearson Education International.
- Benabou, R., and J. Tirole (2006): Incentives and Prosocial Behavior, *American Economic Review* **96**(5), 1652-1678.
- Bjorner, T.B., L.G. Hansen, and C.S. Russell (2004): Environmental labeling and consumers' choice. An empirical analysis of the effect of the Nordic Swan, *Journal of Environmental Economics and Management*, **47**(3), 411-434.
- Brekke, K. A., S. Kverndokk, and K. Nyborg (2003): An Economic Model of Moral Motivation, *Journal of Public Economics* **87** (9-10), 1967-1983.
- Bruvoll, Annegrete, Bente Halvorsen and Karine Nyborg (2002): Households' recycling efforts. *Resources, Conservation and Recycling* **36**, 337-354.
- Camerer, Colin (2003): *Behavioral Game Theory. Experiments in Strategic Interaction*, pp. 43-113, 2003. Princeton University Press/Russell Sage Foundation.
- Cropper, M., and W. Oates (1992): Environmental Economics: A Survey, *Journal of Economic Literature*, June.
- Deci, E. L., and R. M. Ryan (1985): *Intrinsic Motivation and Self-Determination in Human Behavior*, New York: Plenum Press.
- Ellingsen, T., and M. Johannessen (2008): Pride and Prejudice. The Human Side of Incentive Theory, *American Economic Review* **98** (3), 990-1008.
- Fehr, E., H. Bernard and B. Rockenbach (2008): Egalitarianism in young children, *Nature* **454**, 1079-1084.
- Fehr, E., and S. Gächter (2000): Cooperation and punishment in public goods experiments, *American Economic Review* (90), 980-994.
- Fehr, E., and S. Gächter (2002). Altruistic Punishment in Humans, *Nature* **415**, 137-140.

- Fischbacher, Urs, Simon Gächter and Ernst Fehr (2001): Are People Conditionally Cooperative? Evidence from a Public Good Experiment, *Economics Letters* (71), 397-404.
- Fischbacher, U., and S. Gächter (2006). Heterogeneous social preferences and the dynamics of free riding in public goods, Cedex, University of Nottingham.
- Frey, B. S. (1994): Pricing and regulating affect environmental ethics, *Environmental and Resource Economics* 2, 399-414.
- Frey, B. S. (1997). *Not just for the money. An economic theory of personal motivation*. Cheltenham, UK: Edward Elgar.
- Frey, B.S. (1999): Morality and Rationality in Environmental Policy. *Journal of Consumer Policy* **22**, 395-417.
- Frey, B.S., and R. Jegen (2001): Motivation crowding theory, *Journal of Economic Surveys*, 589-611.
- Frey, B. S., and F. Oberholzer-Gee (1997): The Cost of Price Incentives: An Empirical Analysis of Motivation Crowding-Out, *American Economic Review* **87** (4), 746-755.
- Gächter, S. and C. Thöni (2005): Social Learning and Voluntary Cooperation among Like-Minded People, *Journal of the European Economic Association* **3**(2-3), 303–314.
- Gneezy, Uri, og Aldo Rustichini (2000a): Pay Enough or Don't Pay at All, *Quarterly Journal of Economics* 115(3), 791-810.
- Gneezy, U., and Rustichini, A., (2000b): A Fine is a Price, *The Journal of Legal Studies* 29 (1), part 1, 1-17.
- Gunnthorsdottir, A., D. Houser, and K. McCabe (2007): Disposition, History and Contributions in Public Goods Experiments, *Journal of Economic Behavior and Organization* **62**(2), 304–315.
- Hauge, Karen Evelyn (2007): Normative Reflection in Public Good Experiments, unpublished note, Department of Economics, University of Oslo.
- Heinrich, J., R. Boyd, S. Bowles, C. Camerer, E. Fehr, H. Gintis, and R. McElreath (2001): In Search of Homo Economicus: Behavioral Experiments in 15 small-scale societies, *American Economic Review* **91** (2), 73-78.
- Kocher, M.G., P. Martinsson and M. Visser (2008): Does Stake Size Matter for Cooperation and Punishment? *Economics Letters* **99**, 508-511.
- Ledyard, J.O. (1995): "Public Goods: A Survey of Experimental Research", in J.H. Kagel and A.E. Roth (eds.): *The Handbook of Experimental Economics*, Princeton: Princeton University Press.
- Kagel, J.H., and A.E. Roth (eds.) (1995): *The Handbook of Experimental Economics*, Princeton: Princeton University Press.
- Mellström, Carl, and Magnus Johannesson (2008): Crowding Out in Blood Donation: Was Titmuss Right? Forthcoming, *Journal of the European Economic Association*.

Östling, Robert (2009): Economic Influences on Moral Values, *The B.E. Journal of Economic Analysis & Policy (Advances)*, 9(1), Article 2.

Stockholm School of Economics: SSE/EFI Working Paper Series in Economics and Finance no. 635.

Rabin, M. (1993): Incorporating Fairness into Game Theory and Economics, *American Economic Review*, **83** (5), 1281-1302.

Sugden, R. (1984): Reciprocity: The Supply of Public Goods through Voluntary Contributions, *The Economic Journal* **94**, 772-787.