



**University of
Zurich** ^{UZH}

University of Zurich
Department of Economics

Working Paper Series

ISSN 1664-7041 (print)
ISSN 1664-705X (online)

Working Paper No. 177

**Job mission as a substitute for monetary incentives:
experimental evidence**

Lea Cassar

October 2014

JOB MISSION AS A SUBSTITUTE FOR MONETARY INCENTIVES:
EXPERIMENTAL EVIDENCE

LEA CASSAR*

October 28, 2014

Abstract

Are monetary and non-monetary incentives used as substitutes in motivating effort? I address this question in a laboratory experiment in which the choice of the job characteristics (i.e., the mission) is part of the compensation package that principals can use to influence agents' effort. Principals offer contracts that specify a piece rate and a charity—which can be either the preferred charity of the agent, or the one of the principal. The agents then exert a level of effort that generates a profit to the principal and a donation to the specified charity. My results show that the agents exert more effort than the level that maximizes their own pecuniary payoff in order to benefit the charity, especially their preferred one. The principals take advantage of this intrinsic motivation by offering lower piece rates and by using the choice of the charity as a substitute to motivate effort. However, I also find that because of fairness considerations, the majority of principals are reluctant to lower the piece rate below a fair threshold, making the substitution between monetary and non-monetary incentives imperfect. These findings have implications for the design of incentives in mission-oriented organizations and contribute to our understanding of job satisfaction and wage differentials across organizations and sectors.

JEL Classification Codes: C92, J33, M52, M55

Keywords: mission, intrinsic motivation, incentives, experiment

*University of Zürich, Blümlisalpstrasse 10, CH-8006, Zürich, Switzerland. Email: lea.cassar@econ.uzh.ch. I would like to thank in particular Roberto Weber for his excellent supervision. The paper has also benefited from comments and discussions with Gani Aldashev, Bjorn Bartling, Ernesto Dal Bó, Donja Darai, Florian Engl, Florian Englmaier, Ernst Fehr, Holger Herz, David Huffman, Michael Kosfeld, Michel Marchal, Nick Netzer, Marco Ottaviani, Eldar Shafir, Klaus Schmidt, Florian Zimmermann and with participants at Micro Economic Workshop at the University of Munich, at the N.G.O workshop 2014 at Warwick University, at the Workshop in Behavioral Economics and Experimental Research at UNIL, and at the Zurich Workshop of Economics 2013. This project was made possible through the support of a grant from the John Templeton Foundation. The opinions expressed in this paper are those of the author and do not necessarily reflect the views of the John Templeton Foundation. A previous version of this paper circulated under the title “Contract design with motivated agents: Experimental evidence.”

The design of incentives in organizations has always been at the core of economists' research agendas, motivating a large body of the empirical and theoretical literature. To date, however, most of the economic studies have focused on the design of financial incentives, such as pay-for-performance contracts and bonus awards. While these incentives can be a powerful tool in firms where employees' primary goal is to earn money, they may become less effective—or even harmful (Benabou and Tirole, 2003)—in organizations where workers perform tasks that have positive externalities for society, e.g. in organizations involved in the provision of social goods and services. Workers in these organizations are often driven only partly by financial rewards, but also by the mission of their job, i.e., by the overall job design and characteristics and its impact on society. This raises the question of whether principals use the job mission as a substitute for monetary incentives in eliciting effort from motivated agents. This is the central question of this paper.

The answer to this question has implications for a wide set of labor market environments, for instance, for the design of contracts in organizations that are involved in the provision of social goods and services. Given that many of these organizations do not generate profit and, therefore, have limited financial resources at their disposal, non-monetary incentives can play a major role in motivating employees, while, at the same time, reducing the organization's financial expenses. If employees care about the mission of their projects, it may be desirable for an employer to increase the employees' discretion in designing their own projects even though their mission preferences may not be perfectly aligned. For example, employees of a non-profit organization in the field of education may be willing to exert the same level of effort for a lower wage in order to be able to choose the characteristics of the project on which they work: the educational curricula that should be taught, the teaching methods that should be used, the beneficiaries' targeting, and so on. Such “ideological” issues play an important role in the social and development sector (Besley and Ghatak, 1999, 2001).

More generally, this analysis has implication for the design of contracts in organizations where employees solve tasks that involve a certain level of creativity. Motivated journalists, researchers, architects, and computer engineers, can feel very strongly about the products they develop and may be willing to give up financial gains in order to decide which news to cover, which research topics to address, which phone applications to design. Do firms and organizations such as the New York Times, the World Bank and Apple, take advantage of this intrinsic motivation when designing their contracts? Google seems to understand how to make the most of its employees' intrinsic motivation and mission preferences: “... by allowing its engineers to spend 20% of their work week on projects that interest them, Google is able to

tap into the many talents of its employees.”¹

I address the above question in a novel laboratory experiment in which the choice of the job mission is part of the compensation package that a principal can use to motivate an agent. A laboratory experiment is an appropriate method to address this issue for several reasons. First, it allows to perfectly control for self-selection of agents with unobservable characteristics. For example, while a negative wage differential between the public and the private sector may be consistent with public organizations substituting mission and job design for monetary compensations, it may also result from workers with higher abilities self-selecting in the private sector, or from different production functions across sectors. Second, as it will become clearer from the design description, since the mission preferences of the principals and of the agents are not perfectly observable, identification must be achieved by comparing the behavior of the same individuals in different contexts, namely, by a within-subject design. Third, the effort task must be monetary in order to compare the agents’ effort choices with the exact predictions made by different theoretical models. Only laboratory data allow one to perfectly control for all these issues.

In my experiment, pairs of subjects form principal-agent contractual relationships for the development of a project. The principal offers the agent a contract, which consists of a mission and a piece rate, to carry out the project. The agent then chooses a monetary effort level that generates a profit to the principal and a social output. The social output is implemented as a donation to a charity, while the choice of the project mission is implemented as the choice of which charity receives the donation. More specifically, the principal can offer two types of contracts: an open contract, or a closed contract. If the *closed* contract is offered, the effort exerted by the agent generates a donation to the principal’s preferred charity. If the *open* contract is offered, the effort exerted by the agent generates a donation to the agent’s preferred charity.² The agent is paid according to the piece rate specified in the offered contract. After receiving the contract, the agent chooses his effort level, which determines payoffs and donations. Thus, the employer cannot contract directly on effort. He can, however, influence the agent’s effort level through the choice of the piece rate and through the choice of which of the two types of contract to offer, namely, the mission.

Notice that the choice of the project mission is only relevant if it entails a trade-off for the principal, namely, if the principal and the agent prefer different charities. If this wasn’t the

¹<http://www.forbes.com/sites/laurahe/2013/03/29/googles-secrets-of-innovation-empowering-its-employees>

²It is worth pointing out that differently from Fehr et al. (2013) and Charness et al. (2012), neither contracts imply a proper delegation of decision from the principal to the agent. The decision is always made by the principal about whose preferred charity receives the donation. Thus, this paper is not about the intrinsic value of authority.

case, the choice between the open and the closed contract would be trivial and the question of whether the choice of the project mission is treated as substitute to monetary incentives could not be addressed. Thus, to generate a pool of subjects who care about different charities, in the recruitment process students were informed that they would earn slightly less than usual in experiments but that they would be given the opportunity to generate a substantial donation to their favorite charity, which had to be specified in advance.³

In order to identify how agents' intrinsic motivation and mission preferences affect the contracts offered by the principals, I need to compare contracts offered to motivated agents with the contracts that the principals would offer if the agents were not intrinsically motivated and, therefore, did not care about the project mission. For this purpose, the experiment consisted of two treatments. In the *main* treatment, the agents were allowed to choose any desired effort level for any given piece rate. The parameters of the payoff functions were chosen such that an agent who is intrinsically motivated and thus cares about the charity, would choose an effort level that is higher than the piece rate offered in the contract. In the *control* treatment, the agents were not allowed to choose more effort than the level that maximizes their monetary payoff. That is, they could not choose an effort level that is higher than the piece rate specified in the contract. In this treatment any effort premium from agents' intrinsic motivation was, therefore, ruled out. The agents had to behave *as if* they did not care about the charities.

By comparing the piece rates and the frequency of open contracts offered across the two treatments, I am able to test whether principals treat monetary incentives and the choice of the project mission as substitutes to elicit effort from motivated agents. If this is the case, one should observe agents in the main treatment being offered the open contract more often and being paid less than agents in the control treatment. That is, in the presence of an effort premium from agents' intrinsic motivation, principals should be more likely to let the agents work for their preferred charity and they should offer lower monetary incentives compared to situations where this effort premium is absent.

The main results are the following: (i) Agents were motivated to generate a donation under both contracts, but more so when the donation was made to their preferred charity. In other words, conditional on the piece rate, in the main treatment effort exerted under the open contract was higher than under the closed contract, and both these efforts were higher than the effort exerted under any contract in the control treatment. This result shows that the mission of a job matters for motivated workers: the higher the mission-matching, the higher the effort

³This recruitment procedure has, of course, generated a selected pool of subjects. This selection is, however, necessary because the underlying theory is about employers and workers who care about the mission of the project. The design is meant to reflect the environment of mission-oriented organizations rather than purely profit-driven firms.

level. (ii) Given a type of contract, open or closed, the piece rates were lower in the main treatment than in the control treatment. Hence, principals took advantage of the agents' effort premium in the main treatment by offering lower monetary incentives. (iii) The open contract was offered more frequently in the main treatment than in the control treatment. Thus, if faced with motivated agents who can act on their motivation, principals were more willing to compromise on the job mission to save on monetary incentives.

These results are further supported when matching subjects' decisions in the experiment with their choices from an allocation game, which was run at the end of the experiment to measure the charity preferences at an individual level. More specifically, I find that the results apply mainly to those subjects who, according to their decisions in the allocation game, have a stronger preference for their chosen charity compared to other charities.

Viewed jointly, these findings provide evidence that principals treat the choice of the job mission as a substitute for monetary incentives, in order to elicit effort from motivated agents. In the experiment, this resulted in the same *realized* level of effort across treatments but with different combinations of incentives. That is, in the main treatment, principals were able to replace monetary incentives with non-monetary incentives and still induce the same level of effort as in the control treatment.

I also find, however, that principals were heterogenous in their willingness to substitute monetary with non-monetary incentives. This heterogeneity was driven by fairness considerations. I find that the high number of principals who, in a postexperiment questionnaire, reported some fairness concerns about offering a lower wage to a motivated worker, sacrificed profits because they were reluctant to lower the piece rate below a fair threshold. Thus, these fairness concerns acted as a limited liability constraint, making the substitution between the mission and the piece rate imperfect.

The findings of this paper contribute not only to the design of incentives in organizations, but also to our understanding of job satisfaction and wage differentials across organizations and sectors. They may help explain, for instance, why firms such as Google may be able to pay lower wages than their main competitor without becoming any less attractive to talented employees.⁴ They may also explain why workers in mission-oriented sectors, such as welfare and religious services, are paid the lowest wage, but report the highest job satisfaction, across all sectors Pischke (2011). The results of this paper suggest that this is likely not driven only

⁴Circumstantial evidence reported in the New York Times suggests that: "...Google pays less and its health insurance is not as good as Microsoft." Source: <http://bits.blogs.nytimes.com/2007/06/28/google-v-microsoft-whats-the-better-workplace>. Taking a more extreme example, these results provide also an explanation for the recent empirical evidence in Bahney et al. (2013), which shows that insurgent fighters by Al Qa'ida Iraq are paid less than their outside options, even when compared to unskilled labor, and receive an estimated negative risk premium.

by self-selection of happy workers into these industries, but also by specific non-monetary job attributes in these sectors which substitute for low income.

The paper relates to several strands of literature. First, it relates to contract theory models with motivated agents (Besley and Ghatak, 2005; Delfgaauw and Dur, 2007; Cassar, 2014). In a setting where the job mission is exogenous, Besley and Ghatak (2005) show that it is optimal for a principal to offer lower monetary incentives if he is matched with an agent who shares his same mission preferences. Similarly, Delfgaauw and Dur (2007) shows that more motivated workers imply weaker monetary incentives. In a setting where the choice of a project mission is endogenous and where the principal and the agent have misaligned mission preferences, Cassar (2014) shows that it is optimal for the principal to (partly) compromise on the choice of the mission in order to save on financial payments. Thus, overall, the findings of this paper provide experimental support for the predictions of these models. Second, the paper relates to previous experimental studies that show that the matching of the job mission with the workers' preferred mission increases effort (Carpenter and Gong, 2013; Gerhards, 2013; Fehrler and Kosfeld, 2013; Besley and Ghatak, 2013; Tonin and Vlassopoulos, 2014). These studies, however, only focus on the behavior of the agents, and do not address the main question of this paper.⁵ Finally, the paper points to a different "hidden cost of control" which, contrary to Falk and Kosfeld (2006) and Bartling et al. (2012, 2013), does not arise from the perception that the lack of discretion is a signal of the principal's distrust, but from the fact that motivated workers have direct preferences on how to solve their tasks, and these preferences are not always aligned with the ones of their employers.

The remainder of the paper is organized as follows. The next section introduces the theoretical framework and predictions. Section II describes the experimental design. Section III presents the results. Section IV discuss some results in more details. Section V concludes.

⁵Gerhards (2013) also analyzes the principals' choice and finds that principals pay higher piece rates in contracts with the agents' preferred mission than in contracts with a random mission. However, given that in the design by Gerhards (2013) principals often share the same mission preference as the agents, this may simply be the result of the employers wanting to elicit more effort for their preferred mission than for a random mission. Finally, (Nyborg and Zhang, 2013; Nyborg, 2014) shows that firms with a reputation as socially responsible pay significantly lower wages. These studies, however, only report correlations and cannot make any causal argument.

I. THEORETICAL FRAMEWORK

I.1. The Model

The theoretical basis for my experimental design is a simple extension, which I describe below, of the model by Besley and Ghatak (2005) to an environment where the project's mission is endogenous and where the principal and the agent have different mission preferences.

An agent is matched with a principal who offers him a contract for the development of a social project. The contract specifies a wage per unit of effort w , henceforth piece rate, and a project's mission $m \in \{m_P, m_A\}$, where m_P is the preferred mission of the principal, m_A is the preferred mission of the agent, and $m_P \neq m_A$. The agent will then exert a level of effort, e , to develop the project. The utility of the principal from implementing the contract (w, m) is equal to

$$V(e; w, m) = \Pi(e, w) + \gamma_m D(e) \quad m = m_A, m_P \quad (1)$$

$\Pi(e)$ is the profit generated by the project as a function of the agent's effort. This is composed of the monetary output minus the piece rate cost. γ_m is the principal's valuation for the social output of a project with mission m , and $D(e)$ is the social output generated by the project as a function of the agent's effort. It follows that $\gamma_{m_P} > \gamma_{m_A} \geq 0$, that is, the principal values more the social output of a project with his preferred mission than with the preferred mission of the agent. The principal chooses (w, m) that maximizes (1). Given the contract (w, m) , the agent chooses the level of effort that maximizes the following utility:

$$U(w, m; e) = we + \alpha_m D(e) - C(e) \quad m = m_A, m_P \quad (2)$$

where α_m is the agent's intrinsic motivation to develop a project with mission m and $C(e)$ is the cost of exerting effort, which is equal to $\frac{1}{2}e^2$. It follows that $\alpha_{m_A} > \alpha_{m_P} \geq 0$, that is, the agent's intrinsic motivation is higher when he works for his preferred mission than for the preferred mission of the principal. If the agent were not motivated, he would not care about the project mission, i.e. $\alpha_{m_A} = \alpha_{m_P} = 0$.⁶

I.2. Analysis and Theoretical Implications

For sake of notational clarity, let's define w_m as the piece rate paid to the agent if the project's mission m is chosen. Finally, to keep the analysis simple, $\Pi(\cdot)$ and $D(\cdot)$ can be assumed to be

⁶Notice that in principle γ_{m_A} and α_{m_P} could also be negative, that is, the principal and the agent could derive an intrinsic cost from the social output of a project with mission m_A and m_P respectively. This could be the case, for instance, if mission preferences reflect different political or religious views, as in Carpenter and Gong (2013). This would, however, have no fundamental implication for the theoretical predictions. Nevertheless, since the experiment involves charities rather than political parties, for an illustrative point of view it is easier to just assume these parameters to be positive.

linear in effort:

$$\Pi(e, w) = (\pi - w)e \quad (3)$$

$$D(e) = de \quad (4)$$

The agent's optimal level of effort is then

$$e_m^* = w_m + \alpha_m d \quad m = m_A, m_P \quad (5)$$

The agent's optimal level of effort is increasing in the offered piece rate, in the intrinsic motivation from working in a project with mission m , α_m , and in the effort's productivity in generating social output, d . Since $\alpha_{m_A} > \alpha_{m_P}$, motivated agents exert more effort, for any given piece rate, in contracts with their preferred mission. Furthermore, this level of effort is higher than the piece rate. If an agent is not motivated, his optimal level of effort is equal to the piece rate independently of the project mission. This leads to the following prediction:

Prediction 1 *The effort premium from agents' intrinsic motivation is higher when the agents work for their preferred mission than for the preferred mission of the principals.*

Given the optimal level of effort in (5), the optimal piece rate in a contract with mission m is equal to

$$w_m^* = \text{Max} \left\{ 0 ; \frac{1}{2} [\pi + d(\gamma_m - \alpha_m)] \right\} \quad m = m_A, m_P \quad (6)$$

The optimal piece rate is increasing in the effort's productivity in generating profit, π and in the principal's valuation of the social output of a project with mission m , γ_m . This is the *preference channel*: everything else being equal, the more the principal cares about the social output of a project, the higher the effort level he wants to elicit from the agent, and in turn, the higher the piece rate he will offer in the contract. More importantly for this study, the optimal piece rate is decreasing in the agent's intrinsic motivation to work for a project with mission m , α_m . This is the *substitution channel*: everything else being equal, the higher the agent's intrinsic motivation, the more the principal can save on monetary incentives, i.e., the more he can lower the piece rate. The experiment is designed to identify this channel.

These two channels have three implications for contract design. First, since the preference channel is stronger under the contract with the principal's preferred mission ($\gamma_{m_P} > \gamma_{m_A}$) while the opposite holds true for the substitution channel ($\alpha_{m_A} > \alpha_{m_P}$), wages are higher in contracts with the principal's preferred mission, i.e. $w_{m_P}^* > w_{m_A}^*$. This leads to a negative correlation between financial incentives and the alignment of the project mission with the agent's mission preferences. Second, because of the substitution channel, motivated agents are offered lower wages compared to non-motivated agents. If $\alpha_{m_P} > 0$, this is true for both project missions.

Third, this piece rate differential should be higher if the motivated agent works for his preferred mission, as the substitution channel is stronger in this case. In other words, the principal can take more advantage of the agent's intrinsic motivation by lowering the piece rate when the latter works for his preferred mission. This generates the following prediction(s):

Prediction 2 *i) Wages are higher in contracts with the principals' preferred mission; ii) Motivated agents are paid less than non-motivated agents; iii) The piece rate differential in ii is higher if motivated agents work for their preferred mission than for the preferred mission of the principals.*

By replacing the optimal piece rate in the optimal effort function of the agent, I can derive the optimal level of effort as a function of the exogenous parameters only:

$$e_m^* = \frac{1}{2}[\pi + d(\gamma_m + \alpha_m)] \quad m = m_A, m_P \quad (7)$$

Equation (7) suggests that for any project mission, the unconditional level of effort is increasing in both the principal's valuation for social output and the agent's intrinsic motivation. Comparative statics depend on the specific values taken by γ_m and α_m . Since these parameters can take any value, I cannot predict whether the unconditional level of effort will be higher under the contract with the principal's preferred mission or under the contract with the agent's preferred mission. For sufficiently high α_{m_A} , however, it is clear that the level of effort will be higher under the contract with the agent's preferred mission. Since wages are lower under such contracts, this can result in a negative correlation between effort and monetary incentive in a cross-section of organizations.

Finally, I can derive the principal's utility from each of the two contracts to make predictions about the optimal choice of the project's mission. Equation (1) can be rewritten as

$$V(e^*; w^*, m) = (\pi - w_m^*)e_m^* + \gamma_m de_m^* \quad m = m_A, m_P$$

Replacing the optimal level of effort and the optimal piece rate gives

$$V(m) = \frac{1}{4}[\pi + d(\gamma_m + \alpha_m)]^2 \quad m = m_A, m_P \quad (8)$$

As for the unconditional level of effort, the principal's utility from the contract with project's mission m is increasing in both the principal's valuation for social output and the agent's intrinsic motivation for mission m . Which project's mission maximizes the principal's utility depends on the specific value taken by γ_m and α_m , $\forall m = m_A, m_P$. Notice, however, that the higher is the difference between α_{m_A} and α_{m_P} , that is, the stronger the agent feels about his preferred mission, the more likely is $V(m_A)$ to be higher than $V(m_D)$ and, therefore, the

more likely should be the principal to let the agent work for his preferred mission. Now, in the extreme case where $\alpha_{m_A} = \alpha_{m_P}$, i.e. the agent is indifferent about the missions or the agent is not-motivated, the principal is undoubtedly better off by choosing his preferred mission, i.e. $m^* = m_P$. Indeed, if the principal cannot take advantage of the substitution channel, there is no reason why the principal should make a compromise on the mission. Therefore, compared to a situation where the agent is not intrinsically motivated, if the agent is motivated the principal should be more likely to offer the contract with the agent's preferred mission.⁷ The last prediction follows:

Prediction 3 *Principals should be more likely to compromise on the mission the more the agents are motivated.*

Thus, all together, these predictions 2-3 imply that motivated agents should be compensated with higher non-monetary incentives and lower monetary incentives than non-motivated agents.

II. THE EXPERIMENT

The design of the experiment closely follows the theoretical framework described in Section I. The social output generated by the project is implemented as a donation to a charity that subjects can generate in addition to their monetary payoffs. The choice of the project's mission is implemented as the choice of which charity receives the donation. My identification strategy of the causal effect of the agents' intrinsic motivation on the choice of the piece rate and of the project mission is the comparison between the contracts offered to motivated agents and the contracts offered if the agents were constrained to behave as if they were not motivated. To do so, I run the two treatments described in Section II.3.

II.1. Recruitment

The implementation described above only works if subjects care about different charities. If this were not the case, the choice of the project mission, i.e. the choice of which charity receives the donation, would not entail any trade-off for the principal. Thus, the question whether the choice of the project mission is treated as substitute to monetary incentives could not be addressed. To generate such pool of subjects, in the recruitment process students were informed that they may earn slightly less than usual in experiments but that they would be given the opportunity

⁷In a setting where m is continuous, the model would predict that the higher the agent's intrinsic motivation, the larger the compromise that the principal would make on the project mission (Cassar, 2014).

to generate a substantial donation to their favorite charity. Furthermore, in order to sign up, subjects had to specify the name and the website of their favorite charity. This recruitment procedure turned out to be successful as 92 different charities were chosen by 146 subjects. As soon as the subjects arrived in the laboratory, they received a list of all the charities that were chosen by subjects participating in the same experimental session. The list also specified the number of subjects choosing the same charity. Therefore, subjects knew the full distribution of charities choices.

In order to have updated information about subjects' charity preferences, at the beginning of the experiment subjects were asked to choose (again) their preferred charity from the list. Subjects' specific charity preferences were, however, kept confidential throughout the experiment. In other words, subjects knew that all the preferred charities of other subjects must be in the list but, when interacting with another subject, they did not know which specific charity in the list that other subject had chosen. As it will become clearer in the next section, this design choice had several advantages, including avoiding the risk, *ex-ante*, of matching subjects with preferences for the same charity.

II.2. Experimental design

After the instructions were read aloud, half of the subjects were assigned to the role of principals, the other half the role of agents. The same role was kept throughout the experiment, which consisted of multiple periods. At the beginning of each period, every agent was randomly and anonymously re-matched to a principal.

In each period, the sequence of actions was as follows. The principal chose a piece rate w in the integer set $\{1, 2, \dots, 9\}$, and whether to offer a closed or an open contract. If the closed contract was offered, the effort exerted by the agent generated, in addition to the principal's profit, a donation to the principal's preferred charity. If the open contract was offered, the effort exerted by the agent generated, in addition to the principal's profit, a donation to the agent's preferred charity. Thus, in terms of the model described in Section I, the closed contract corresponds to the choice of a project mission equal to m_P , while the open contract corresponds to the choice of a project mission equal m_A .

The agent then chose a costly level of effort e in the integer set $\{1, 2, \dots, 9\}$ which determined payoffs and donations in that period. To rule out binding budget constraints, at the beginning of each period every agent was endowed with 35 points. This allowed the agent to exert the maximum level of effort even if the principal chose a piece rate of 1. Both the principal's profit and the level of donation depended linearly on effort as in (3) and (4). Specifically, $\pi = 10$ and $d = 20$. This choice of parameters, along with the convex effort function, ensured that it

was not convenient neither for the principal nor for the agent to maximize own income and donate later. Finally, at the beginning of each period a base income of 20 points was given to every principal in order to get roughly the same monetary payoffs under the optimal choices for agents and principals who did not care about charities, i.e. for $e = w = 5$ in both contracts.

As subjects may vary in the extent to which they care about the charities and this is not observable, the strategy method was necessary to elicit the principal's choice of piece rate and the agent's choice of effort for both the closed and open contract. More specifically, I elicited the agent's effort choice for each possible piece rate in each contract. To elicit the principal's beliefs about the agent's effort, the principal also specified an expected level of effort for each of the two contracts given the chosen piece rate. This elicitation was, however, not incentivized and the expected effort level was not revealed to the agent. The principal would then choose which contract, between the open and the closed, to offer to the agent. The chosen contract was implemented with 75% probability, whereas with 25% probability the other contract was implemented. This random implementation of contracts was necessary to make the choice of piece rate incentive compatible even in the principal's least preferred contract. Furthermore, the fact that agents did not know with certainty whether the implemented contract corresponded to the contract chosen by the principal, and that the agents' effort in each contract was elicited ex-ante through the strategy method, make it very unlikely that the agents reciprocated principals who chose to offer the open contract.⁸ At the end of each period, payoffs realized and became observable to the principal and the agent. To improve learning, the principal observed not only the level of effort and the donation realized, but also the level of effort chosen by the agent in the contract that was not implemented. Similarly, at the end of each period, the agent observed not only the piece rate offered in the implemented contract, but also the piece rate offered in the remaining contract. He also observed which of the two contracts the principal had chosen to offer. Principals and agents were then randomly re-matched.

As pointed out earlier, neither the principals nor the agents were revealed the preferred charity of the subjects they were matched with. It follows that the principal's choice between the closed and the open contract was equivalent to a choice between a contract that generated a donation to his preferred mission with certainty and a contract that generates a donation to any of the charity specified in the list with some probability.⁹ Similarly, an agent knew that

⁸Reciprocity towards the choice of the open contract would have, indeed, the same implications as an increase in α_{m_A} , so it could have acted as a potential confound. On the contrary, note that reciprocity towards high wages, by reducing the principal's benefit from saving on monetary incentives, works against my hypothesis. Similarly, reciprocity only towards high wages in the open contract leads to the opposite prediction regarding the optimal choice of wages, namely, the principal may want to set higher wages in the open contract than in the closed contract.

⁹This probability was known to the subjects because, as mentioned above, they knew the full distribution of charity choices.

his effort level in the closed contract would generate a donation to any charity specified in the list with a given probability, whereas his effort level in the open contract would generate a donation to his preferred charity with certainty. Thus, everything else being equal, it is clear that the principal’s valuation of social output was higher under the closed contract while the agent’s motivation was higher under the open contract. Theoretical predictions were, therefore, unchanged.

While this design feature did not affect the theoretical predictions, it had two main advantages. First, this design avoided the risk, ex-ante, of having too many pairs of principal-agent preferring the same charity. Second, it eliminated any heterogeneity across periods and treatments in the agent’s intrinsic motivation under the closed contract and in the principal’s valuation of output in the open contract. In terms of the model, this was equivalent to keeping α_{m_P} and γ_{m_A} of each individual constant across periods and treatments. Indeed, such unobservable heterogeneity across periods and treatments could have added substantial noise to the data and acted as a potential confound.

II.3. Treatments

The experiment involved two treatments. Each subject participated in both treatments. In the *main* treatment the agent was free to choose any desired level of effort. As shown in equation (5), if the agent cared about the charity he would choose an effort level that is higher than the piece rate offered in the contract. In the *control* treatment, the agent was not allowed to exert more effort than the level that maximized his material payoff. That is, the level of effort could not be higher than the piece rate specified in the contract. In this treatment any effort premium from agents’ intrinsic motivation was, therefore, ruled out.

To keep the notation consistent with the theory described in Section I, let w_{m_A} and e_{m_A} be, respectively, the piece rate and the effort level chosen in the open contract, while let w_{m_P} and e_{m_P} be, respectively, the piece rate and the effort level chosen in the closed contract. Treatment variations are summarized in Table I.

Table I: Treatment variation

Treatment	Effort
Main	$e_{m_A} \in \{1, 2, \dots, 9\} ; e_{m_P} \in \{1, 2, \dots, 9\}$
Control	$e_{m_A} \leq w_{m_A} ; e_{m_P} \leq w_{m_P}$

To ensure that both the principals and the agents understood the payoff implications of their choices, a payoff table was given in the instructions where, in red, was the effort choice

that maximizes the agent's income for any given piece rate - thus, the table diagonal was in red - and, in blue, was the effort choice that maximizes the donation for any given piece rate, i.e. an effort level of 9 for any given piece rate. Thus, based on this payoff table, it was clear to both the principals and the agents that if the latter cared about the charity, in the main treatment they would choose a higher effort level than the piece rate.

Six sessions were run in total. To control for potential order effects, the order in which the two treatments were played were systematically modified across sessions. Each treatment lasted 10 periods. So overall, every subject made decisions for 20 periods. Payments to subjects and donations to charities were made according to the outcome of one out of the 20 periods selected at random.

II.4. Elicitation of Preferences over Charities

It is important to point out that subjects are likely to have heterogenous preferences over the charities. In terms of the model, this means that the agent's parameters α_{m_A} and α_{m_P} and the principal's parameters γ_{m_P} and γ_{m_A} are individual specific. Therefore, to add control and run robustness checks, it is useful to get an independent measure of these parameters at an individual level.

For this purpose, I use an allocation game with concave allocation function. At the end of the experiment, but before the subjects knew which period had been randomly selected to count for payments and donations, subjects were asked to allocate 40 points between themselves, their preferred charity, and a charity selected at random from the list. At the time of making the decision, however, subjects did not know which charity would be randomly selected from the list. So their information was analogous to the information they had when making decisions during the experiment. Let's define p_{ij} as the number of points allocated by subject i to a particular recipient j . To induce non-linearity and avoid corner solutions, we apply the following non-linear transformation to the payoff of recipient j : $P_{ij} = 7 * \sqrt{p_{ij}}$. That is, recipient j receives a payoff of P_{ij} points if subject i allocates p_{ij} points to him/her. Since every subject must allocate all 40 points, it follows that $\sum_j p_{ij} = 40 \forall i$.

According to this allocation function, any subject who cares about his preferred charity, i.e. any agent with $\alpha_{m_A} > 0$ or any principal with $\gamma_{m_P} > 0$, should allocate some points to his/her preferred charity. The optimal number of points allocated to the preferred charity is increasing in α_{m_A} , or equivalently in γ_{m_P} . Similarly, any subject who cares about the charities in the list as a whole, i.e. any agent with $\alpha_{m_P} > 0$ or any principal with $\gamma_{m_A} > 0$, should allocate some points to the random charity. The optimal number of points allocated to the random charity is increasing in α_{m_P} for the agents, or equivalently in γ_{m_A} for the principals.

The advantage of using an allocation game with concave allocation function is that it rules out potential corner solutions that may arise from efficiency concerns. So this method generates an approximate measure at an individual level of (i) the intensity with which subjects care about their preferred charities, i.e. of the size of α_{m_A} and γ_{m_P} ; (ii) the importance of subjects' preferred mission compared to other missions, i.e. $\alpha_{m_A} - \alpha_{m_P}$ and $\gamma_{m_P} - \gamma_{m_A}$. The latter measures to what extent a principal, or an agent, is *mission-driven* and is particularly relevant for comparative statics across types of contract.

II.5. Questionnaire on fairness

Finally, at the end of the experiment subjects answered a very short questionnaire. The purpose of the questionnaire was to elicit subjects' fairness perceptions about the substitutability between monetary and non-monetary compensations. Previous survey evidence by Kahneman et al. (1986) reveals, indeed, that fairness considerations may affect wage-setting decisions.

More specifically, subjects were asked the following two questions: i) "Suppose you run an organization with two employees: employee A and employee B. These employees are assigned the exact same task. Furthermore, they are identical in every aspect: in particular, they have the same ability and are equally productive. There is only one difference: employee A enjoys the task more than employee B does. Would you pay them the same wage?" Subjects could choose between three options: "same wage", a "higher wage to employee A", and a "higher wage to employee B"; ii) "Suppose you run an organization and you are in the process of hiring a new employee. During the job interview, you realize that the candidate for the position would really enjoy the job so he/she would accept to do the same amount of work even for a lower wage than the one you initially planned to offer. Would you offer the candidate a lower wage than the one you initially planned?" After having answered this question, subjects were asked whether they thought it was fair to offer a lower wage.

II.6. Procedural details

Overall 146 subjects participated in six sessions, 73 in the role of principals and 73 in the role of agents. The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007). Our subject pool consisted primarily of students at the University of Zurich and at the Federal Institute of Technology in Zurich and were recruited using the software hroot (Bock et al., 2012). Subjects were paid based on the number of points generated in one period selected at random from the experiment and based on the number of points they allocated to themselves in the allocation game. Donations to charities were also made accordingly. In addition, each subject received a show-up fee of 10 SFr. On average, an experimental session lasted 1 hour

and 45 minutes. The average payment was 35 SFr and the average donation was 36 SFr per pair of subjects.

II.7. Hypotheses

The analysis of the theoretical model in Section I generates a set of predictions that I now restate as hypotheses to be tested in the experiment.

First, the model predicts that conditional on piece rate, the effort premium from agents' intrinsic motivation is higher in contracts with the agents' preferred mission. Prediction 1 translates to the following hypothesis:

Hypothesis 1 *In the main treatment, conditional on piece rate, effort is higher in the open contract than in the closed contract. Furthermore, effort in the main treatment should be higher than in the control treatment.*

This first hypothesis tests if the project mission is effective in motivating effort provision. Notice that this hypothesis needs to be satisfied in order to test the remaining hypotheses.

Next, conditional on the hypothesis above being satisfied, the model makes predictions about the level of piece rate offered by principals. Specifically, it predicts that wages are higher in contracts with the principals' preferred mission. Indeed, in such contracts, there is (i) a *stronger* preference channel: the principal wants to elicit a higher donation for his preferred charity than for any other charity; (ii) a *weaker* substitution channel: the agent needs higher monetary incentives to provide effort for the principals' preferred charity than for his own preferred charity. Prediction 2a translates to the following hypothesis:

Hypothesis 2 *Wages are higher in the closed contract than in the open contract.*

This second hypothesis tests if, contrary to the results in Gerhards (2013), there is a negative correlation between financial incentives and the alignment of the project mission with the agents' mission preferences. Testing this hypothesis alone, however, can not answer whether principals take advantage of the agents' intrinsic motivation by saving on monetary incentives. To address this question, I need to disentangle the substitution channel from the preference channel. My identification strategy is to compare the piece rate offered by principals within a specific type of contract across the two treatments. By comparing wages within the same type of contract, the preference channel is constant, and therefore, any observed piece rate differential across treatments must come from the substitution channel: while the substitution channel should be active in the main treatment, in the control treatment, since the effort premium from

intrinsic motivation is ruled out by design, the substitution channel is switched off. This is equivalent to Prediction 2b, which translates to the following hypothesis:

Hypothesis 3 *Conditional on the type of contract, wages are lower in the main treatment than in the control treatment.*

This hypothesis tests if, given a certain project mission, the principal takes advantage of the agent’s intrinsic motivation to work for that mission by lowering the piece rate. In other words, this tests if the principal would pay a motivated agent a lower piece rate than a non-motivated agent.

The size of the substitution channel will depend, of course, on the size of the effort premium resulting from the agent’s intrinsic motivation in Hypothesis 1. Since I expect the effort premium to be higher if the agent works for his preferred mission, the next hypothesis follows from Prediction 2.iii:

Hypothesis 4 *The piece rate differential across treatments is higher in the open contract than in the closed contract.*

Finally, the model makes predictions about the principal’s choice of the project’s mission. More specifically, if the agent is not motivated, the principal has no reason to choose a different mission from his preferred one. On the contrary, if the agent is motivated, the principal may want to compromise on the mission in order to save on monetary incentives. The next hypothesis follows from Prediction 3:

Hypothesis 5 *The principal is more likely to offer the open contract in the main treatment than in the control treatment.*

Hypotheses 3-5 together test if the principal treats the choice of the project mission as a substitute for monetary incentives in motivating effort provision.

III. RESULTS

I start by describing the results from the allocation game to test whether the recruitment process was successful in generating the right pool of subjects. These results will also be used for robustness checks in the subsequent analyses. I then report the agents’ effort choices. Indeed, in order to address the central question of this paper, it is necessary to find the predicted effort differences across contracts and treatments. Finally, I turn to the central part of this paper, namely, to the principals’ choices of piece rates and contracts.

III.1. Descriptive statistics from the allocation game

The results from the allocation game are illustrated in Figure I. Recall that the subjects had to allocate 40 points between themselves, their preferred charity, and one charity in the list selected at random. The figure shows the number of points allocated by subjects to their preferred charity and to the random charity in the list. By subtracting these numbers from the 40 points endowment, one can infer how many points subjects kept for themselves. Circles below (above) the 45-degree line indicate higher (lower) number of points allocated to one's preferred charity compared to the random charity. Circles below (above) the line $y = 20 - x$ indicate higher (lower) number of points allocated to oneself compared to both the charities. The size of the circles reflect the number of observations. The biggest circle represents 16 observations while the smallest circle represents one observation.

As can be seen, with the exception of one subject, all other subjects allocated equal or more points to their preferred charity than to the random charity in the list. Furthermore, only a minority of subjects, approximately 11 percent, acted as a pure profit maximizers by keeping all the money for themselves. Roughly a fifth of the subjects seemed to be indifferent between all charities in the list.

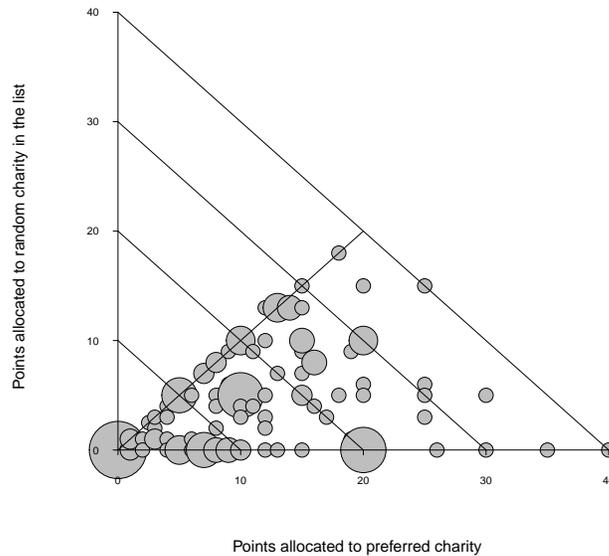


Figure I: Allocation Choices

While this is only an approximation of subjects' preferences over charities, and taking into account that these preferences were elicited only at the end of experiment, it still provides some informative evidence that, overall: i) Subjects cared about the charities; ii) The charities were

not perfect substitutes, but subjects strictly preferred their chosen charity. Thus, the recruitment process was successful in generating the right pool of subjects to address the questions of this paper.

III.2. Agents' effort choices

I start by analyzing the agents' effort choices for different wages and contracts in the two treatments. Consistent with Hypothesis 1, I find:

Result 1 *Within the main treatment, effort is higher in the open contract than in the closed contract. And for both types of contract, effort in the main treatment is higher than in the control treatment.*

Evidence for Result 1 can be seen in Figure II. This figure shows the effort choices per piece rate and type of contract elicited through the strategy method in each treatment. The black line represents the 45 degrees line. Thus, any point on the right of the black line corresponds to a (positive) effort premium from intrinsic motivation.

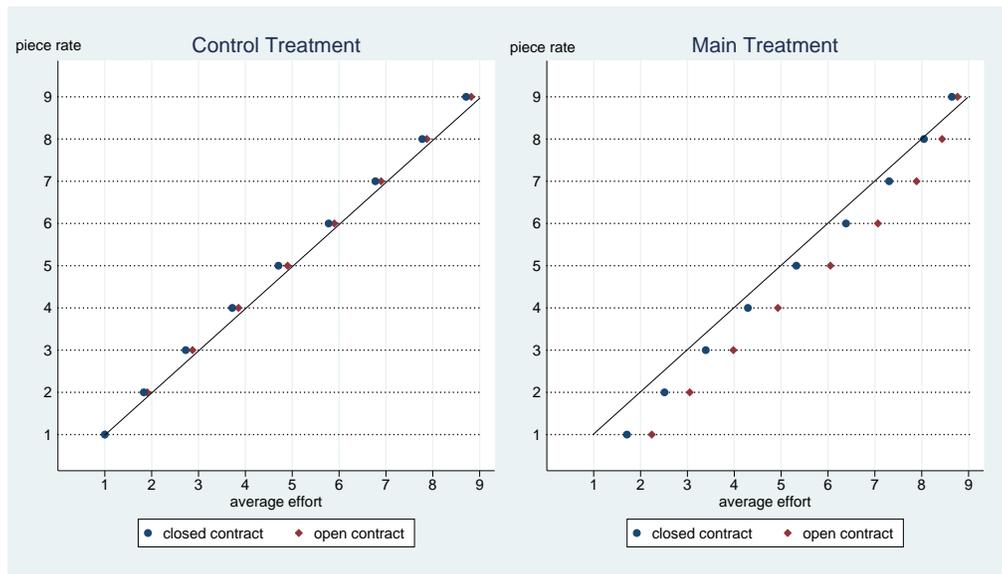


Figure II: Effort Choices

As predicted, because of the imposed restriction on the agents' choice set, in the control treatment the average effort in both contracts is equal or slightly lower than the piece rate. On the contrary, in the main treatment average effort in both contracts is always higher than the piece rate - with the obvious exception for a piece rate equal to 9. Furthermore, within the main treatment, for any given piece rate average effort is roughly half unit higher in the

open contract than in the closed contract. To assess the statistical significance of these effort differences, I use the one-sided clustered version of the signed-rank test proposed by Datta and Satten (2008), which controls for potential dependencies between observations. Within the main treatment, for each given piece rate, individual and period, I calculate the difference between the effort choice in the open contract and in the closed contract. This generates 10 observations per subject, per piece rate. Clustering at the individual level, I find that for any piece rate level, the effort choice in the open contract is significantly higher than in the closed contract ($p < 0.01$). Using the same test to assess the statistical significance of the effort differences across treatments, I find that for any piece rate level in a specific type of contract, the effort choice in the main treatment is significantly higher than the effort choice in the control treatment - and also higher than the piece rate ($p < 0.01$). This holds for both the open and the closed contract.¹⁰

Table II: Regressions

	(1)	(2)	(3)	(4)
	effort	piece rate	open contract	beliefs effort
piece rate	0.942*** (0.014)			0.686*** (0.050)
piece rate*open	-0.014 (0.008)			-0.053 (0.035)
open contract	0.179*** (0.046)	-0.755*** (0.156)		0.137 (0.224)
main treatment	0.510*** (0.109)	-0.684*** (0.131)	0.116*** (0.029)	0.309*** (0.091)
main*open	0.422*** (0.083)	0.056 (0.149)		0.703*** (0.118)
constant	0.041 (0.088)	6.041*** (0.152)	0.298*** (0.042)	1.520*** (0.299)
Round FE?	Yes	Yes	Yes	Yes
Session FE?	No	No	No	No
Individual FE?	Yes	Yes	Yes	Yes
Adj. R^2	0.878	0.125	0.024	0.468
Observations	26280	2920	1460	2920

Standard errors are clustered at the individual level. All regressions use a fixed-effect model. The dependent variables in regression (1), (2), (3), (4) are, respectively, the agent’s effort level, the principal’s belief about the agent’s effort level, the piece rate offered by the principal, the type of contract chosen by the principal. More specifically, in regression (4), the dependent variable equals “1” if principals chooses the open contract and “0” if the principal chooses the closed. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

¹⁰Figures A.1-A.9 in the Appendix also report, for any given piece rate, the time path of the effort level chosen in the open and in the closed contracts across treatments. It can be seen that Result 1 is robust throughout the time periods and wages - with the exception, of course, for the piece rate equal to 9 where there can’t be any effort premium from intrinsic motivation.

Similar results are found in the OLS regression, which is reported in Table II. It can be seen from regression (1) that conditional on piece rate, average effort is 0.51 unit points significantly higher in the main treatment than in the control treatment. Furthermore, as shown in the interaction of main treatment and open contract, this treatment difference is 0.42 unit points significantly higher in the open contract than in the closed contract. Thus, the overall effort premium from being matched with a motivated agent who works for his preferred mission is almost one additional unit of effort.¹¹ Results are robust to using a Tobit model, to clustering at the session level, or to restricting attention to the first period only, i.e., where potential spillover effects are fully ruled out.

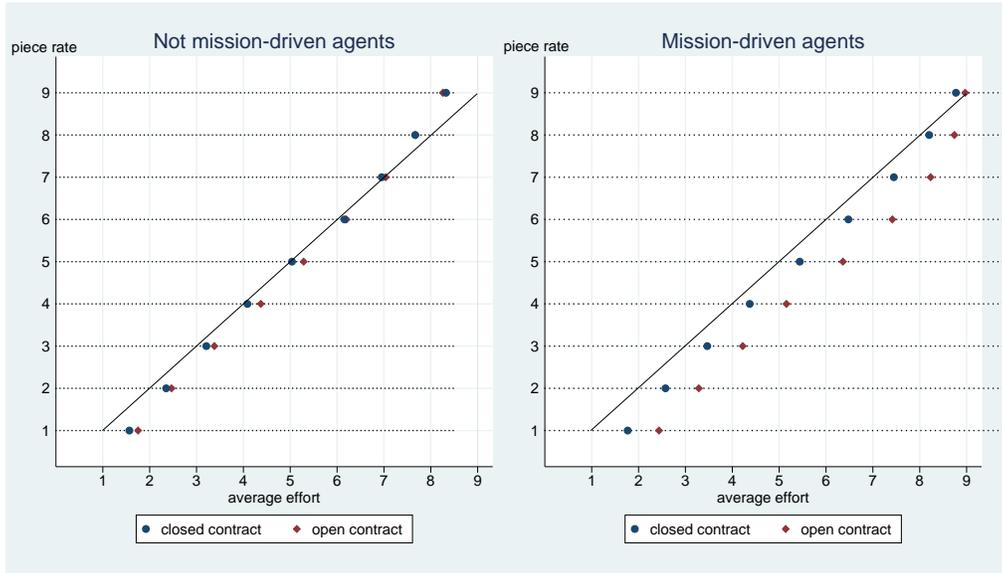


Figure III: Average effort in the main treatment by different agents

Finally, Figure III uses data from the allocation game to distinguish between the effort levels chosen by agents with different charity preferences. As Hypothesis 1 relies on the assumption that agents care more about their chosen charity than about any other charity in the list, i.e. $\alpha_{m_A} > \alpha_{m_P}$, I distinguish between those agents whom I define as *mission driven*, i.e. those who, in the allocation game, allocated more points to their preferred charity than to the random charity, and those who are *not mission driven*, i.e. who allocated the same number of points, including 0, to their preferred charity and to the random charity. As can be seen, in the main treatment, for any given piece rate, the difference in effort levels between the open and the closed contract is approximately half a unit larger for agents who are mission-driven than for

¹¹Note that if subjects in the control treatment were really non-motivated, the coefficient on the variable open contract should not be significant, as they would not care about any charity. So this is probably due to a minority of workers who disliked some charities in the list, and therefore, in the control treatment and closed contract they chose a level of effort lower than the piece rate. This has no implication for the interpretation of the results, because if anything, it should make it harder to find a treatment effect.

those who are not. For the latter, this difference almost disappears. Not surprisingly, the mission of a job is not an effective incentive tool if agents are not motivated or if they don't have specific mission preferences.

These results show that, overall, agents were motivated to generate a donation to a charity, but more so if the donation was directed to their preferred charity rather than to the preferred charity of the principal.

III.3. Principals' piece rate choices

I now analyze the piece rate offered by principals in different contracts and treatments. Consistent with Hypothesis 2, I find that:

Result 2 *Wages are higher in the closed contract than in the open contract.*

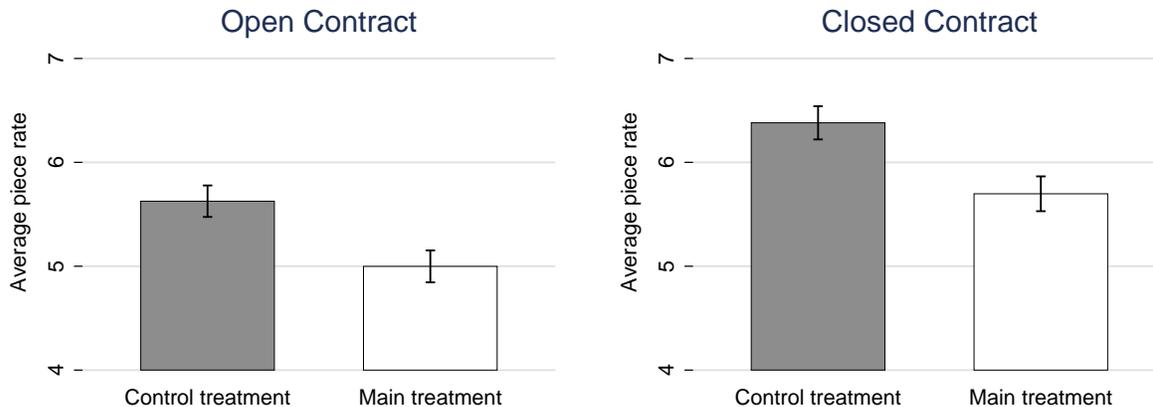


Figure IV: Piece rate choices

Evidence for Result 2 can be seen in Figure IV. This figure compares the average piece rate offered across treatments and contracts. It can be seen that in both treatments, the average piece rate in the closed contract is higher than in the open contract (signed-rank test with clusters $p < 0.01$). Notice, furthermore, that this piece rate differential should be higher for principals who are mission-driven ($\gamma_{m_P} > \gamma_{m_A}$) than for those who are not ($\gamma_{m_P} = \gamma_{m_A}$). Indeed, since the latter do not care about the charities, or are indifferent between them, the preference channel will not be higher in the closed contract. Consistent with this argument, Figure V shows that the piece rate differential between contracts is approximately 0.6 units higher for mission-driven principals than for principals who are not mission-driven (rank-sum test with clusters $p < 0.05$).

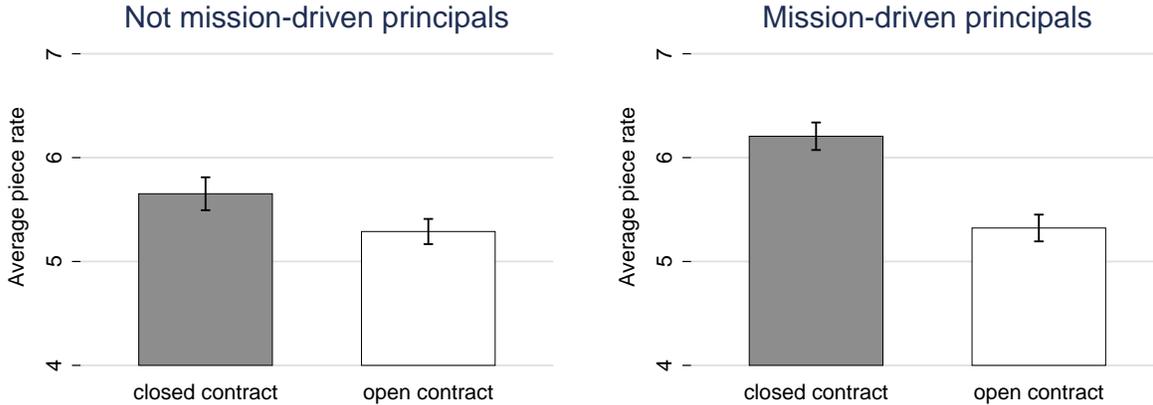


Figure V: Piece rate choices by different principals

Figure IV also reveals the first main result of this paper. Consistent with Hypothesis 3, I find:

Result 3 *For both types of contract, wages are lower in the main treatment than in the control treatment.*

It can be seen that in both the open and the closed contract, the average piece rate in the control treatment is approximately 0.7 unit points (or 14 percent) higher than in the main treatment. Both these differences are highly significant (signed-rank with clusters $p < 0.01$). Since the preference channel is constant within the same type of contract, these piece rate differentials between treatments identify the substitution channel. Thus, this result shows that principals take advantage of agents' intrinsic motivation by lowering the piece rate by approximately 14 percent.¹²

Another important observation to take from Figure IV, however, is that the piece rate differential between treatments is the same in both types of contract. This is confirmed by the signed-rank test with clusters. So in others words, in contrast with the theoretical predic-

¹²One possible alternative interpretation of the result that wages in the open contract are higher in the main treatment than in control treatment could be that principals care about the agents and do not want to prevent them from generating a high donation to their favorite charity by setting a low piece rate in the control treatment. Indeed, recall that the piece rate offered by the principal acts as an upper-bound on the effort that the agent can exert and, thus, on the donation that can be generated. This interpretation, however, cannot explain the piece rate difference across treatments that is found in the closed contract, where the donation is directed to the principal's preferred charity. Furthermore, if one believes that in the allocation game, the number of points allocated by a principal to the preferred charity of a subject taken at random is correlated with how much a principal may care about other subjects - and thus about the agents-, I find that Result 3 does not change when I distinguish between principals who allocated some points to the random charity and those who did not. The results are available upon request.

tions, the substitution channel is not higher in the open contract than in the closed contract. Hypothesis 4 is, therefore, rejected.

Result 4 *There is no difference in the piece rate differential between treatments across the two contracts.*

Similar results are found in regression (3) in Table II. It can be seen that wages in the open contract are 0.76 unit point significantly lower than in the closed contract. Furthermore, wages in the main treatment are 0.68 unit point lower than in the control treatment. Both these differences are highly significant. The interaction between open contract and main treatment is, however, positive and insignificant. This confirms that the treatment effect is the same in both contracts. Explanations for this result are discussed in Section IV. Again, results are robust to using a Tobit model, to clustering at the session level, or to restricting attention to the first period only, where potential spillover effects are fully ruled out.¹³

III.4. Principals' contract choices

Next, I analyze the principals' contract choices across treatments. Consistent with Hypothesis 5, I find:

Result 5 *The open contract is more frequently offered in the main treatment than in the control treatment.*

Evidence for Result 5 can be seen in the left of Figure VI. In the main treatment the open contract is offered approximately 40 percent of the time, whereas in the control treatment the open contract is offered approximately 28 percent of the times.¹⁴ This difference is highly significant (signed-rank test with clusters $p < 0.01$).

Additional evidence for Result 5 can be found in the regression (4) reported in Table II. It can be seen that the probability of the open contract being offered is approximately 12 percent higher in the main treatment than in the control treatment ($p < 0.01$). The marginal effect from a Probit model gives the same result.¹⁵

¹³Figure A.10 in the Appendix also reports the time path of the piece rate offered in the open and in the closed contracts across treatments. It can be seen that throughout all time periods, for both types of contracts, wages are higher in the main treatment than in the control treatment. Furthermore, within each treatment, wages are higher in the closed contract than in the open contract.

¹⁴Notice that based on the predictions from standard contract theory models, the principal should never offer the open contract in the control treatment. Indeed, as any effort premium is ruled out by design, there is no reason to offer the open contract in the control treatment. Thus, this result is probably driven by some of social preferences of the principals towards the agents.

¹⁵Figure A.11 in the Appendix also reports the time path of the contract choice across treatments. It can be seen that, with the exception of the first period, the frequency of open contracts being offered is always higher in the main treatment than in the control treatment.

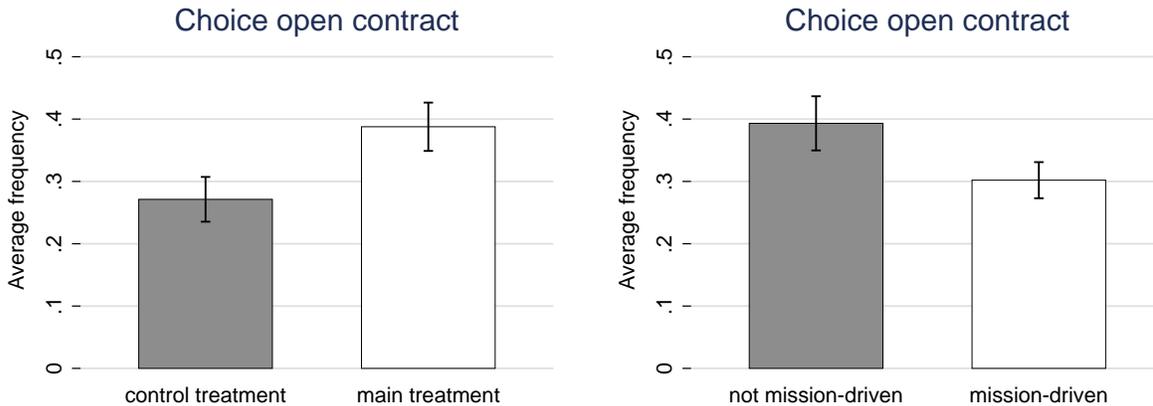


Figure VI: Contract choices by treatment and different principals

Finally, notice that principals who are not mission-driven should be more likely to offer the open contract than principals who are mission-driven. Indeed, since the formers are indifferent between the charities, they don't make any mission compromise by letting the agents generate a donation to their preferred charity. The right side of Figure VI shows exactly this. Open contracts are offered 10 percent more frequently by principals who are not mission-driven than by principals who are mission-driven. However, this difference is not significant (rank-sum test with clusters $p = 0.110$).

Results 3 and 5 together provide evidence that principals treat the choice of the project mission as substitute for monetary incentives. Motivated agents are more likely to work for their preferred missions and are paid less. This answers the central question of the paper.

III.5. Realized effort, piece rate and profit

I now turn my attention to the realized effort and piece rate choices. As it can be seen in Figure A.12, the average realized effort is approximately equal to 6 in both treatments. This means that, on average, the same amount of donation was generated across treatments. However, as it also appears in Figure A.12, the average realized piece rate is approximately 10 percent lower in the main treatment than in the control treatment. Thus, in the main treatment, principals were able to replace monetary incentives with non-monetary incentives, i.e., with letting the agents work for their preferred charities, and still induce the same level of effort as in the control treatment. As it can be seen in Figure A.13, these financial savings from agents' intrinsic motivation increases principals' average profit in the main treatment by approximately 10 percent compared to the control treatment. On the contrary, agents' average income is

approximately 9 percent lower in the main treatment than in the control treatment.

IV. Discussion

Result 4 contradicts the theoretical predictions from section I. Surprisingly, the piece rate differential between treatments is not higher in the open contract than in the closed contract. However, given that the agents' intrinsic motivation is higher under the open contract, the amount that a principal can save in terms of piece rate from being matched with a motivated agent should be higher under the open contract than under the closed contract. Why is this not the case?

Based on Figure A.14 in the appendix and regression (4) in Table II, this does not seem to depend on principals' erroneous beliefs about agents' intrinsic motivation, as the formers correctly anticipate a higher effort premium in the open contract than in the closed contract for any given piece rate. It can be seen from regression (4) that conditional on the piece rate, average expected effort is 0.3 unit points significantly higher in the main treatment than in the control treatment. Furthermore, as suggested by the coefficient of the interaction term, this treatment difference is 0.7 unit points significantly higher in the open contract than in the closed contract. So, consistent with the actual agents' effort choices, the overall expected effort premium from being matched with a motivated agent who work for his preferred mission is approximately one additional unit of effort. Hence, principals correctly anticipate that agents' are intrinsically motivated to generate a donation to a charity, particularly so in the open contract.

A potential explanation, consistent with the data, is that many principals may be sacrificing profits in the open contract because they are reluctant to offer a piece rate below what I call the *standard* piece rate, that is, the piece rate that a principal would offer if neither him nor the agent cared about the charities, namely, if $\gamma_m = \alpha_m = 0$, $i \in \{A, D\}$. Given the parameters chosen in the experiment, this corresponds to a piece rate of 5. Notice, that based on equation (6), in the main treatment the piece rate offered by the principals in the open contract should be, on average, lower than the standard piece rate. Indeed, α_{m_A} should be, on average, higher than γ_{m_A} , and therefore, in the main treatment the average piece rate in the open contract should be lower than 5.¹⁶ Importantly, as it is evident from equation (6), the same predictions do not hold for the closed contract. The reason is that in such contracts principals are more

¹⁶This is confirmed by the data from the allocation game. On average, agents allocate to their preferred charity more than twice (10 points) as much as the principals allocate to a random charity in the list (4.3 points).

motivated than the agents to generate a donation, as the latter is made to their preferred charity. That is, γ_{m_P} should be, on average, higher than α_{m_P} .¹⁷ Similarly, since in the control treatment the substitution channel is switched off, the piece rate should never be lower than 5 in neither type of contracts. Thus, if, as I hypothesize, such *fairness constraint* exists, it would only be binding for the open contract in the main treatment, but not in the closed contract nor in the control treatment.

Consistent with this argument, as shown in Figure IV, in the control treatment or in the closed contract, the average piece rate is significantly higher than 5, whereas in the open contract in the main treatment the average piece rate is just equal to 5. A similar picture emerges if we look at the distribution of wages chosen in different treatments and contracts (Figure A.15). While the frequency of wages below 5 in the open contract in the main treatment (30 percent) is significantly higher than in the closed contract (17 percent) or than in the control treatment (11 percent), it is still quite low compared to the theoretical predictions. I, therefore, conjecture that principals are reluctant to set a piece rate below the standard piece rate.

Taking this conjecture as given, the next question is why are the principals reluctant to set a piece rate below the standard piece rate. Inequity aversion as defined in Fehr and Schmidt (1999) does not seem to be the reason because, given the parameters chosen in the experiment, inequity averse principals would prefer wages that are equal or lower than the standard piece rate, but not higher. Neither is there evidence, from Figure II or Figure A.14, of negative reciprocity nor of expected negative reciprocity by principals with respect to wages below the standard piece rate in the open contract.¹⁸

Such reluctance seems related to a different fairness standard: offering a piece rate below the standard piece rate is a clear sign of exploiting workers' intrinsic motivation, and therefore, many principals who consider this to be unfair are more reluctant to lower the piece rate beyond that level. Subjects' responses to a survey conducted at the end of the experimental study are consistent with this argument. When asked if, in the role of employers during a job interview, they would offer a lower piece rate than the one they initially thought to offer if they realized that the job candidate would really enjoy doing this job, only 22 subjects replied they would offer a lower piece rate, while 124 subjects replied they would not offer a lower piece rate, most of them for fairness reasons.

By matching subjects' survey responses with the decisions made in the experiment, I get the

¹⁷This also is confirmed by the data from the allocation game. On average, principals allocate to their preferred charity almost 3 times (11.7 points) as much as the agents allocate to a random charity in the list (4 points).

¹⁸Given a piece rate below the standard piece rate, only 8 percent of the observed effort levels are lower than the piece rate offered in the open contract.

following results. Regressions reported in Table A.1 show that in the main treatment those 14 percent of principals who do not perceive as unfair to offer a lower piece rate to the motivated job candidate are more likely in the open contract - but not in the closed contract- to offer a piece rate below the standard piece rate compared to the other principals ($p < 0.1$). This difference is equally significant using a clustered version of the rank-sum test proposed by Datta and Satten (2005). Furthermore, for this sub-sample of principals, Figures A.16 and A.17 show that: (i) In the main treatment, throughout all time periods, the average piece rate offered in the open contract is lower than the standard piece rate. (ii) In almost every period, the piece rate differential between treatments is higher in the open contract than in the closed contract.¹⁹ Given the sub-sample small size, and given that the experiment was not designed to test this specific hypothesis, this interpretation should be, however, taken with caution.

So the overall picture that emerges from this analysis is that principals are heterogeneous in the extent to which they are willing to take advantage of agents' intrinsic motivation by saving on wages. Most principals seemed to be willing to take advantage of agents' intrinsic motivation as long as the offered piece rate was still higher or equal than the standard piece rate. Contrary to the theoretical predictions, only one third of the principals were willing to offer a motivated agent a piece rate below the standard piece rate in the open contract. These results provide experimental support for previous survey evidence by Kahneman et al. (1986), which suggests that fairness considerations can act as a constraint on the employers' profits. In this experiment, fairness considerations acted as a limited liability constraint. This constraint made the substitution between monetary and non-monetary incentives *imperfect*.

V. CONCLUDING REMARKS

Recently, much attention has been devoted to study how agents respond to different monetary and non-monetary incentives.²⁰ There is no evidence, however, on how principals combine these incentives in order to induce effort. This paper reports the results from an experiment that was designed to test, for the first time, how principals write contracts when agents care about the mission of their job and when the principals have two instruments, the piece rate and the choice of the project mission, to influence the agents' effort. Therefore, this study also created a new paradigm that can be used for future laboratory experiments with motivated

¹⁹However, potentially due to the lack of power in such a small sample, I cannot establish statistical significance.

²⁰The latter include not only the job mission, but also status incentives and social recognition (Kosfeld and Neckermann, 2011; Ashraf et al., 2014), corporate social responsibility (Koppel and Regner, 2013), and the job impact (Kosfeld et al., 2014)

agents and misaligned mission preferences between principals and agents.

On the whole, my results provide evidence for the validity of the theoretical predictions from contract theory models with motivated agents (Besley and Ghatak, 2005; Delfgaauw and Dur, 2007; Cassar, 2014). I show that principals take advantage of agents' intrinsic motivation by lowering the piece rate and use the choice of the project mission as a substitute to motivate effort. Because of fairness considerations, however, this substitution remains imperfect.

But, of course, some questions remain open. First, the experiment is designed to test the validity of such predictions in a fully "neutral" environment, namely, in the absence of additional features of labor relations that may play a role in a natural environment. In reality, the matching of employers and workers is not exogenous, but rather, workers and employers with similar mission preferences will try to match and form a long-term contractual relationship. This selection and long-term interaction may, of course, enhance the development of social ties between the employers and the workers and, in turn, increase the effect of fairness considerations relative to a one-shot game. Whether these stronger fairness considerations would translate into higher wages or into better alignments of the job mission with the workers' preferred mission remains, however, an open question.

Finally, while the experiment did not allow for any competitive forces, in a natural environment competition between workers plays a crucial role in the determination of wages. More specifically, competition is likely to act against fairness considerations and reinforce the substitutability between monetary and non-monetary incentives: why should a non-profit organization hire a new employee if it can count on the effort of a motivated volunteer who is willing to work for free? These extensions are left for future work.

APPENDIX

Tables

Table A.1: Piece rate choices of principals with different survey responses

	(1)	(2)	(3)
	RE	Probit	RE
fair to pay lower wage	0.205*	0.574*	0.100
	(0.123)	(0.327)	(0.086)
constant	0.386***	-0.323	0.179**
	(0.093)	(0.271)	(0.070)
Round FE?	Yes	Yes	Yes
Session FE?	Yes	Yes	Yes
Individual FE?	No	No	No
Adj. R^2			
Observations	730	730	730

Standard errors are clustered at the individual level. Regression (1) and (2) include only observations within the main treatment and *open* contract. Regression (3) includes only observations within the main treatment and *closed* contract. The dependent variable is binary and takes value 1 if the offered piece rate is below 5 and value 0 otherwise. The independent variable is a dummy variable equal to 1 if, in the survey conducted at the end of the experiment, the principal reported that he would offer a lower wage to a motivated job candidate, and 0 otherwise. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

Figures

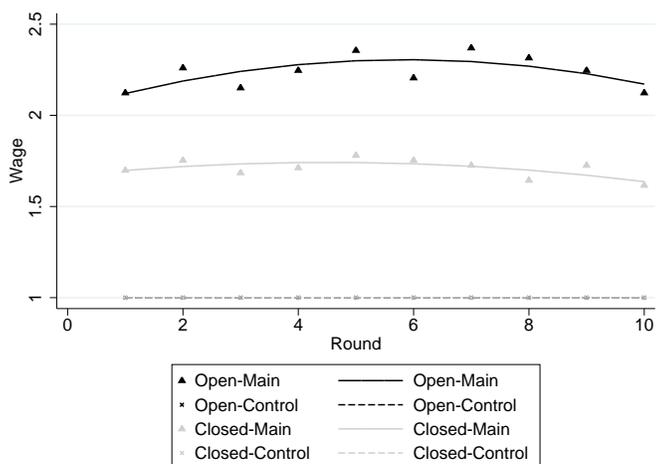


Figure A.1: **Effort over time for a piece rate= 1.** Averages shown by the dots for the respective type of contract and treatment, quadratic time trend shown by the lines for each type of contract and treatment. In total, 73 subjects participated in each type of contract and treatment.

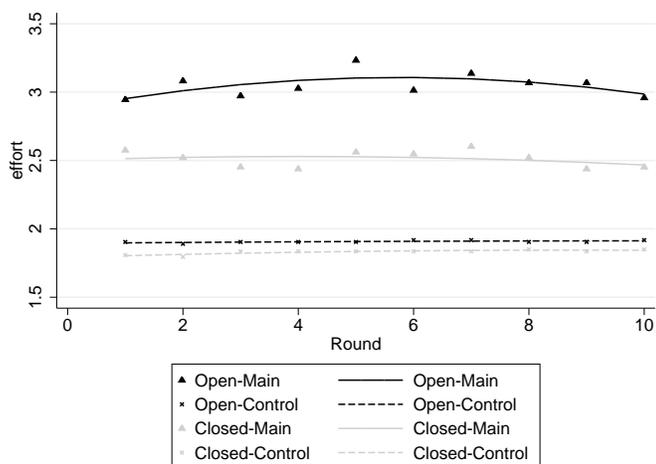


Figure A.2: **Effort over time for a piece rate= 2.** Averages shown by the dots for the respective type of contract and treatment, quadratic time trend shown by the lines for each type of contract and treatment. In total, 73 subjects participated in each type of contract and treatment.

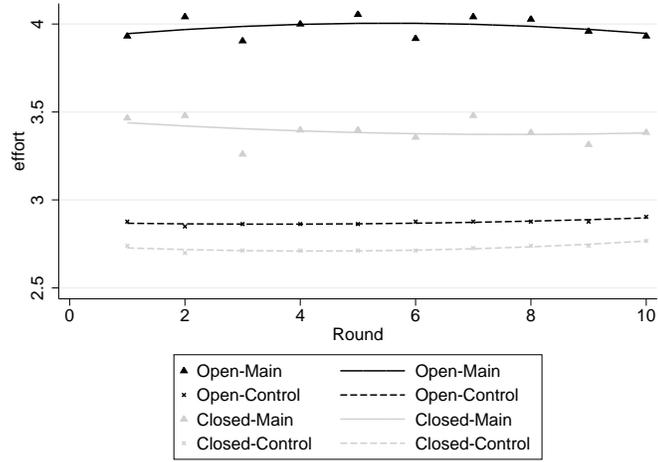


Figure A.3: **Effort over time for a piece rate= 3.** Averages shown by the dots for the respective type of contract and treatment, quadratic time trend shown by the lines for each type of contract and treatment. In total, 73 subjects participated in each type of contract and treatment.

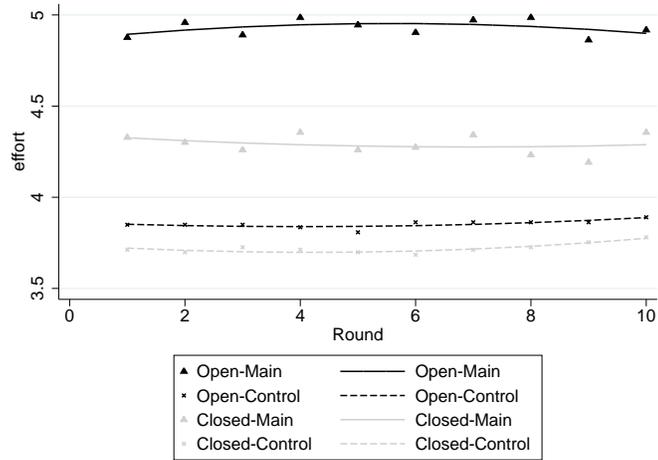


Figure A.4: **Effort over time for a piece rate= 4.** Averages shown by the dots for the respective type of contract and treatment, quadratic time trend shown by the lines for each type of contract and treatment. In total, 73 subjects participated in each type of contract and treatment.

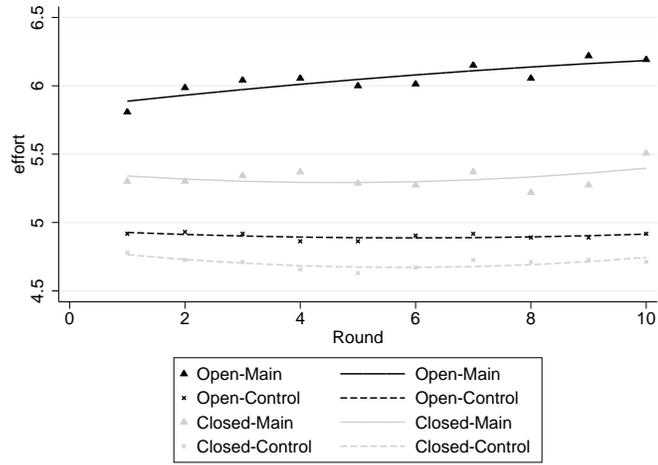


Figure A.5: **Effort over time for a piece rate= 5.** Averages shown by the dots for the respective type of contract and treatment, quadratic time trend shown by the lines for each type of contract and treatment. In total, 73 subjects participated in each type of contract and treatment.

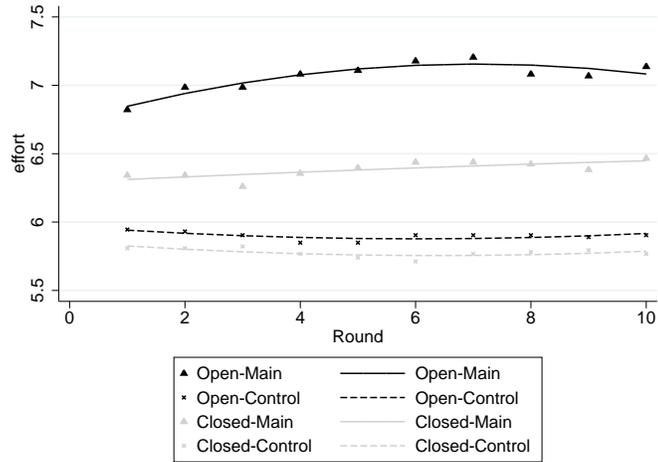


Figure A.6: **Effort over time for a piece rate= 6.** Averages shown by the dots for the respective type of contract and treatment, quadratic time trend shown by the lines for each type of contract and treatment. In total, 73 subjects participated in each type of contract and treatment.

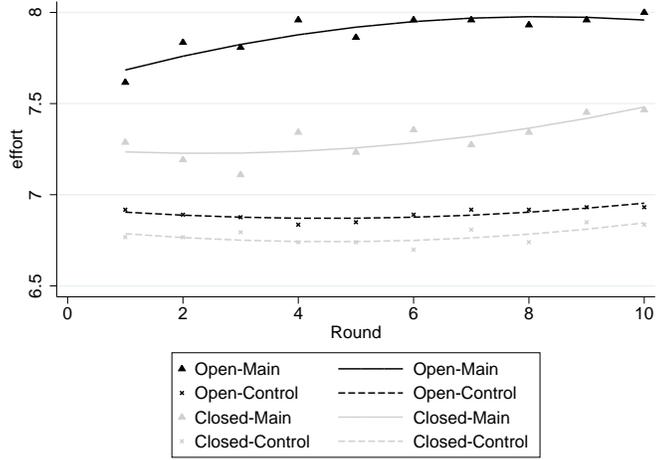


Figure A.7: **Effort over time for a piece rate= 7.** Averages shown by the dots for the respective type of contract and treatment, quadratic time trend shown by the lines for each type of contract and treatment. In total, 73 subjects participated in each type of contract and treatment.

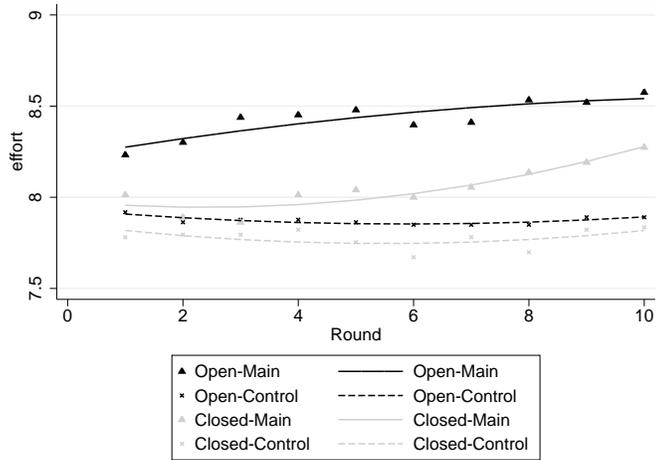


Figure A.8: **Effort over time for a piece rate= 8.** Averages shown by the dots for the respective type of contract and treatment, quadratic time trend shown by the lines for each type of contract and treatment. In total, 73 subjects participated in each type of contract and treatment.

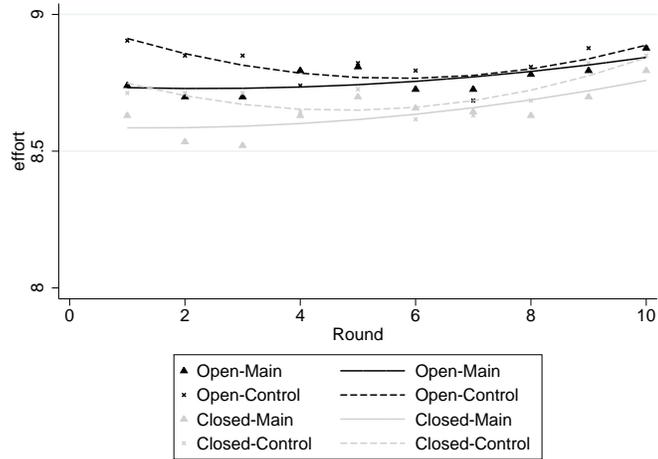


Figure A.9: **Effort over time for a piece rate= 9.** Averages shown by the dots for the respective type of contract and treatment, quadratic time trend shown by the lines for each type of contract and treatment. In total, 73 subjects participated in each type of contract and treatment.

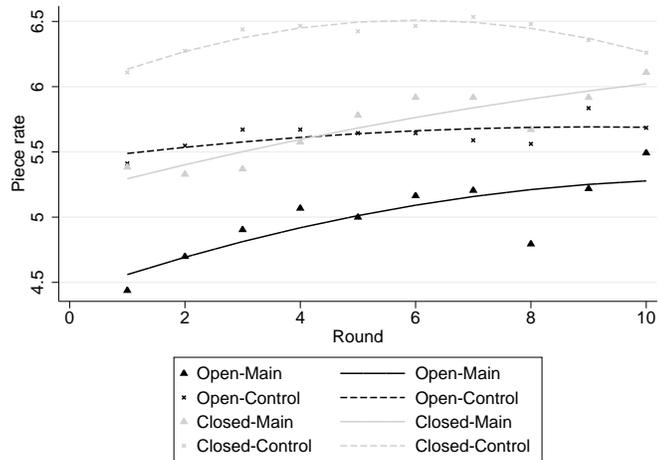


Figure A.10: **Piece rates over time.** Averages shown by the dots for the respective type of contract and treatment, quadratic time trend shown by the lines for each type of contract and treatment. In total, 73 subjects participated in each type of contract and treatment.

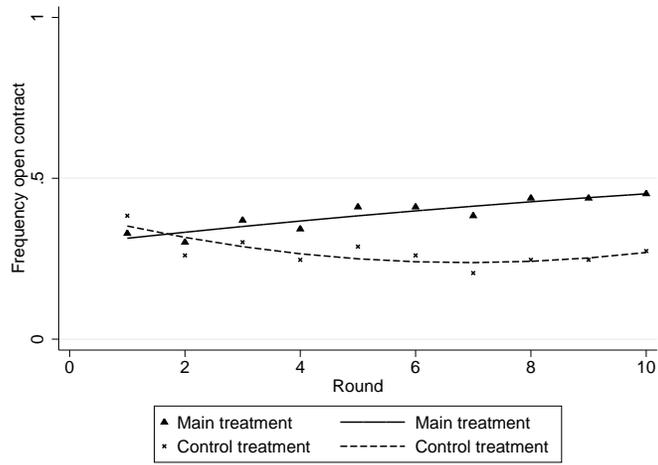


Figure A.11: **Contract choice over time.** Averages shown by the dots for the respective treatment, quadratic time trend shown by the lines for each treatment. In total, 73 subjects participated in each treatment.

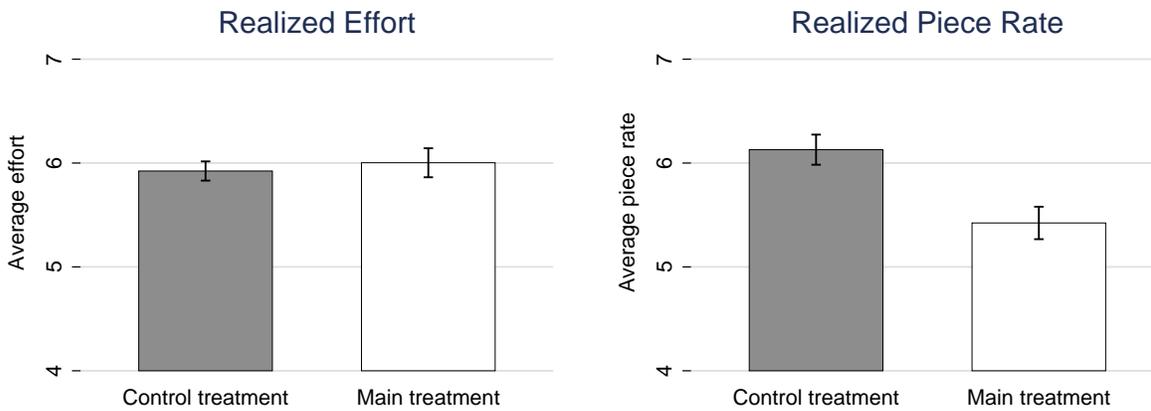


Figure A.12: Average realized effort and piece rate

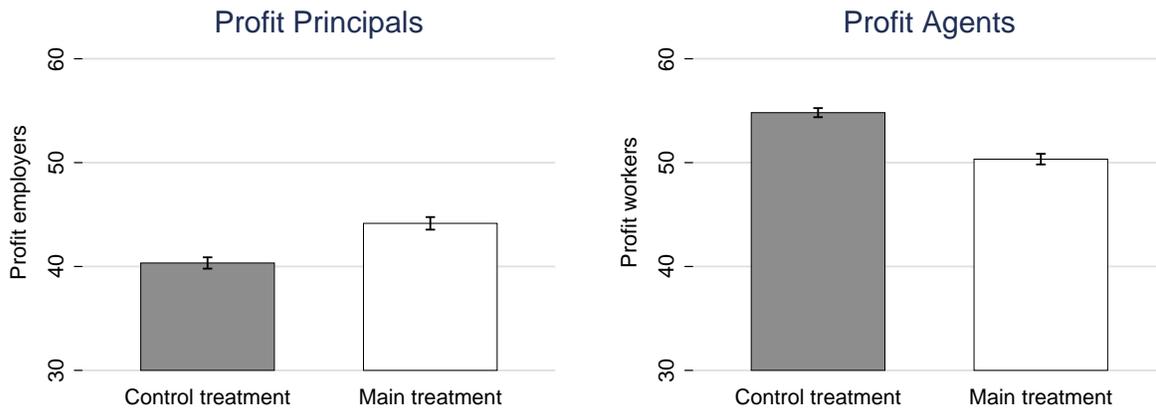


Figure A.13: Average realized profit

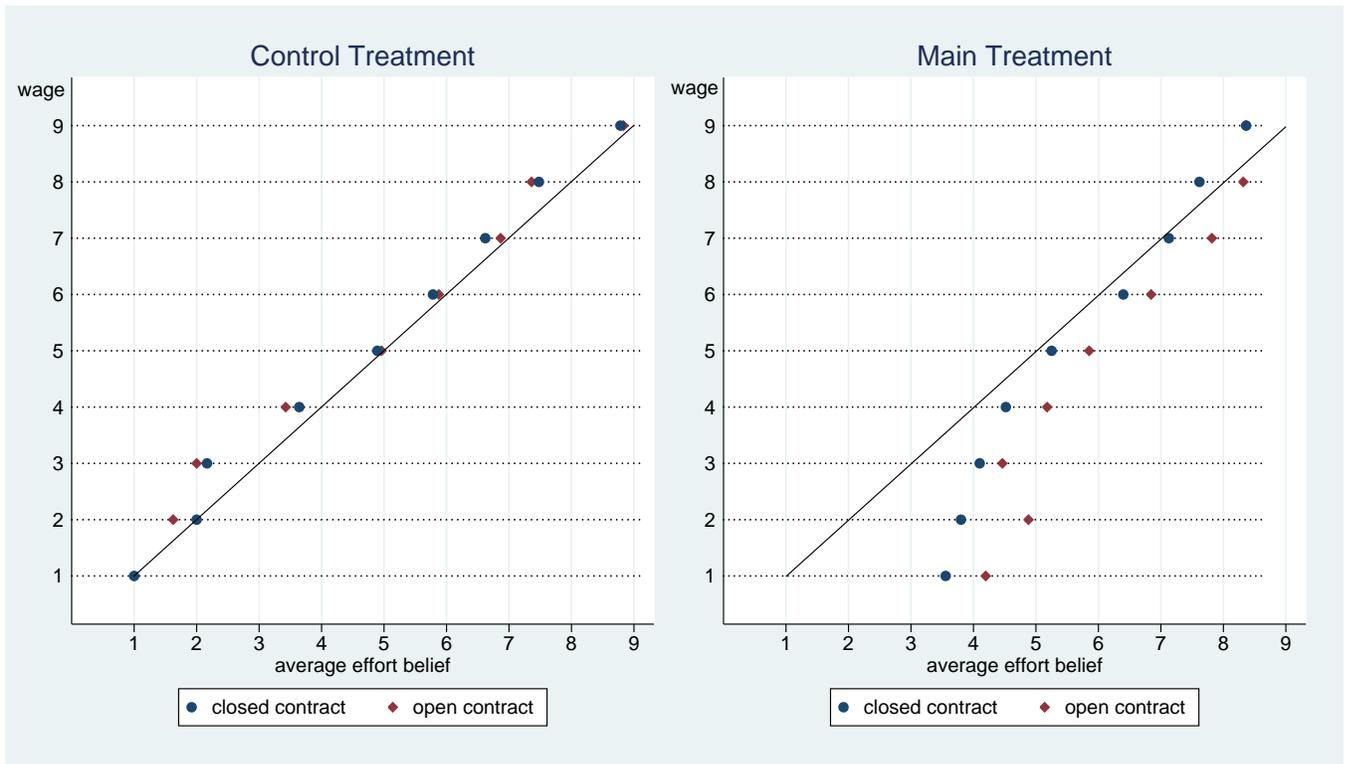


Figure A.14: Beliefs about effort

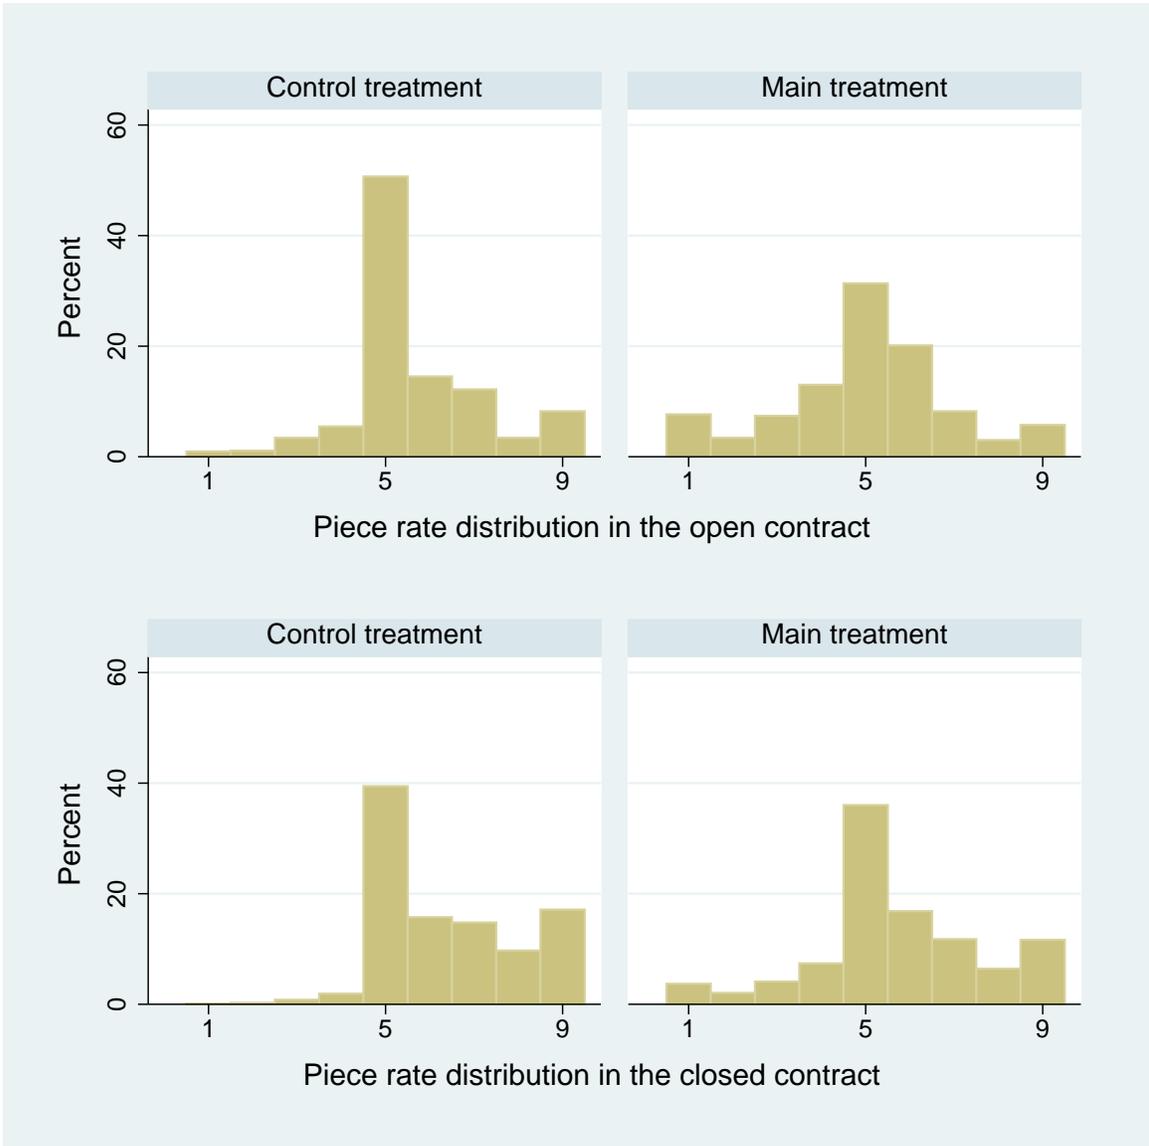


Figure A.15: Piece rate distribution

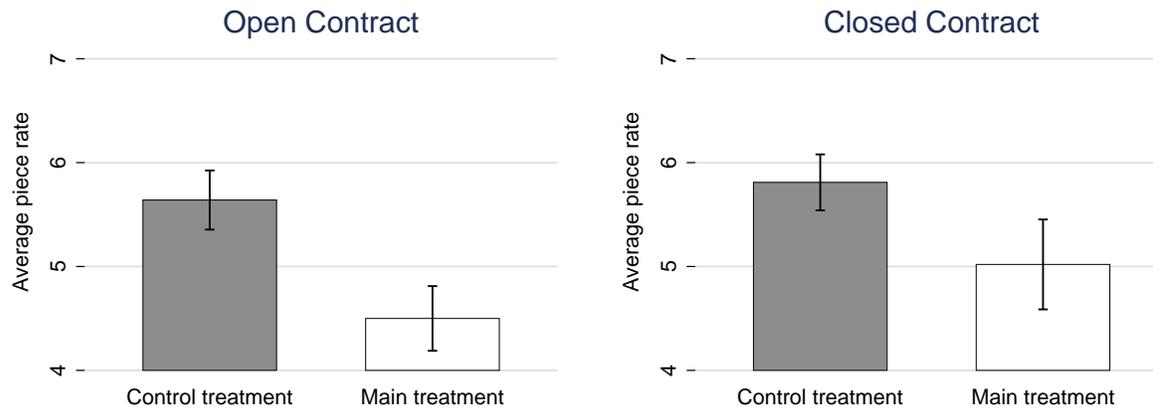


Figure A.16: **Piece rate difference across contracts for sub-sample of principals.** Average piece rate offered by 14% of principals ($n = 10$) who, in a survey conducted at the end of the experiment, said that they would offer a lower wage to motivated job candidate.

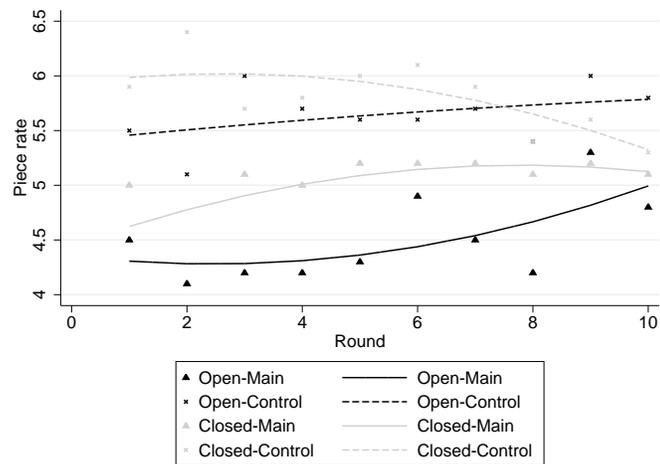


Figure A.17: **Piece rates over time for sub-sample of principals.** Averages shown by the dots for the respective type of contract and treatment, quadratic time trend shown by the lines for each type of contract and treatment. Only observations from the principals who, in the survey, reported that they would offer a lower wage to a motivated job candidate are included. In total, 10 subjects participated in each type of contract and treatment.

References

- Ashraf, N., O. Bandiera, and J. Kelsey (2014). No margin, no mission? a field experiment on incentives for pro-social tasks. *Journal of Public Economics*. Forthcoming.
- Bahney, B. W., R. K. Iyengar, P. B. Johnston, D. F. Jung, J. N. Shapiro, and H. J. Shatz (2013). Insurgent compensation: Evidence from iraq. *American Economic Review* 103(3), 518–22.
- Bartling, B., E. Fehr, and K. M. Schmidt (2012). Screening, competition, and job design: Economic origins of good jobs. *American Economic Review* 102(2), 834–64.
- Bartling, B., E. Fehr, and K. M. Schmidt (2013). Discretion, productivity, and work satisfaction. *Journal of Institutional and Theoretical Economics* 169(1), 4–22.
- Benabou, R. and J. Tirole (2003). Intrinsic and extrinsic motivation. *The Review of Economic Studies* 70(3), pp. 489–520.
- Besley, T. and M. Ghatak (1999). Public-private partnerships for the provision of public goods: Theory and an application to ngos. STICERD Development Economics Discussion Paper Series, DEDPS 17.
- Besley, T. and M. Ghatak (2001). Government versus private ownership of public goods. *Quarterly Journal of Economics* 116(4), 1343–1372.
- Besley, T. and M. Ghatak (2005). Competition and incentives with motivated agents. *American Economic Review* 95(3), 616–636.
- Besley, T. and M. Ghatak (2013). Profit with Purpose? A Theory of Social Enterprise with Experimental Evidence. Working paper.
- Bock, O., A. Nicklisch, and I. Baetge (2012). hroot: Hamburg registration and organization online tool. H-Lab Working Paper No. 1.
- Carpenter, J. and E. Gong (2013). Motivating agents: How much does the mission matter? Working paper.
- Cassar, L. (2014). Optimal contracting with endogenous project mission. Working paper.
- Charness, G., R. Cobo-Reyes, N. Jimenez, J. A. Lacomba, and F. Lagos (2012). The hidden advantage of delegation: Pareto improvements in a gift exchange game. *American Economic Review* 102(5), 2358–79.
- Datta, S. and G. Satten (2005). Rank-sum tests for clustered data. *Journal of the American Statistical Association* 471(1), 908–915.
- Datta, S. and G. Satten (2008). A signed-rank test for clustered data. *Biometrics* 64(2), 501–507.

- Delfgaauw, J. and R. Dur (2007). Signaling and screening of workers' motivation. *Journal of Economic Behavior & Organization* 62(4), 605–624.
- Falk, A. and M. Kosfeld (2006). The hidden costs of control. *The American Economic Review* 96(5), pp. 1611–1630.
- Fehr, E., H. Herz, and T. Wilkening (2013). The lure of authority: Motivation and incentive effects of power. *American Economic Review* 103(4), 1325–59.
- Fehr, E. and K. M. Schmidt (1999, August). A Theory Of Fairness, Competition, And Cooperation. *The Quarterly Journal of Economics* 114(3), 817–868.
- Fehrler, S. and M. Kosfeld (2013). Pro-social missions and worker motivation: An experimental study. Forthcoming in *Journal of Economics Behavior and Organization*.
- Fischbacher, U. (2007). z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics* 10, 171–178.
- Gerhards, L. (2013). How (not) to incentivize motivated agents - an experiment with employees from a non-profit organization. Working paper.
- Kahneman, D., J. L. Knetsch, and R. Thaler (1986). Fairness as a Constraint on Profit Seeking: Entitlements in the Market. *American Economic Review, American Economic Association* 76(4), 728–41.
- Koppel, H. and T. Regner (2013). Corporate social responsibility in the work place - experimental evidence on csr from a gift-exchange game. Working Paper.
- Kosfeld, M. and S. Neckermann (2011). Getting more work for nothing? symbolic awards and worker performance. *American Economic Journal: Microeconomics* 3(3), 86–99.
- Kosfeld, M., S. Neckermann, and X. Yang (2014). Profit with Purpose? A Theory of Social Enterprise with Experimental Evidence. Working paper.
- Nyborg, K. (2014). Is corporate social responsibility associated with lower wages? *IZA World of Labor*.
- Nyborg, K. and T. Zhang (2013). Is corporate social responsibility associated with lower wages? *Environmental and Resource Economics* 55(1), 107–117.
- Pischke, J.-S. (2011). Money and happiness: Evidence from the industry wage structure. NBER Working Paper No. 17056.
- Tonin, M. and M. Vlassopoulos (2014). Corporate Philanthropy and Productivity: Evidence from an Online Real Effort Experiment. Working paper.