

IZA DP No. 7442

Organizations, Diffused Pivotality and Immoral Outcomes

Armin Falk
Nora Szech

June 2013

Organizations, Diffused Pivotality and Immoral Outcomes

Armin Falk

*CENs, University of Bonn
and IZA*

Nora Szech

University of Bamberg

Discussion Paper No. 7442
June 2013

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Organizations, Diffused Pivotality and Immoral Outcomes

This paper studies how organizational design affects moral outcomes. Subjects face the decision to either kill mice for money or to save mice. We compare a Baseline treatment where subjects are fully pivotal to a Diffused-Pivotality treatment where subjects simultaneously choose in groups of eight. In the latter condition eight mice are killed if at least one subject opts for killing. The fraction of subjects deciding to kill is higher when pivotality is diffused. The likelihood of killing is monotone in subjective perceptions of pivotality. On an aggregate level many more mice are killed in Diffused-Pivotality than Baseline.

JEL Classification: C91, D01, D03, D23, D63

Keywords: morality, pivotality, experiment, organization, responsibility

Corresponding author:

Armin Falk
Center for Economics and Neuroscience
University of Bonn
Nachtigallenweg 86
53127 Bonn
Germany
E-mail: armin.falk@uni-bonn.de

1. Introduction

This paper studies the critical role of institutions in affecting moral decisions. In particular, we aim at identifying a specific institutional mechanism, the diffusion of pivotality, and its potential to promote immoral outcomes. Diffused pivotality provides an excuse for engaging in immoral activities simply because an actor may perceive himself as irresponsible for the outcome. There are many mechanisms to lower feelings of pivotality. Possibly one of the simplest mechanisms is to organize people into groups and to implement a group decision rule that allows people to easily think that their decision is not decisive. A design of this type is used, e.g., in the practice of firing squads. Typically these squads consist of a group of executors rather than a single person. This reduces feelings of pivotality, as for an executor it is unclear whether his shot becomes the fatal shot or not. Apparently, this makes it easier to participate in the execution.¹

Diffused pivotality enables a simple “replacement logic” that is ubiquitous in markets. Traders may prefer concluding a trade themselves to letting another trader perform the same transaction, even if outcomes are morally questionable (Sobel (2007)). For example, many people complain about detrimental working conditions including low health or safety standards, suffering of animals that are kept in bad conditions, or environmental damage. At the same time they buy cheap fashion, electronics or food products, favoring the occurrence of outcomes they generally object to. A frequent “excuse” is “if I don’t buy, another buyer will”.

Replacement arguments do not only apply to markets. They also apply to organizations that manage to create an environment where an actor comes to believe that even if he does not engage in a particular activity, someone else will. A particularly striking example for the potential of organizations to produce evil outcomes is the organization of the Holocaust (Darley (1992), Arendt (1963), Lifton (1986)). Lifton (1986) interviewed German doctors stationed in Auschwitz. They were operating in a nightmarish environment with one of their objectives being to “select” between prisoners who would be allowed to live and those who would be gassed right away. Being ordinary doctors this activity was clearly morally terrible and self-contradictory to them. Nevertheless they engaged in the activities. One of the

¹ Moreover, members of firing squads are often randomly issued a gun containing a blank cartridge. This diffuses responsibility about who fired the fatal shot even more.

frequently made justifications for the obvious evil was that the horrible machinery would go on, regardless of whether or not a particular doctor would continue to participate. Arguments of being unable to stop ongoing moral crime were also used as excuses in the Nuremberg Trials for having participated in atrocity under the Nazi Regime (see, e.g., Crawford (2007) and the references therein).

Morality is an elusive term. Its content is culturally determined and time and space contingent. There exists, however, a basic consensus that harming others in an unjustified and intentional way is considered as immoral (Haidt and Kesebir, 2010). It is this notion that motivates our research paradigm. In the experiments subjects were facing the trade-off between receiving money and killing a mouse versus not receiving money and saving the life of a mouse.² Importantly, mice used in the experiment were so-called “surplus” mice, which would all have been killed without our intervention (see Section 2). Subjects were informed about this default in a post-experimental debriefing.

To study the role of pivotality on immoral outcomes we analyze two treatments. The Baseline treatment implements a simple binary choice according to which subjects either receive zero euros for saving the mouse (option A) or 10 euros for killing the mouse (option B). In Baseline subjects are fully pivotal. This condition therefore serves as a comparison benchmark for the Diffused-Pivotality treatment. In the latter, eight subjects simultaneously decide between option A and option B. As in Baseline a subject receives zero euros for choosing option A and 10 euros for choosing option B, irrespective of the other subjects’ choices. If at least one subject chooses option B, however, eight mice are killed. Thus, if a subject believes that it is likely that at least one other subject chooses option B, he may no longer feel pivotal. Low chances of being pivotal provide an excuse to choose option B, as choosing B would unlikely change the outcome but guarantee a payoff of 10 euros.

Note that with a deontological (e.g., Kantian) moral conception, feelings of pivotality should not affect moral decisions: it is the immoral situation per se one should object to, no matter whether one is decisive or not. With a utilitarian conception of morality, however, perceptions of pivotality become crucial. Our main finding indicates that feelings of pivotality can easily be changed by institutional design. The fraction of subjects choosing option B is significantly higher in Diffused-

² The study was approved by the Ethics Committee of the University of Bonn.

Pivotality than in Baseline. Moreover, the likelihood that a subject chooses to kill mice is decreasing in his belief of being pivotal. At the aggregate level all mice were killed in Diffused-Pivotality. From an organizational “efficiency” point of view, killing mice in the latter treatment was also cheaper than in Baseline. The cost for killing one mouse was ten euros in Baseline, but less than six euros in Diffused-Pivotality. These results demonstrate the power of organizations to promote immoral outcomes. Since subjects were randomly assigned to the two conditions we can assess the causal effect of a simple organizational change on the production of killing. The “same” people produce different levels of immoral outcomes when acting in different institutions.

The general issue this paper contributes to is to understand why “ordinary” people endowed with given moral values engage in activities they would generally object to; and how organizational design affects these outcomes. Research in social psychology suggests that people engage in “psychosocial manoeuvres – often aided by the institutions [...], which absolve them from moral responsibility for harmful acts” (Haidt and Kesebir, 2010, p. 812). In other words, moral behavior is malleable and institutions play a crucial role in promoting or limiting such behaviors. Our focus is on the role of institutions in shaping perceptions of pivotality.

Other institutional mechanisms that have been shown to produce “unfair” outcomes in bargaining contexts are delegation or moral “wobble rooms”, as discussed, e.g., in Bartling and Fischbacher (2012), Hamman et al. (2010) and Dana et al. (2007).³ Falk and Szech (2013) analyze moral outcomes in buyer-seller relationships in bilateral and multilateral market situations. A related literature in social psychology concerns the so-called bystander effect (see, e.g., Latané and Darley (1968) or for a recent overview Fischer et al. (2011)). Typical bystander experiments study helping behavior in response to a staged emergency (e.g., the experimenter becomes injured). What sets our study apart is that in our experiment, even if a subject opts for the moral outcome, he remains uncertain whether the moral outcome is implemented or not.⁴ Note also that in our experiment, consequences are real, incentives are exactly specified and the mechanism (pivotality) is explicitly measured.

³ On the effects of institutions on values see also Bowles (1998).

⁴ This is in contrast to the helping decision in a bystander experiment where, clearly, if a subject opts for helping, the person in need receives help. Furthermore, in a standard bystander experiment, while deliberating about helping or not, subjects

The remainder of the paper is organized as follows. Section 2 describes the design of the experiment together with procedural details and hypotheses. Results are presented in Section 3, and Section 4 concludes.

2. Design and Hypotheses

In this section we first discuss our paradigm, before we present treatments, procedural details and hypotheses.

Paradigm: A key challenge in studying immoral outcomes is to define a choice paradigm that involves the harm of others in an unjustified way. We therefore suggest the “mouse paradigm”, which involves the trade-off between killing a mouse and receiving money versus saving a mouse life and receiving no money. In both treatments subjects were explicitly informed that their mouse is a young and healthy mouse, which will live for about two years if the subject saves it. For illustrative purposes, we presented subjects the picture of a mouse on an instruction screen (Figure 1). We guaranteed subjects that mice, if saved, would live in an appropriate, enriched environment, jointly with a few other mice. The experimenters purchased mice that were saved due to subjects’ decisions. They are now kept in an appropriate, enriched environment, precisely as stated in the instructions.

Subjects were also informed about the killing process, in case they decided to kill their mouse. In the instructions they read the following passage: “The mouse is gassed. The gas flows slowly into the hermetically sealed cage. The gas leads to breathing arrest. At the point at which the mouse is not visibly breathing anymore, it remains in the cage for another 10 minutes. It will then be removed.” As part of the instructions subjects also saw a short demonstration video of the killing process. In the video five mice first move vividly in the cage, then they successively slow down. Eventually they die, with their hearts beating visibly heavy and slow.

It is important to stress that the mice used in the experiment were so-called “surplus” mice: These mice were bred for animal experiments, but turned out to be

observe that others do not help either. Feelings of guilt for not helping can hence be shared. In our simultaneous-move set-up, such social learning is impossible. Opting for killing, subjects do not know whether other subjects opt for killing as well. In addition, in a bystander experiment, participants need to realize that their help is required.

unsuited for study. They were perfectly healthy, but keeping them alive would have been costly. It is standard to gas such mice. This is common practice in laboratories conducting animal experiments. Thus, as a consequence of our experiment, many mice that would otherwise all have died were saved. Subjects were informed about this default in a post-experimental debriefing.⁵



Figure 1: Picture of a mouse as presented in the instructions of the experiment.

Treatments. To study the role of pivotality we need to contrast two decision contexts, one where subjects are fully pivotal with one where pivotality is diffused, while keeping financial incentives and expected overall consequences identical. To this end we study two treatments. In **Baseline**, subjects faced a simple binary choice, labeled option A and option B. Option A implied that the mouse would survive and that the subject would receive no money. Option B implied the killing of the mouse and receiving 10 euros. This treatment informs us about the fraction of subjects who are willing to kill the mouse for 10 euros if they are fully pivotal. In **Diffused-Pivotality** subjects were matched in groups of eight. As in Baseline, each subject faced a binary choice option, again labeled option A and option B. Option A implied that a subject would receive no money. If a subject chose option B, he would receive 10 euros. These payoff consequences were independent of other subjects' decisions. All subjects chose simultaneously. If at least one subject chose option B, *eight* mice were killed. Thus the group of eight subjects was facing the same financial incentives and was endowed with the same number of mice as in Baseline. The difference between

⁵ Note that while *perceptions* of the situation may have changed due to this information, *consequences* were exactly the same and as stated in the instructions. In future research it would be interesting to explore whether using an alternative framing would affect decisions in response to institutional changes differently (compare evidence on the so-called omission-commission bias, e.g., Spranca et al. (1991)).

treatments concerns the organizational design. While in Baseline each subject was pivotal, this was not necessarily the case in the Diffused-Pivotality condition.

In both treatments we elicited beliefs right after subjects had taken their decision. To allow for comparison we asked the same question in both treatments. Subjects were asked to assume a group of eight subjects and to answer the question how many other subjects they think had chosen option B. Subjects could enter any number from 0 to 7 and were paid one euro for a correct estimate. In Diffused-Pivotality we also asked subjects to indicate the probability that all other seven group members had chosen option A. In this case their decision would be pivotal for the survival of the mice. Subjects were asked to enter an integer percentage number, i.e., higher percentages indicate a higher likelihood that the subject perceived himself as pivotal.

Procedures. A total of 252 subjects, mainly undergraduate university students from all majors, took part in the experiments. 124 subjects participated in Baseline and 128 subjects in Diffused-Pivotality, respectively. Every subject participated only in one treatment condition. We used z-Tree as experimental software (Fischbacher, 2007); subjects were recruited using the software ORSEE (Greiner, 2003). At the beginning of an experimental session, participants received detailed information about the rules and structure of the experiment. In all treatments, the experiment started only after all participants had answered several control questions correctly.

To reduce possible communication between subjects across sessions, the experiments were run in short time. They took place between May 3 and 4 in 2012 and were conducted in six different rooms at the “Beethovenhalle”, the largest concert hall in Bonn. We set up six parallel, computerized labs in these rooms. Subjects received payments according to the rules of the experiment and an additional show-up fee of 20 euros to compensate for the remote location. In both treatments subjects received their payments in a sealed envelope outside the room where the experiment had taken place. This way, neither other subjects nor the experimenter handing over the envelopes knew what a particular subject had earned. This procedure was explained in the instructions.

To ensure credibility, we stated right at the beginning that all statements made in the instructions were true, as is standard in economic experiments, and that all consequences of subjects’ decisions would be implemented exactly as written in the

instructions. We emphasized orally that the experimenters personally guarantee the truthfulness of the instructions. Subjects were also invited to send us an email if they wanted to discuss the study.

Hypotheses. Our main hypothesis is that the killing rate is higher in Diffused-Pivotality than in Baseline due to simple replacement logic. If subjects in the former treatment believe that other group members may engage in the immoral activity and choose option B, they are no longer necessarily pivotal and may choose option B as well. We further hypothesized that the likelihood of choosing option B should decrease in the likelihood of perceived pivotality.

While intuitive, two important qualifications with respect to the pivotality hypothesis should be made. First, even if pivotality is diluted in Diffused-Pivotality, the consequences from an individual perspective are much more severe as well. While in Baseline subjects decide about the life of one mouse, subjects in Diffused-Pivotality potentially cause the death of eight mice. Second, replacement arguments corresponding to a low perception of pivotality provide a moral excuse only from a consequentialist or utilitarian moral perspective, as advocated for example by Jeremy Bentham and John Stuart Mill. If subjects follow some deontological moral principle (e.g., the Kantian Categorical Imperative), they should stick to their morally preferred option, regardless of the behavior of others. The empirical relevance of these two moral conceptions, which have been the main combatants in occidental moral philosophy for the last centuries, has been studied using the so-called trolley problem put forward by Philippa Foot (see also, e.g., Thomson (1976), Greene et al. 2004). The quandary to be resolved in this problem is to either follow the deontologically warranted option (and not to throw a switch that will divert a trolley and kill one person) or the option preferred from a consequentialist perspective (killing the person to save five others). Evidence from this literature suggests that both moral approaches are empirically relevant and that the extent to which people follow the one or the other largely depends on situational and emotional factors (see, e.g., the distinction between the trolley problem and the footbridge problem, which assigns actors a more active role, leading to more rule based behaviors in an otherwise identical set-up). In sum, while diluted pivotality would suggest a higher killing rate in Diffused-Pivotality than in Baseline, the effect will largely depend on perceptions of the

severeness of the killing (eight vs. one) and on the extent to which subjects follow a utilitarian moral approach.

3. Results

Our first result concerns the comparison of individual killing rates between the two treatments. In Baseline 45.9 percent of subjects chose option B. In Diffused-Pivotality this fraction was 58.6 percent, an increase of almost 30 percent. The difference is significant ($p=0.022$, Two sample test of proportions, one-sided). Thus, a simple organizational change increases the likelihood of killing in a causal and significant way.

If perceptions of pivotality are responsible for immoral behavior in Diffused-Pivotality, we should observe that the likelihood of subjects choosing to kill their mouse should decrease in the belief of being pivotal. This is exactly what we find. Recall that we asked subjects about the probability (in percentage points) that all other group members had chosen option A. Figure 2 shows the fraction of subjects choosing option B depending on this belief. The four categories displayed in Figure 2 are based on quartiles of the belief distribution with respective percentage values of 0-3.5; 3.5-10; 10-35 and 35-100. The figure shows a clear negative relation between subjective perception of pivotality and the likelihood of choosing option B.




Figure 2: Fraction of subjects choosing option B depending on belief of being pivotal.

Figure 3 shows the full distribution of pivotality beliefs. It displays the cumulative distribution of beliefs separately for subjects choosing option A and option B, respectively. The figure shows that beliefs vary strongly between subjects suggesting a considerable amount of uncertainty about pivotality. Moreover, beliefs differ vastly depending on the option chosen. For example, while among those who kill the mouse 31 percent of subjects believe that the likelihood of being pivotal is zero, the corresponding value for those who save the mouse is only 9.4 percent. Likewise, subjective pivotality estimates of below or equal 20 percent hold for 90.7 percent of subjects who decided to kill, while the respective value is only 39.6 percent for those who decided to save mice.




Figure 3: Cumulative distribution of beliefs of being pivotal in the Diffused-Pivotality treatment, separately for subjects choosing option A or B, respectively. *Probability* refers to the estimated probability that all other group members have chosen option A; n = 128.

In light of Figures 2 and 3 it is not surprising that the likelihood of killing depends statistically significantly on perceptions of pivotality. In a linear probability regression where we regress the likelihood of choosing option B on subjective pivotality, the respective coefficient is negative and highly significant. Assuming a linear relationship the coefficient implies that a change in perceived pivotality of 10 percent decreases the likelihood of killing mice by 8.4 percent (see Table 1, column 1)⁶. Figure 3 suggests that especially very low values of perceived pivotality coincide with choosing option B. Therefore we tested for convexity: if we include a squared term of the percentage value, both coefficients are significant at the 1-percent level (see Table 1, column 2). Thus, the likelihood of killing depends strongly on the belief of being pivotal and option B is chosen particularly frequently if subjects perceive themselves as very unlikely to be pivotal. This finding suggests that in particular larger groups should give way to immoral behavior, enabling many subjects to believe that they are most likely not pivotal anyway. The importance of subjective

⁶ We obtain the same results in Probit regressions.

perceptions of pivotality can also be derived from results reported in column 3. In this specification we use subjects' responses to the question how many other subjects they think had chosen option B.⁷ We constructed a dummy, which takes a value of 1 if a subject stated the point belief that no other subject has chosen option B, which would render the subject's decision pivotal. 17 subjects (13.3 percent) stated this belief. Among these 17 subjects only one subject chose option B. The negative and significant coefficient of the dummy implies that relative to subjects with other point beliefs the likelihood of choosing option B is about 61 percent lower for subjects who believe that they are pivotal.

	Kill Mice (1)	Kill Mice (2)	Kill Mice (3)
Pivotality probability	-.0084*** (.0012)	-.0224*** (.0049)	
Pivotality probability squared		.00016*** (.00005)	
Belief that nobody chooses B			-.6078*** (.1174)
Constant	.7910*** (.0478)	.8871*** (.05667)	.6667*** (.0428)
Observations	128	128	128
R-squared	0.2733	0.3207	0.1754

Table 1: OLS regression coefficient estimates, with binary outcome (option B: kill mice vs. option A: save mice) as dependent variable and standard errors in parentheses. Data comes from the Diffused-Pivotality treatment. *Pivotality probability* is the belief that all other group members have chosen option A (in percent). *Belief that nobody chooses B* is a dummy, which takes value 1 for the point belief that all other subjects have chosen option A. *** indicate significance at the 1-percent-level.

We have shown that organizations can promote individual willingness to engage in immoral outcomes. We conclude this section in briefly discussing the overall, aggregate consequences for moral outcomes. While in Baseline 45.9 percent of mice were killed, all mice were killed in all groups of the Diffused-Pivotality treatment.

⁷ Note that we did not elicit a complete belief distribution. This means that even if a subject stated a low number of participants choosing B, there may have been considerable uncertainty about the exact number of participants, i.e., pivotality, when opting for B.

Thus if the goal of an organization is to increase immoral activities it can easily achieve this goal if it introduces replacement and thus limits pivotality. Killing is not only much more pronounced in Diffused-Pivotality, it is also more “efficient” in the following sense: the money needed to kill one mouse is 10 euros in Baseline but only 5.86 euros in Diffused-Pivotality.

4. Concluding Remarks

This paper has shown how a specific organizational feature, diffused pivotality, affects moral transgression both on an individual and an aggregate level. In support of our pivotality hypothesis, diluting pivotality increases individual willingness to engage in morally suspect ways. This effect is the stronger the more pronounced are individuals’ beliefs of not being pivotal. Our findings also suggest that reduced notions of pivotality lower moral concerns. On an aggregate level the effects are fatal in the sense that diluting pivotality leads to a killing rate of 100 percent, rendering the killing more “efficient” from an organizational point of view. While the focus of this paper is to highlight possible negative consequences of organizational design on moral behavior, the reverse inference is of course intended as well. Our findings suggest that if an organization aims at promoting morality it should reduce diffusion of pivotality and instead attribute individual responsibility to its members.

Additional suggestive evidence for the role of pivotality comes from an explorative control treatment. In this treatment everything was identical to Diffused-Pivotality except that in the control treatment, at least *three* subjects had to choose option B in order to kill eight mice. Thus in this treatment a subject was pivotal if at least two other subjects had chosen option B, rather than no other subject as in Diffused-Pivotality. In this control condition the killing rate is 53.1 percent (n=64), i.e., right in between Baseline and Diffused-Pivotality. The difference is neither statistically significant with respect to Baseline nor with respect to Diffused-Pivotality. However, the fact that the killing rate is higher than in Baseline but lower than in Diffused-Pivotality is in line with the pivotality hypothesis. In the control condition, perceived pivotality is lower than 100 percent, as in Baseline, but higher than in Diffused-Pivotality: while the average subjective assessment of pivotality in Diffused-Pivotality is 24.4 percent, the respective value is 32.5 percent in the control condition ($p < 0.001$, Wilcoxon rank sum test, two-sided).

A possible concern in interpreting beliefs, which we explicitly acknowledge, is the potential endogeneity of beliefs, e.g., due to reasons of false-consensus effects (Marks and Miller, 1987)⁸, or driven by subjects trying to justify behavior ex post in an attempt to maintain a positive self-image (Bénabou and Tirole, 2006). Eliciting and interpreting beliefs is notoriously difficult in this respect. To limit this problem we incentivized beliefs about the number of other participants choosing option B in both treatments. These beliefs were quite accurate on average. We did not incentivize the pivotality belief for the following reasons. A possible remuneration would have depended on the distance to a binary outcome; either the subject turned out to be pivotal in his group or not. We felt such a remuneration scheme would have been difficult to understand for many subjects. Furthermore, many subjects may have chosen intermediate pivotality percentages to diversify against the only two possible outcomes of 0 percent and 100 percent. Reassuringly, however, both types of beliefs (number of other participants choosing option B and percentage of being pivotal) are highly correlated (Spearman rank correlation: -0.5214 , $p < 0.001$, $n = 192$). Moreover, all our results are qualitatively the same (also in terms of significance) if we use beliefs about the number of other participants choosing option B instead of using the (more appropriate) pivotality belief (see also results in column 3 of Table 1).

As briefly discussed above, pivotality provides a moral excuse only from a consequentialist or utilitarian moral perspective. In this respect it is interesting to note that in Diffused-Pivotality about 10 percent of subjects chose option A despite holding a belief that the chance of being pivotal is exactly zero. Thus from a consequentialist moral perspective, these subjects had no reason not to choose option B and cashing in 10 euros. Instead, these subjects seem to have followed a deontological moral principle, sticking to their morally preferred option regardless of consequentialist considerations.

⁸ A finding that suggests a limited role of false consensus as source of beliefs is the following. Suppose that due to false consensus subjects who choose option A believe all other subjects choose option A as well. If this were the case, the fraction of subjects who choose A and think all others choose option A should be identical in Baseline and in Diffused-Pivotality. If, however, beliefs of being pivotal have an independent impact on subjects' decisions, the fraction of subjects who choose option B while believing that everybody else chooses option A should be lower in Diffused-Pivotality. This is what we find. 11.1 percent of subjects who hold this belief choose option B in Baseline, but only 5.9 percent do so in Diffused-Pivotality.

References

- Arendt, H. (1963) *Eichmann in Jerusalem: A Report on the Banality of Evil*. Viking.
- Bartling, B., Fischbacher, U. (2012) Shifting the Blame: On Delegation and Responsibility. *Review of Economic Studies* 79, 67-87.
- Bénabou, R., Tirole, J. (2006) Incentives and prosocial behavior. *American Economic Review* 96, 1652-78.
- Bowles, S. (1998) Endogenous Preferences: The cultural consequences of markets and other economic institutions. *Journal of Economic Literature* 36, 75-111.
- Crawford, N. (2007) Individual and Collective Moral Responsibility for Systemic Military Atrocity. *The Journal of Political Philosophy*, 187 – 212.
- Dana, J., Weber, R.A., Kuang, J.X. (2007) Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory* 33, 67-80.
- Darley, J.M. (1992) Social Organization for the Production of Evil. *Psychological Inquiry* 3, 199-218.
- Falk, A., Szech, N. (2013) Morals and Markets. *Science*, 2013, 340, 707-711.
- Fischbacher, U. (2007) z-Tree: Zurich Toolbox for ready-made economic experiments. *Experimental Economics* 10, 171.
- Greene, J.D., Nystrom, L.E., Engell, A.D., Darley, J.M., Cohen, J.D. (2004) The Neural Bases of Cognitive Conflict and Control in Moral Judgment. *Neuron* 44, 389–400.
- Greiner, B. (2003) An online recruitment system for economic experiments (in *Forschung und wissenschaftliches Rechnen. GWDG Bericht 63, Ges. für Wiss. Datenverarbeitung, Göttingen*, pp. 79–93).
- Foot, P. (1967) The Problem of Abortion and the Doctrine of Double Effect. *Oxford Review* 5, 5-15.
- Haidt, J., Kesebir, S. (2010) Morality. *Handbook of Social Psychology*, 5th Edition. chap. 22. Fiske, S., Gilbert, D., Lindzey, G. (Eds.). Wiley.
- Hamman, J., Loewenstein, G., Weber, R.A. (2010) Self-interest through delegation: An additional rationale for the principal-agent relationship. *The American Economic Review* 100, 1826-46.
- Latané, B., Darley, J.M. (1968) Group inhibition of bystander intervention in emergencies. *Journal of Personality and Social Psychology* 10, 215-21.
- Lifton, R.J. (1986) *The Nazi Doctors: Medical Killing and the Psychology of Genocide*. Basic.
- Marks, G., Miller, N. (1987) Ten years of research on the false-consensus effect: An empirical and theoretical review. *Psychological Bulletin* 102 (1): 72-90.
- Sobel, J. (2010) Do markets make people selfish? Discussion Paper, Economics Department, University of California.
- Spranca, M., Minsk, E., Baron, J. (1991) Omission and Commission in Judgment and Choice. *Journal of Experimental Social Psychology* 27 (1): 76-105.
- Thomson, J. (1976) Killing, Letting Die, and the Trolley Problem. *The Monist* 59, 204-17.