Selected Presentation at the 2020 Agricultural & Applied Economics Association Annual Meeting, Kansas City, Missouri, July 26-28

Incentive Alignment and Reward Strength in Pay-for-Performance Contracts

Jose G. Nuno-Ledesma*

Abstract

I report an experiment on multitask agency theory. I evaluate hypotheses related to the effect of incentive misalignment on effort provision and per-period payoffs. Subjects are divided in two groups with different degrees of misalignment between performance and quality. These two variables determine the payoffs received by the agent and the principal, correspondingly. I vary the efficiency of the contract within subjects distorting its piece rate: switching to provide weak incentives in the mild misalignment group and strong incentives when the tension is severe. As expected, agents allocate effort in a way that increases the rewarded performance metric to the detriment of quality. Providing inefficiently powered incentives reduces profit although it remains over the point predictions in both misalignment groups. However, distorting the contract does not reduce the observed earnings of the agent when the performance-quality tension is mild. These outcomes have implications for the relative larger frequency of low-powered incentive contracts over high-powered agreements in the field.

PRELIMINARY. PLEASE, DO NOT CITE.

Keywords: Experiment; incentive alignment; pay-for-performance contracts; multiple tasks

JEL classification: C9, D8, J3

*jnuno@uoguelph.ca. Assistant Professor. University of Guelph. Department of Food, Agricultural & Resource Economics. J.D. MacLachlan Building. 50 Stone Road East. Guelph, Ontario, Canada N1G 2W1.

Data collection for this paper was conducted while I was a Ph.D. student at Purdue University. I thank Steven Y. Wu for valuable suggestions and input, and Tim Cason for insightful comments and for providing access to Purdue's Vernon Smith Experimental Economics Laboratory. I gratefully acknowledge financial support from the Jim and Neta Hicks Small Grants Program which is internally allocated by Purdue Univerity's department of Agricultural Economics. The data that support the findings of this study are available from the corresponding author upon reasonable request.

1 Introduction

Pay-for-performance (PFP) contracts or "merit pay" plans are often used with the intention to align the actions of an agent with the goals of a principal. In the typical PFP scheme, when quality is difficult to measure, the worker's payoff is partly determined by an easy-to-obtain metric that does not perfectly correlate with quality (the variable desired by the principal). The use of expedient evaluation criteria results in either the encouragement of activities irrelevant to quality, null or mild reward for other necessary activities, or a combination of these. Intuitively, the agent invests more on tasks closely measured by the performance metric regardless of their contribution to quality. For example, merit-based pay may lead teachers to devote excessive time to producing a single outcome, such as high standardized test scores, rather than focus holistically on the education process. In agricultural contracts, rewarding volume goals may cause suppliers to decrease focus on food safety.

Standard multitask agency theory describes an inverse relation between incentive provision (measured by the magnitude of the contract's piece rate) and the degree of performance-quality misalignment. In settings with mild misalignment, high-powered incentives ought to be the norm. In the field however, contracts with strong incentives are not as prevalent as expected (Williamson, 1985). In this context, my objective is twofold. Firstly, I contribute to the contract theory literature by providing a robustness check of comparative statics derived from the canonical multitask principal agent model. Secondly, I investigate an understudied aspect of PFP schemes and its role in the observed prevalence of contracts with low-powered incentives: the interaction between the strength of incentive provision and the performance-quality tension. Relying on data from an economic experiment, I find that when agents experience a low degree of misalignment and low-powered incentives (as opposed to the ideal high-powered contract), the profit generated does not drop as much as the theory predicts and the agent's payoff is not affected. Thus, I argue that contracts with mild incentive intensity are more likely to emerge in situations with high correlation between performance and quality.

The logic is simple: from the point of view of the principal, smaller profit from weaker-than-ideal incentives is preferred over the null benefit resulting from the deal not taking place. If the agent agrees to this contract, then a weakly powered arrangement is born. For an inefficiently-powered contract to emerge, the threat of discarding the efficient agreement should be credible. An instance when this is the case, is a contract that exposes the agent to exploitation. The agent might be required to invest in assets or skills that are not easily transferable. In these cases, an opportunistic principal may dis-honor promises or reduce payments finding justification within the gaps of an incomplete contract. Anticipating this behavior the parties may decide not to sign a contract. For example, a poultry grower (the agent) may fear exploitation by a large food processor (the principal) if the agent is required to invest in infrastructure that cannot be easily used for purposes other than growing chicken (e.g. growing houses). Because both parties anticipate the possibility of unfair treatment, the agreement may either not be a (simple incomplete) linear contract, or not take place at all. However, the deal may take place with weak incentives. For example, by decreasing the proportion of the agent's payoff determined by the use of the non-transferable asset.

Laboratory experiments are particularly useful to evaluate conclusions derived from principal-agent models. This is because agency theory occupies itself with studying the implications of the presence of private information on contract design. Consequently, researchers interested in developing empirical tests of agency theory are challenged by an inherent data availability problem that can be circumvented with help from controlled experiments. This methodological advantage has been pointed out by Hoppe and Schmitz (2015) and Huck et al. (2011).

In my experiment, human subjects are randomly assigned to one of two conditions with different degrees of tension between performance and quality: the first group with mild and the second with severe misalignment. In addition, each subject is exposed to both an efficient and a distorted performance-tied piece rate. This allows me to compare the distortions resulting from i) giving strong incentives to the agent when weak incentives ought to be provided and, conversely, ii) providing weak incentives when strong ought to be offered.

I find that the results largely corroborate the direction of theoretical predictions. With mild misalignment, more effort is exerted with the efficient piece rate than when it is distorted; performance and quality are also larger with the efficient rate. With severe misalignment, efforts, quality, and performance are larger with the distorted rate. Surprisingly, manipulating the contract to offer low-powered incentives when there is mild misalignment results in a lower-than-expected loss of profit but not a reduction in net wage. This is because the agents performance falls at a lower rate than predicted. In this situation, agents consistently exert more effort than is strictly necessary in both activities.

The rest of the paper is structured as follows. The next section provides a literature review; the following part provides and overview of the principal-agent model used to derive the testable hypotheses; the fourth section of the paper introduces the experimental design, lists the hypotheses, and explains the experimental procedures; the fifth part discusses the findings, and the last concludes.

2 Related Literature

As a result of the wide applicability of the theory, the literature leveraging multitask agency is abundant.[1] Notwithstanding the practical relevance of the theory reflected by its wide presence in academic research, there is limited direct testing of its main models. This paper contributes to the relatively small but growing set of works directly testing predictions from seminal contract-theoretic models. I begin this review by surveying the articles within this set that more closely relate to the work in this paper.

One early experimental study of the problem of hidden action with a single activity is DeJong et al. (1985), where the authors conduct an experimental investigation regarding how different institutional remedies ameliorate the moral hazard problem. Hoppe and Schmitz (2015) evaluate and find support for a fundamental hypothesis from adverse selection theory: that sellers separate demand from privately-informed buyers. Huck et al. (2011) provides the first experimental test of the theory of deferred compensation (Lazear, 1979), a labor economics model that nonetheless is intimately related to the theory of incentive design. In the first experimental study of multitasking principal-agent relationships, Fehr and Schmidt (2004) explore the influence of fairness on the decisions made by the contracting parties. Fehr and and Schmidt observe that when both principal and agent are human, subjects split gains from trade generously in accordance to the theory of inequity aversion (Fehr and Schmidt, 1999). My work is different from that of Fehr and Schmidt (2004) because, by design, I exclude the confounding effects of other-regarding behaviors. Thus, if agents act in a way that differs from theory predictions, this can be attributed to changes in either

[1] For articles discussing advances in the theoretical and empirical contract-theoretic literature, see Dewatripont et al. (2000) and Masten and Saussier (2000).

alignment or incentive intensity solely.

As a result of their work aiming to explain the phenomenon of comparable quality levels across piece rate and flat-compensation schemes, Al-Ubaydli et al. (2014) find evidence supporting the theory's intuition that workers systematically neglect under-rewarded but critical components of output. Using a natural field experiment, Hong et al. (2018) find that workers do increase their productivity in ways that make them more likely to receive larger monetary rewards to the detriment of difficult-to-measure quality dimensions. I obtain additional evidence in this direction. I find that agents consistently concentrate in tasks better reflected by the performance criterion while neglecting others. This is robust to changes in the degree quality-performance tension and to variations in the contract's piece rate magnitude.

Studying the interaction between degree of misalignment and the performance piece rate is important in the context of an observed prevalence of contracts offering low-powered incentives even when theory suggests that high-powered incentives ought to be more prevalent. Williamson (1985) argues that provision of weak incentives is the result of opportunism and incomplete contracts. Because in most cases it is impossible to write contracts that anticipate every contingency, there is always scope for opportunism causing some deals not to take place at all.

Contracts with weak incentives might emerge because the principal cannot oversee all possible actions taken by the worker. Holmström and Milgrom (1991) show that high-powered incentives may result in dysfunctional effort allocations depending on the principal's monitoring ability. With strong incentives and weak monitoring, the risk-averse agent shifts his effort towards activities better captured by the performance metric to the detriment of poorly-tracked tasks. Baker (1992) shows that low-powered incentive contracts emerge even when the agent is risk-neutral. Meng and Tian (2013) explore another explanation. They argue that the combination of moral hazard and adverse selection (such as heterogeneous agents differing in their level of risk aversion) can produce contracts with low-powered incentives.

I explore an additional factor that can explain the phenomenon: the degree of misalignment between quality and performance. In my experiment, when agents experience a low degree of misalignment and artificially low-powered incentives, profit does not drop as much as expected and net wage remains unaffected. I argue that contracts with mild incentive intensity are more likely to emerge in situations with high correlation between performance and quality. If providing weaker incentives results in a relatively small loss in earnings, this is preferred for both parties over not signing a contract and receiving no payoff.

3 Theory

Even though the contribution of this paper is not theoretical, it is useful to briefly discuss the main features of the two-task agency model from which I derive the testable hypotheses. I rely on the multitask model with noisy performance as introduced by Holmström and Milgrom (1991). The specific functional forms in this paper rely on Gibbons (2005). However, the agent in my discussion is risk-averse, unlike in Gibbons where the worker is risk-neutral.

Consider a relationship where the agent is expected to undertake two activities A and B. In the model, a_i represents the level of effort invested in task $i \in \{A,\ B\}$ and $a_A \geq 0$ and $a_B \geq 0$. The agent's total remuneration or wage is the sum of two components. The first component is constant while the second part varies with performance. Formally, wage can be written as follows:

$$w(p) = \alpha + \beta p \tag{1}$$

Where α is the constant part of the agent's remuneration and can take both positive and negative values. If it is positive, this component can be interpreted as a base salary paid to the worker regardless of performance. If negative, α could be read as a rent paid by the agent to the principal. The variable part of the contract is determined as the product of a piece rate β and a performance metric p which in turn is defined as:

$$p = \gamma_A a_A + \gamma_B a_B + \eta \tag{2}$$

Where the normally distributed random variable η has variance σ_η^2 and expectation $\mathbb{E}(\eta) = 0$. This noise represents events that impact the outcome of the performance metric and is beyond the control of either contracting party. The expenses incurred by the agent when exerting effort are determined by a well-behaved cost function $c(a_A, a_B)$. The agent's net wage subtracts cost from the remuneration: $U = w(p) - c(a_A, a_B)$. Importantly, the principal's earnings are defined as $\pi = [q - w(p)]$ and depend on both the agent's performance p and the quality variable q. Quality is formally defined as follows:

$$q = \delta_A a_A + \delta_B a_B + \varepsilon \tag{3}$$

Where the normally distributed random variable ε has variance σ_ε^2 and expectation $\mathbb{E}(\varepsilon) =$

0. This random variable affects the quality variable and cannot be manipulated by neither the principal nor the agent. The two noise terms η and ε are independent.

The vectors $\vec{\gamma} = [\gamma_A, \gamma_B]$ and $\vec{\delta} = [\delta_A, \delta_B]$ determine how much of a contribution each task makes to the performance metric and to quality, respectively. Incentive misalignment is absent when the vectors overlap (e.g. when $\gamma_A = \delta_A$ and $\gamma_B = \delta_B$) because the performance criterion perfectly captures the contributions to quality made by the agent's work. When this is not the case, tension between the principal and the agent emerges and the output generated by the worker does not exactly match the quality variable. Figure 1 illustrates a case when $\gamma_A \neq \delta_A$ and $\gamma_B \neq \delta_B$. In the figure, θ stands for the degree of misalignment between the performance outcome and quality. In the positive quadrant of the Cartesian system, a larger θ implies a poorer relation between performance and quality. In general, if $\gamma_i < \delta_i$, then the worker will neglect activity i and invest low levels of effort on it; on the other hand, if $\gamma_i > \delta_i$ the worker will tend to disproportionately concentrate his/her effort on task i. In the situation depicted in figure 1 for instance, the contract encourages a provision of effort in activity A beyond what is strictly required by the principal, to the detriment of activity B.

[Figure 1 about here]

I proceed now to discuss how the degree of misalignment directly influences the magnitude of both salary and piece rate (α and β, respectively) in the contract. When designing the compensation scheme, the principal needs to anticipate the agent's response to the contract. Taking α and β as given, the agent faces the following problem:

$$\underset{a_A,\, a_B}{\text{maximize}}\ V(a_A, a_B|\alpha,\ \beta) = -e^{-r[w(p)-c(a_A,a_B)]}\int_{-\infty}^{\infty} e^{-r\beta\eta}\phi(\eta)\mathrm{d}\eta \tag{4}$$

Where $\phi(\cdot)$ is the normal density function. To make sure that the agent will voluntarily decide to take part in the agreement, the principal needs to set both salary and piece rate such that the agent's Certainty Equivalence equals a reservation value which, for simplicity, I assume to be zero:

$$\text{CE}(a_A, a_B) = \alpha + \beta(\delta_A a_A + \delta_B a_B) - \frac{r}{2}\beta^2\sigma_\eta^2 - c(a_A, a_B) = 0 \tag{5}$$

The level of effort chosen by the agent when solving the utility maximization problem (4) is to be taken by the principal as the incentive compatibility constraints in his/her own maximization problem. If the cost function is $c(a_A, a_B) = \frac{1}{c}(a_A^2 + a_B^2)$, then the optimal levels of effort are $a_i = \frac{\beta}{2}\gamma_i c$ for $i \in \{A, B\}$. This results in the following maximization program for the principal:

$$\begin{aligned} &\underset{\alpha,\ \beta}{\text{maximize}} \quad \pi = (\delta_A a_A + \delta_B a_B + \varepsilon) - \alpha - \beta(\gamma_A a_A + \gamma_B a_B + \eta) \\ &\text{subject to:} \begin{cases} a_A = \frac{\beta}{2}\gamma_A c \\ a_B = \frac{\beta}{2}\gamma_B c \\ \text{CE}(a_A, a_B) = 0 \end{cases} \end{aligned} \tag{6}$$

Using equation 5, the salary α can be expressed in terms of the other parameters in the program. This reduces the dimensionality of the problem and leaves the piece rate β as the only choice variable for the principal. It is straightforward to find out that the efficiency maximizing piece rate is:

$$\beta^* = \frac{||\vec{\delta}||}{||\vec{\gamma}||}\cos(\theta) \tag{7}$$

Where $\|\cdot\|$ is notation for the Euclidean norm, and θ denotes the degrees of separation between the vectors $\vec{\gamma}$ and $\vec{\delta}$. There are two factors that interact to define β^*: scaling and misalignment. Scaling is captured by the ratio $\frac{||\vec{\delta}||}{||\vec{\gamma}||}$, and describes the form in which relative length of the vectors affects the magnitude of the piece rate. The larger $\vec{\gamma}$ is relatively to $\vec{\delta}$, the smaller the efficient piece rate is. Incentive misalignment is captured by $cos(\theta)$.

3.1 Parameters

Table 1 shows the set of parameters that subjects experience throughout the experimental sessions. The table also displays the optimal effort allocation, performance outcome, quality level, and final payoffs for the given parameters. The upper panel describes an environment where the incentive misalignment is mild ($\theta = 17.1°$), while the lower panel refers to a setting with severe misalignment ($\theta = 70.8°$). I hold scaling constant across treatments at $\frac{||\vec{\delta}||}{||\vec{\gamma}||} = 5.7$. For both misalignment conditions, the table includes parameters and outcomes for contracts with both the "efficient" and a non-efficient "distorted" piece rate. The efficient rate maximizes payoffs and is larger in the environment with lesser tension between performance and quality. The distorted rate in the condition with severe (mild) misalignment corresponds to the efficient rate for the environment with mild (severe) misalignment. Thus, the proportional change in the magnitude of the reward rate is comparable in both groups. In the experiment and this numerical exercise, the random variables η and ε are uniformly distributed over the range $[-10, 10]$. For an agent with constant absolute risk aversion and normally distributed risks, equation 5 holds exactly. In my experiment with this parameters, the relationship holds approximately. Uniform noise is a typical feature in laboratory studies because it facilitates the understanding of experimental instructions.[2]

[Table 1 about here]

I now turn to briefly explain the mechanism driving the numerical outcomes. With mild misalignment, the contribution to quality of activity B is about three times higher than that of activity A; however, the contract provides incentives in a manner that the agent is encouraged to exert only about 1.3 times more effort in activity B compared to A. As a result, the contract results in an under-provision of activity B relative to A.

The problem of poorly allocated effort is accentuated in the environment with severe misalignment. In this setting, task B contributes to quality about seven times more than task A, however, the contract rewards activity A over B in a way that one unit of effort in the earlier contributes five times more to the performance metric than a unit of effort invested in the latter. Thus, in this case, both the over-provision of effort in activity A and the neglect of task B are pronounced.

To observe the agent's effort investments when presented with inefficiently-powered contracts, I swap the piece rates across misalignment conditions. In the condition with mild performance-quality tension, I include a treatment where the scheme features a small piece rate. Similarly, for the group with severe misalignment, I add a treatment with an artifi-

[2]The impact of this simplifying feature is negligible. The assumption does not affect the direction of the changes in effort distribution, and the point predictions in all tables and figures come from the exercise with uniform noises. Neither of the testable hypotheses are affected by this choice.

cially large piece rate. In all cases, the salary is adjusted accordingly to keep the agent's participation constraint binding.

As shown in table 1, providing weak incentives when there is mild misalignment results in lower effort provision for both tasks, although the relative investments remains the same with more effort directed toward activity B relative to task A. Naturally, this results in reductions in both performance and quality, and consequently a sensible loss in profit for the seller. On the other hand, providing strong incentives when there is severe misalignment naturally results in higher effort exertion for both tasks, which translate into higher performance an quality; however, the greater quality does not offset the increase in the cost associated with the agent's better performance, thus the principal profit suffers a reduction.

After this brief description of the numerical example, I continue to discuss the experimental design and results. The objectives are to corroborate the directional outcomes from the model and observe whether changes in per-period payoffs (profit and net wage) are as drastic as expected. Evidence suggesting that subjects submit decisions as described by the numerical example are taken as corroboration of the model's comparative statics. Additionally, if the drop in per-period payoffs following the piece rate distortion is less pronounced in either of the misalignment conditions, I argue that such inefficiently-powered contracts are more likely to emerge in the field because these "errors" are less expensive than what the model's predictions would suggest, as long as the agent's payoff ensures participation.

4 Experimental Design and Hypotheses

In the experiment, two groups of subjects where exposed to a single misalignment condition each. Within subjects, I vary the magnitude of the piece rate. The role of the principal is automated. This is for two reasons: first, to facilitate causal identification by removing other-regarding behaviors and second, to allow for experimental manipulation of the contract's piece rate. This decision permits a greater degree of experimental control because it eliminates the confounding effects of inequity aversion, reputation, and other behaviors involving level-k thinking strategies. Between subjects, I can confidently attribute the differences in effort provision to the particular degree of misalignment subjects in that group encounter. Within subjects these differences are attributable to modifications in the contract terms.

Treatments are different in two ways: first, they either have mild (M) or severe (S) misalignment, and second, the exogenous contract is either efficient (E) or distorted (D). Thus, there are four treatments in total: ME, MD, SE, and SD (which stand for Mild-Efficient, Mild-Distorted, Severe-Efficient, and Severe-Distorted). The distorted piece rate in the mild misalignment condition corresponds to the efficient rate in the severe misalignment group, and vice versa. I pay close attention to the effort choices made by the agents, but also discuss performance outcomes, quality metrics, and the earnings of both parties. The testable hypotheses derived from the model are listed below.

Recall that the objectives of this paper are two: i) provide a robustness check of comparative statics from the multitask agency model described earlier, and ii) observe changes in payoffs when the piece rate is distorted, given a degree of misalignment. The first two hypotheses are related to my first objective.

Hypothesis 1: Effort. The ordering of the effort level dedicated to activity A is ME>MD in the environment with mild misalignment, and SE<SD in the environment with severe misalignment. Likewise, the corresponding orderings of the effort levels dedicated to activity B are ME>MD and SE<SD.

Hypothesis 2: Metrics. Regarding the level of performance, the orderings with mild and severe levels of misalignment are ME>MD and SE<SD, correspondingly. In what respects to generated quality, the orderings in the environments with mild and severe degrees of misalignment are ME>MD and SE<SD, respectively.

The third hypothesis is related to my second aim. In addition to corroborate the directional effects predicted by the model, the purpose of testing this hypothesis is to explore a reason for the prevalence of low-powered contracts in the field. For example consider the group with mild misalignment, if during the sessions with inefficient piece rates the observed profit remains over the point prediction but net wage remains unaffected, this would suggest that weakly-powered contracts could be more prevalent in the field because these inefficiently-powered agreements would not be as costly for the principal.

The principal earnings are half of the story, however. For an inefficiently-powered contract to be observed in the field, the agent's earnings are also important. In particular, the agent ought to receive . Thus, the combination of impacts on payoffs that would make an inefficiently-powered contract more likely to appear in the field is the following: a lower-than expected reduction in profit compared to the optimal contract, and either an increase or a null impact on wage. The parameter combination with which the experiment is run, however, favors a large decrease in both per-period profit and per-period wage in both misalignment conditions.

Hypothesis 3: Payoffs. The ordering of the expected profit is ME>MD in the environment with mild misalignment, and SE>SD in the environment with severe misalignment. On the other hand, the ordering of the wage is ME>MD in the environment with mild misalignment, and SE>SD in the environment with severe misalignment. Experimental payoffs are equal to the point predictions in table 1.

4.1 Procedures

All sessions were conducted between October and November 2018. Subjects were recruited from the general student population of a large American university via e-mail solicitations managed with ORSEE (Greiner, 2015). All sessions proceeded as follows: first, the experimenter distributed printed instructions to all subjects in the room and then read them aloud.[3] To test their understanding of the instructions, all participants were asked to answer a quiz and a monetary reward was provided if and only if the subject correctly answered all of the questions. Following the quiz, the subjects played 32 "decision rounds". These were divided into two blocks of 16 rounds each. One block of decision rounds presented subjects with contracts with the efficiency-maximizing piece rate, while the other block featured distorted contracts. To test for order effects, for each of the two incentive misalignment conditions, there were two sessions where the initial block of decision rounds presented subjects with efficient contracts and two sessions where the initial block featured the inefficient scheme. I found no order effects.[4] At the end of the session, subjects were paid according to the pro-

[3]For a copy of the instructions, feel free to contact the author.

[4]I used the Kolmogorov-Smirnov test for equality of distributions on the average effort choice for both activities for each subject. All p-values were considerably over 0.1, suggesting there is no evidence of order

cedure described below. On average, subjects earned \$24.30 including a \$5.00 participation fee and the quiz payment.[5]

All sessions were performed through computers and the experimental interface was implemented using oTree (Chen et al., 2016). All human subjects were assigned to take the role of the agent, while their computers offered exogenous contracts. The subject's decision occurred during "decision rounds". During the experimental sessions, the agent and the principal were named "purple" and "gold" player, respectively. The effort exerted in activity $i \in A, B$ was named "decision number i". Performance and quality were denominated "outcome" and "quality", correspondingly. Finally, the random variables η and ε were denoted by the letters p, and g, respectively.

At the beginning of a block, subjects were informed about contract terms (salary and piece rate). During each decision round, a given subject went through the following stages: first, the decision phase required the participant to input two integer numbers between 0 and 60: their "decision number A" and "decision number B". Following this phase, the random variables impacting performance and quality were drawn by the computer program, and the metrics were computed. The subject was then informed about payoffs and the value of all variables. After this, a new decision period started anew. Subjects had access to a physical "earnings-tracking sheet" where they could register their decisions and outcomes. In addition, during the decision phase period, subjects had access to an on-screen calculator where they could explore with different decision numbers and different values for the random variables; the calculator returned the costs associated with each decision number, total

effects.

[5] \$24.34 in the sessions with mild misalignment, and \$24.25 in sessions with severe misalignment. All cash payments were in USD.

cost (the sum of costs incurred), performance, quality, and the associated payoffs for both principal and agent.

In total, 88 human subjects participated in the experiment. Each participant generated 16 observations per block. Each subject participated only in one session. During a given session, there were eleven human subjects and they experienced only one degree of misalignment (either mild or severe). The materials used during the sessions, including sample instructions, are in the appendix.

I induce CARA utility with risk aversion coefficient of 0.002 using the method of Berg et al. (1986) and Roth and Malouf (1979). This method is fairly common in experiments involving contract-theoretic games. The general procedure is the following: during each decision period, subjects play for points. At the end of the session, two rounds for each block are randomly chosen (four rounds in total). The points earned by the participant during each of these periods are converted into a probability of winning a big prize of $4.50 or a small prize of $0. Thus, the maximum monetary payoff (excluding the participation fee and the quiz earnings) is $18.00. The points-to-probability conversion reflects the preferences the experimenter wants to induce. In this case, to induce risk aversion, the points-to-probability conversion is concave, making the first points more valuable. To facilitate the induction of a particular risk preference, this payment procedure allows me to use the optimal coefficients in each of the equations used for the experimental design.

5 Results

I organize the results around the hypotheses formulated in the section describing the experimental design. Recall that during half of the sessions, subjects were presented with the block of periods featuring the efficient contract first, followed by the environment with a distorted piece rate. For the rest of the sessions, this order was inverted. In all of the figures below, data from blocks with the same piece rate condition are appended at the point indicated by the vertical lines; thus, periods 1 to 16 in the graphs correspond to the rounds played with an efficient rate, while periods 17 to 32 correspond to the periods where participants were exposed to the distorted contract.

5.1 Hypothesis 1: Effort

With mild misalignment, the model predicts that effort in both activities drops when passing from an efficient to an inefficient contract. With severe performance-quality tension, on the other hand, I hypothesised effort in both tasks to increase when switching contracts. Regarding the relative magnitude of efforts given the chosen parameters, a distorted piece rate results in larger changes for activity B, with mild misalignment, and activity A, when the tension is severe.

These patterns are confirmed by the experimental data. Panels a) and b) in figure 2 show the effort allocated to both activities averaged across all subjects in the mild and severe misalignment conditions, correspondingly. The horizontal lines show the theoretically optimal levels; dashed for activity A, and the dash-dot pattern for B. In both misalignment conditions, throughout all the periods, the average effort provision for both activities is close to the point predictions and reacts as expected when changing the piece rate from efficient to distorted. This is true for all sessions. Table 2 shows the average effort per session. Clearly, subjects submitted efforts for both activities consistently close to the point predictions.

[Figure 2 and table 2 about here]

This directional result is further validated by the regression analysis shown in table 4. The columns report models where the observed level of effort provision is the dependent variable. The independent variables include an indicator variable taking the value of one if the observation belongs to a treatment with a distorted piece rate (*distorted*), a time trend indicating the period during which the effort was observed (*period*), and an intercept (*constant*). These first two columns report estimations using only data from sessions with mild misalignment, while the last columns use data from the condition with severe misalignment. The sign of the variable *distorted* denotes the direction of the effect of distorting the performance-tied rate, while the magnitude of the coefficient is the estimated impact of introducing an inefficient contract on effort allocation.

[Table 4 about here]

According to the estimates and as predicted by the theory, with mild misalignment the effort invested in both activities with inefficient contracts is significantly lower compared to the provisions made with optimal piece rates. Likewise, in the setting with severe incentive misalignment, agents increase their effort for both activities. The magnitude of the coefficient also aligns with the theoretical predictions. The numeric exercise indicates that introducing an inefficient contract would result in a larger drop in effort allocated in activity B, compared to A, with mild misalignment; on the other hand, a more pronounced increase in effort provision in activity A, compared to B, would be expected in the environment with severe misalignment. As shown in table 4, the econometric estimates closely align with these numerical predictions.

Albeit close, the submitted efforts are not exactly equal to the point predictions. The coefficients reported in table 4 help us estimate the effort exerted by the typical subject. To illustrate how the magnitude of the effects can be derived from the estimations, consider effort exerted in activity A by subjects in treatment ME. From the parameters reported in the first column of table 4, we can calculate that the effort exerted by the typical subject in period 8 (half the number of periods played in this condition) amounted to 17.92, while the theoretical optimal is 16.41.

Table 3 shows the estimated exerted efforts per treatment and the theory predictions in parentheses. With mild misalignment, subjects submitted i) more effort in activity A than expected, ii) less effort than predicted for activity B with efficient contract, and iii) more effort in activity B with the inefficient piece rate. On the other hand, with severe misalignment, subjects i) submitted effort decisions for task A not statistically different than the predicted when the contract featured an efficient piece rate, ii) exerted less effort in A than predicted with the distorted rate, and iii) submitted more effort than expected for activity B.

5.2 Hypothesis 2: Metrics

According to the model, with mild misalignment, both the performance outcome and the quality metric decrease following the implementation of an inefficient contract. On the other hand, providing high-powered incentives in the environment with severe misalignment is hypothesized to result in increased levels of both performance and quality.

Figure 3 shows average observed performance and quality metrics across periods for each treatment. Both, observed performance and quality in all four conditions with efficient contracts are close to the point predictions. However, with mild misalignment, performance and quality when the contract includes an inefficient piece rate do not fall to the exact predicted values in the majority of decision rounds.

[Figure 3 about here]

Table 5 reports econometric estimates of the impact of exposing participants to contracts with inefficient piece rates. The hypotheses are supported by the data. Adopting contracts with inefficiently weak incentives in the setting with mild misalignment results in significantly lower quality and performance levels. On the other hand, exposing subjects to contracts with artificially high-powered incentives when the degree of misalignment is severe causes an increase in both performance and quality metrics. The previously discussed improvement of effort allocation as the sessions continue in the environment with mild misalignment, and it is reflected in the gradual decline of both performance and quality in treatment MD.

[Table 5 about here]

Thus far, I have shown that the experimental data aligns with theoretical predictions regarding effort allocation, quality, and performance. Effort decisions are close but not statistically equal to the point predictions. Similarly, the more important deviation from the expected metric outcomes is that, with mild misalignment, performance and quality are substantially larger than expected in a number of decision rounds. In the following subsection, I address the implications that these patterns have for per-period payoffs.

5.3 Hypothesis 3: Payoffs

The model indicates that, with mild misalignment, profit earned by the principal falls after the adoption of a contract with an inefficiently low piece rate. Similarly, with severe misalignment, the principal's profit also decreases after the agent is presented with a contract featuring an inefficiently large piece rate. The expected effect on the agent's payoff is negative for both alignment conditions.

Figure 4 shows average payoffs across periods for each treatment. With efficient contracts, observed profit is close to the point prediction in both misalignment conditions. Profit falls when introducing an inefficient contract, albeit it remains consistently over the point predictions. Although the visual evidence is less compelling, the agent's per-period net wage seems to fall only for the condition with severe misalignment.

[Figure 4 about here]

Econometric estimates of the impact on per-period payoffs of distorting the piece rate are reported in table 6. As expected, the principal's profit decreases in both misalignment conditions. The decline is more pronounced with severe misalignment. Regarding the agent's earning, its reduction following the adoption of distorted contracts is not statistically significant when there is mild misalignment, but the impact is significant when the misalignment is severe.

[Table 6 about here]

When misalignment is mild, manipulating the bonus to offer low-powered incentives results in a lower-than-expected loss of profit because the agents performance falls at a lower rate than predicted. In periods with low-powered incentives, agents consistently exert more effort than the strictly necessary in both activities, although at significantly lower levels than during the periods with the efficient performance bonus. When the degree of misalignment is severe, theory recommends low-powered incentives. When I distort the performance bonus with the intention to reduce incentives, the principals profits do fall but remain above the theoretical prediction. This is because the average agents performance does not increase as much as expected.

Given these outcomes, what can be said about the prevalence of inefficiently-powered contracts in the field? I argue that milder-than-optimal incentives are more likely to emerge in situations with a weak tension between performance and quality. This is because providing weak incentives would result in a reduction in profit that nonetheless is less pronounced than what the theory predicts, and the wage received by the agent is comparable to what he/she would be paid with the efficiently-powered contract. This outcome is preferable over the alternative of no earnings for either party.

The other type of inefficiently-powered contract features strong incentives when there is a severe degree of misalignment between performance and quality. Agreements with these characteristics are less likely to be observed in the field not only because profit would be smaller (possibly negative, albeit over the point prediction), but because the agent's net wage would also be lower. The payoff received by the agent with an efficient contract satisfies his/her participation constraint, thus lowering it makes the no-participation alternative more attractive for the worker.

6 Conclusion

I report an experiment designed to evaluate comparative statics from the principal-agent model with multiple tasks and performance-based rewards. In two groups with different degrees of misalignment between performance and quality, I vary the efficiency of the contract within participants. Human subjects take on the role of agents facing exogenous contracts. The data largely support the model's directional predictions. With both efficient and inefficient contracts, subjects distribute effort across the activities in a way that enhances their performance outcome to the detriment of quality. This holds regardless of the degree of misalignment between performance and quality. In consequence, all outcome variables follow the directions predicted by the theory, providing evidence in support of the multitask principal-agent model. Effort is always exerted so as to maximize the rewarded performance metric (as opposed to the level of quality) even when the contract is modified to alter the provision of incentives and encourage either more effort in neglected tasks, or discourage effort from over-provided activities.

This exercise is informative in light of an observed prevalence of reward systems offering low-powered incentives in the field, even when theory suggests that high-powered contracts should be more common. The insight is simple and involves the payoffs of the principal and the agent. Starting with the principal, it is likely for profit to become negative (or at least drop significantly) when an artificially large piece rate is introduced in a setting with severe misalignment. When the tension between performance and outcome is mild however, there is more room for error. An artificially low piece rate is not as likely to eliminate profit completely. With respect to the agent, the efficient contract provides a payoff that makes him/her indifferent between participating or opting out from trade. With the distorted contract, the observed net wage does not fall in the condition with mild misalignment, but it decreases in the group with severe misalignment.

One natural extension of this work would be to include human subjects to adopt the principal's part. This project would investigate the interaction between degree of misalignment and other-regarding behaviors. For example, Letting the principal choose the incentive structure á la Al-Ubaydli et al. (2014) may illuminate a potential interplay between incentive misalignment and effort provision in activities not closely reflected in the piece rate. Another extension would be to limit the amount of effort that agents can exert. In the present experiment, there is no theoretical upper bound in the levels of effort agents can provide, imposing a limit may result in an exercise informative for cases where effort may be interpreted as time allocated to tasks, attention, or any other naturally constrained resource.

References

Al-Ubaydli, O., S. Andersen, U. Gneezy, and J. A. List (2014, February). Carrots that look like sticks: Toward an understanding of multitasking incentive schemes. *Southern Economic Journal*, 538–561.

Baker, G. P. (1992, June). Incentive contracts and performance measurement. *Journal of Political Economy 100*(3), 598–614.

Berg, J. E., L. A. Daley, J. W. Dickhaut, and J. R. OBrien (1986, May). Controlling Preferences for Lotteries on Units of Experimental Exchange. *The Quarterly Journal of Economics 101*(2), 281.

Chen, D. L., M. Schonger, and C. Wickens (2016, March). oTree - An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance 9*, 88–97.

DeJong, D. V., R. Forsythe, R. J. Lundholm, and W. C. Uecker (1985). A Laboratory Investigation of the Moral Hazard Problem in an Agency Relationship. *Journal of Accounting Research, 23*, 81.

Dewatripont, M., I. Jewitt, and J. Tirole (2000, May). Multitask agency problems: Focus and task clustering. *European Economic Review 44*(4-6), 869–877.

Fehr, E. and K. M. Schmidt (1999, August). A Theory of Fairness, Competition, and Cooperation. *The Quarterly Journal of Economics 114*(3), 817–868.

Fehr, E. and K. M. Schmidt (2004, October). Fairness and incentives in a multi-task principal-agent model. *Scandinavian Journal of Economics 106*(3), 453–474.

Gibbons, R. (2005, January). Incentives Between Firms (and Within). *Management Science 51*(1), 2–17.

Greiner, B. (2015, July). Subject pool recruitment procedures: organizing experiments with ORSEE. *Journal of the Economic Science Association 1*(1), 114–125.

Holmström, B. and P. Milgrom (1991, January). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, and Organization 7*(special), 24–52.

Hong, F., T. Hossain, J. A. List, and M. Tanaka (2018, April). Testing the Theory of Multitasking: Evidence from a Natural Field Experiment in Chinese Factories. *International Economic Review 59*(2), 511–536.

Hoppe, E. I. and P. W. Schmitz (2015, January). Do sellers offer menus of contracts to separate buyer types? an experimental test of adverse selection theory. *Games and Economic Behavior 89*, 17–33.

Huck, S., A. J. Seltzer, and B. Wallace (2011, April). Deferred compensation in multi-period labor contracts: An experimental test of lazears model. *American Economic Review 101*(2), 819–843.

Lazear, E. P. (1979). Why is there mandatory retirement? *Journal of political economy 87*(6), 1261–1284.

Masten, S. and S. Saussier (2000). Econometrics of Contracts: an Assessment of Developments in the Empirical Literature on Contracting. *Revue d'économie industrielle 92*(1), 215–236.

Meng, D. and G. Tian (2013, January). Multi-task incentive contract and performance measurement with multidimensional types. *Games and Economic Behavior 77*(1), 377–404.

Roth, A. E. and M. W. Malouf (1979). Game-theoretic models and the role of information in bargaining. *Psychological Review 86*(6), 574–594.

Williamson, O. E. (1985). *The economic institutions of capitalism: firms, markets, relational contracting*. Free Press.

Tables and Figures

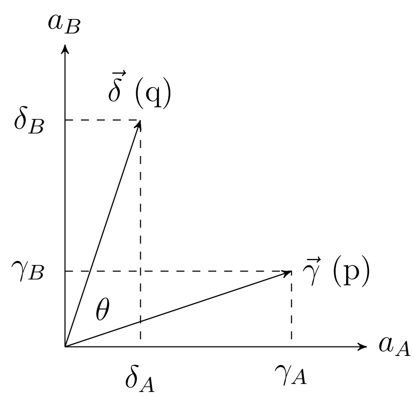

Figure 1: Graphical depiction of misalignment between performance (p) and quality (q).

Table 1: Parameters used in the experiment

Environment	Parameters	Efforts	Metrics	Payoffs
Mild misalignment				
Efficient piece rate	$p = 3a_A + 4a_B + \eta$	$a_A = 16.41$	$p = 136.74$	Profit = 374.93
	$q = 9.7a_A + 27a_B + \varepsilon$	$a_B = 21.88$	$q = 749.86$	Net Wage = 1
	$\eta, \varepsilon \sim$ Unif$[-10, 10]$			
	$\alpha = -372.93,\ \beta = 5.47$			
Distorted piece rate	$p = 3a_A + 4a_B + \eta$	$a_A = 5.65$	$p = 47.09$	Profit = 213.76
	$q = 9.7a_A + 27a_B + \varepsilon$	$a_B = 7.53$	$q = 258.22$	Net Wage = 0.12
	$\eta, \varepsilon \sim$ Unif$[-10, 10]$			
	$\alpha = -44.22,\ \beta = 1.88$			
Severe misalignment				
Efficient piece rate	$p = 10a_A + 2a_B + \eta$	$a_A = 18.83$	$p = 205.87$	Profit = 165.74
	$q = 8a_A + 58a_B + \varepsilon$	$a_B = 3.77$	$q = 369.15$	Net Wage = 18.95
	$\eta, \varepsilon \sim$ Unif$[-10, 10]$			
	$\alpha = -184.34,\ \beta = 1.88$			
Distorted piece rate	$p = 10a_A + 2a_B + \eta$	$a_A = 54.69$	$p = 568.82$	Profit = −484.55
	$q = 8a_A + 58a_B + \varepsilon$	$a_B = 10.94$	$q = 1072.01$	Net Wage = 1
	$\eta, \varepsilon \sim$ Unif$[-10, 10]$			
	$\alpha = -1554.56,\ \beta = 5.47$			

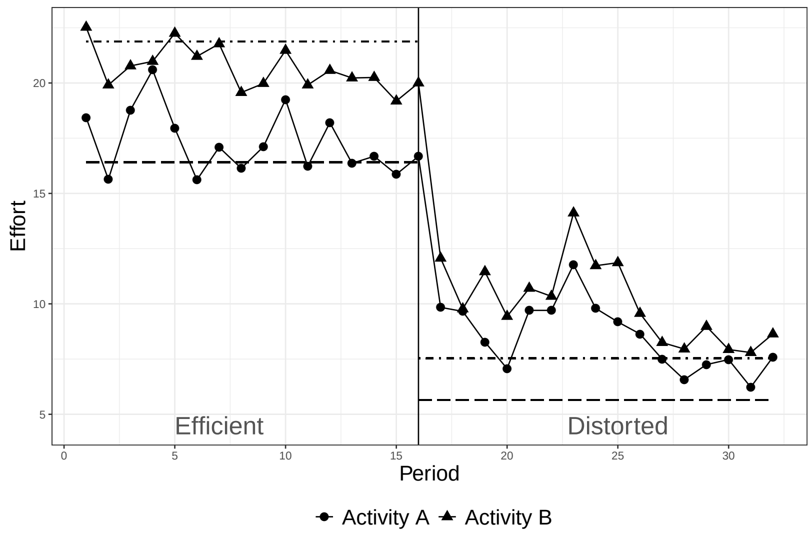

(a) Mild misalignment

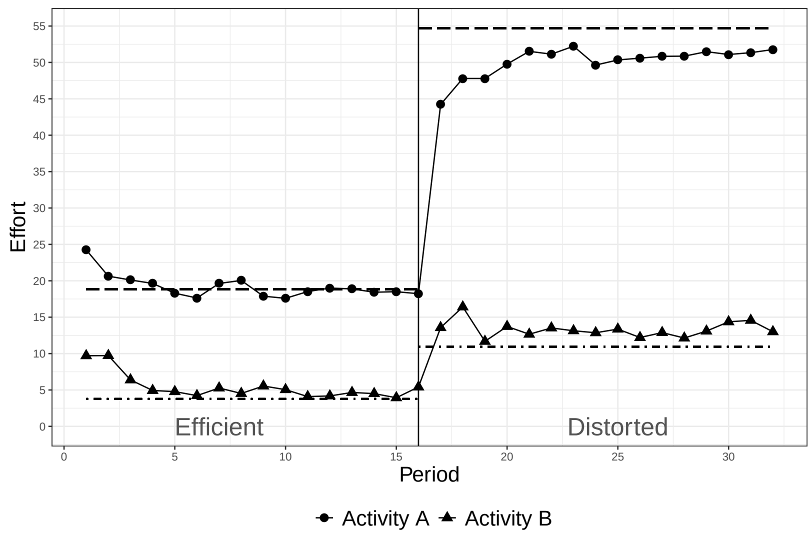

(b) Severe misalignment

Figure 2: Per-period average effort in activities A and B by degree of misalignment. Dashed lines correspond to theoretical optima.

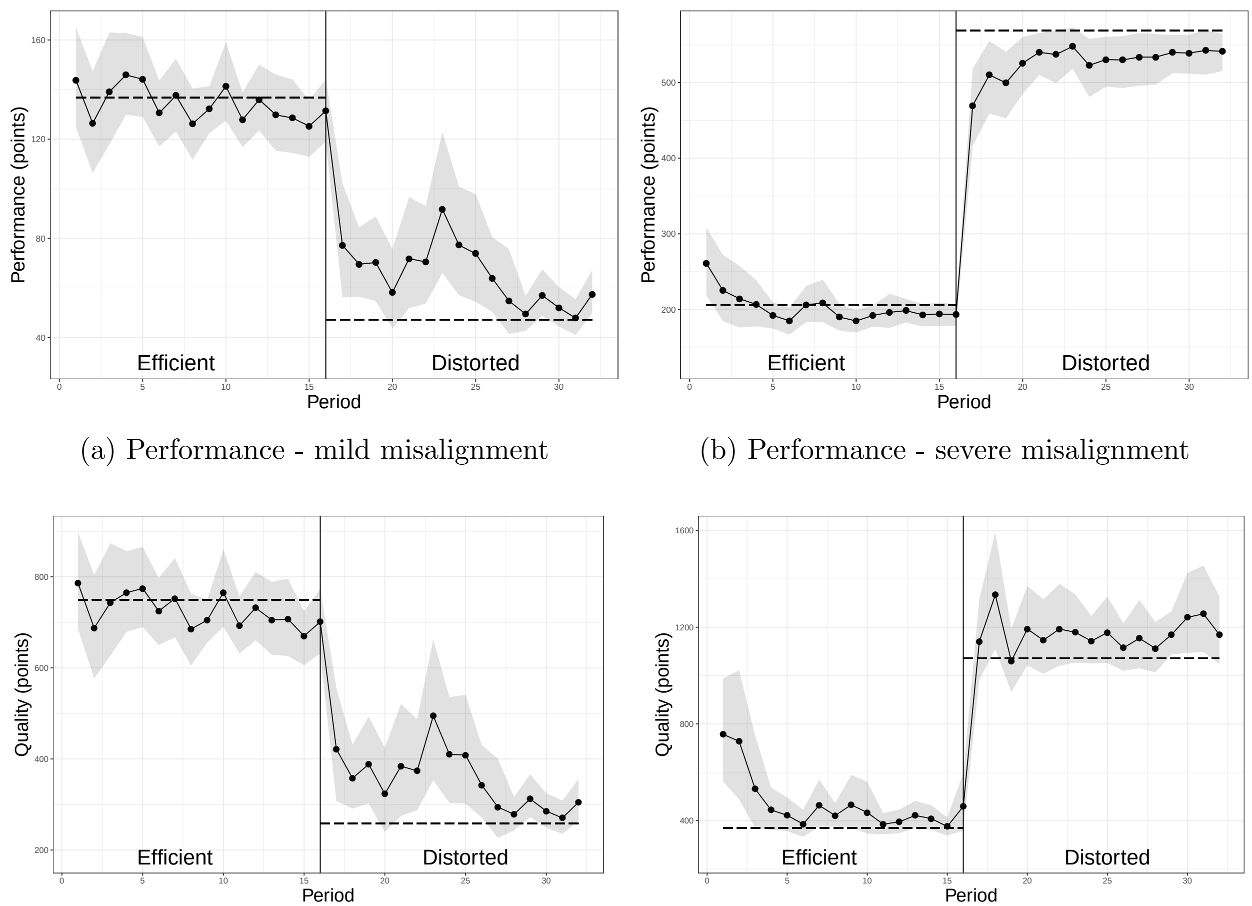

(a) Performance - mild misalignment

(b) Performance - severe misalignment

(c) Quality - mild misalignment

(d) Quality - severe misalignment

Figure 3: Per-period average metrics by degree of misalignment. Dashed lines correspond to theoretical optima. Shaded area is the 95th percentile bootstrap confidence interval.

Table 2: Session-level average effort by treatment

		Efficient		Distorted	
Misalignment	Session	Activity A	Activity B	Activity A	Activity B
Mild	**Theory**	**16.41**	**21.88**	**5.65**	**7.53**
	Session I	16.91	20.48	10.24	11.68
	Session II	17.10	21.15	8.24	10.27
	Session III	17.98	21.63	7.73	9.46
	Session IV	17.14	19.36	7.84	8.72
	All sessions	17.28	20.66	8.51	10.03
Severe	**Theory**	**18.33**	**3.76**	**54.69**	**10.94**
	Session I	19.70	3.87	51.71	11.52
	Session II	21.01	6.71	49.74	11.35
	Session III	17.32	6.46	49.05	15.02
	Session IV	18.78	4.67	50.09	15.38
	All sessions	19.20	5.43	50.15	13.32

Table 3: Observed efforts by treatment

Treatment	Activity A	Activity B
ME	17.36 (16.41) $*$	20.75 (21.88)$*$
MD	8.58 (5.65) $***$	10.13 (7.53) $***$
SE	19.17 (18.83)	5.50 (3.77)$***$
SD	50.12 (54.69) $***$	13.39 (10.94) $***$

Theoretical predictions in parentheses. Estimations for the 8th period.$*$ and $***$ denote the difference between the observed value and the theory benchmark is statistically different at the 5% and 1%, respectively.

Table 4: Impacts of incentive distortion by degree of misalignment

	Mild misalignment		Severe misalignment	
	Activity A	Activity B	Activity A	Activity B
Distorted	-8.772***	-10.623***	30.946***	7.892***
	(0.451)	(0.476)	(0.502)	(0.383)
Period	-0.147***	-0.177***	0.042	-0.136***
	(0.048)	(0.051)	(0.054)	(0.041)
Constant	18.541***	22.171***	18.841***	6.589***
	(0.631)	(0.679)	(0.989)	(0.752)
N	1408	1408	1408	1408

$*$ Pr $<$ 0.1, $**$ Pr $<$ 0.05, $***$ Pr $<$ 0.01. Models estimated using multi-level random effects (at the session and subject levels). Robust standard errors clustered at the session level. The explanatory dummy variables *Distorted* denotes whether the observation belongs to the treatment with distorted bonus rate. Period is a time trend.

Table 5: Impacts of incentive distortion by degree of misalignment - Metrics

	Mild misalignment		Severe misalignment	
	Quality	Performance	Quality	Performance
Distorted	-371.686***	-68.955***	705.372***	325.251***
	(15.633)	(2.887)	(22.783)	(5.128)
Period	-6.277***	-1.196***	-7.535***	0.148
	(1.695)	(0.313)	(2.471)	(0.556)
Constant	778.251***	144.291***	532.423***	201.197***
	(22.231)	(4.097)	(39.573)	(9.095)
N	1408	1408	1408	1408

$*$ Pr $<$ 0.1, $**$ Pr $<$ 0.05, $***$ Pr $<$ 0.01. Models estimated using multi-level random effects (at the session and subject levels). Robust standard errors clustered at the session level. The explanatory dummy variables *Distorted* denotes whether the observation belongs to the treatment with distorted bonus rate. Period is a time trend.

Table 6: Impacts of incentive distortion by degree of misalignment - Per-period payoffs.

	Mild misalignment		Severe misalignment	
	Principal's profit	Agent's wage	Principal's profit	Agent's wage
Distorted	-89.270***	-8.528	-430.35***	-74.053***
	(8.087)	(14.101)	(27.836)	(12.877)
Period	-2.668***	7.289***	-12.964***	10.064***
	(0.877)	(1.529)	(3.019)	(1.396)
Constant	386.873***	-145.469***	382.292***	-149.943***
	(11.234)	(22.547)	(69.263)	(26.341)
N	1408	1408	1408	1408

$*$ Pr $<$ 0.1, $**$ Pr $<$ 0.05, $***$ Pr $<$ 0.01. Models estimated using multi-level random effects (at the session and subject levels). Robust standard errors clustered at the session level. The explanatory dummy variables *Distorted* denotes whether the observation belongs to the treatment with distorted bonus rate. Period is a time trend.

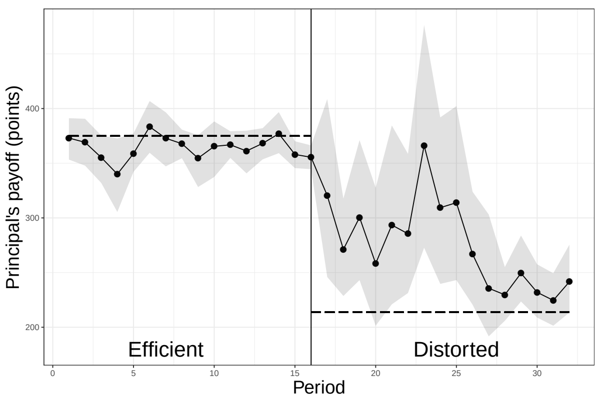

(a) Principal's profit - mild misalignment

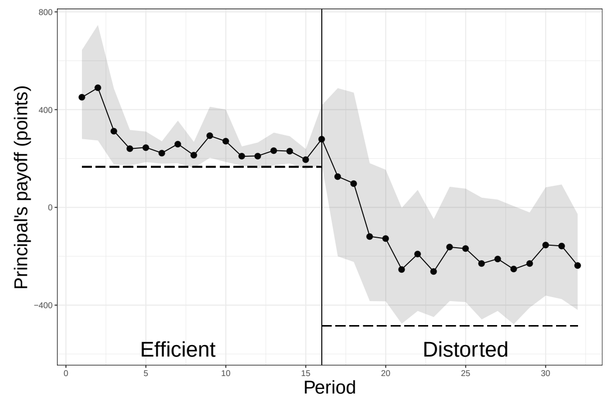

(b) Principal's profit - severe misalignment

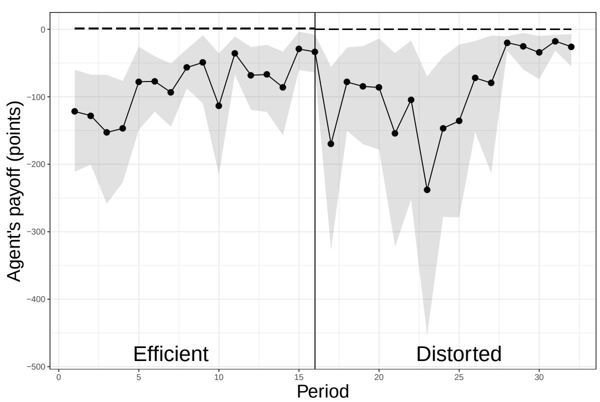

(c) Agent's net wage - mild misalignment

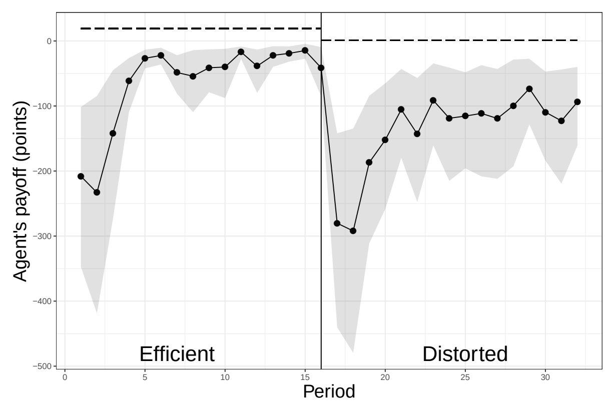

(d) Agent's net wage - severe misalignment

Figure 4: Per-period average payoffs by degree of misalignment. Dashed lines correspond to theoretical optima. Shaded area is the 95th percentile bootstrap confidence interval.