

Are You Paying Your Employees to Cheat? An Experimental Investigation

C. Bram Cadsby*

Department of Economics

University of Guelph

Fei Song

Ted Rogers School of Business Management

Ryerson University

Francis Tapon

Department of Economics

University of Guelph

*We would like to thank the Social Sciences and Humanities Research Council of Canada for generous research support through grants 410-2001-1590 and 410-2007-1380. We are also grateful to J. Atsu Amegashie, Jeremy Clark, Jim Cox, and Bradley Ruffle for very helpful comments and to Amy Peng for help with the statistical analysis.

Are You Paying Your Employees to Cheat? An Experimental Investigation

ABSTRACT

We compare misrepresentations of performance under a target-based compensation system with those under both a linear piece-rate system and a tournament-based bonus system using a laboratory experiment with salient financial incentives. An anagram game was employed as the experimental task. Results show that productivity, defined as the number of correct words a participant created during the seven experimental rounds, was similar and statistically indistinguishable under the three pay-for-performance schemes. In contrast, whether one considers the number of over-claimed words, the number of work/pay periods in which over-claims occur, or the number of participants making an over-claim at least once, target-based compensation produced significantly more cheating than either of the other two systems. Moreover, consistent with Schweitzer et al. (2004), cheating is more likely under a target-based scheme the closer a participant's actual production is to the target. The larger amounts of cheating under target-based compensation support Jensen's (2003) argument that such schemes encourage cheating and should be eliminated in favor of other types of performance pay.

JEL Classification Codes: C91, J33, M52.

Keywords: Misrepresentation, cheating, guilt, experiment, compensation, target, tournament, piece-rate, pay-for-performance.

How people make decisions involving compliance with ethical guidelines and organizational regulations has been an important focus of research in such diverse fields as economics, philosophy, psychology, accounting, law, and management. Prior work has identified a number of important factors that affect such compliance decisions by individuals (see Ford and Richardson, 1994, and Loe et al., 2000, for comprehensive reviews). These determinants include gender (Ambrose and Schminke, 1999; Glover et al., 2002), self-presentation concerns (Covey et al., 1989), stage of moral development (Trevino and Youngblood, 1990), ethical framework (Schminke et al., 1997), social norms (Donaldson and Dunfee, 1994, Gino and Bazerman, 2009; Gino, Ayal and Ariely, 2009), organizational culture (Chen et al., 1997, Pierce and Snyder, 2008), ethical training (Delaney and Sockell, 1992), the use of ethics codes (Trevino and Youngblood, 1990; Weaver et al., 2000), attitudes and behavior of friends and relatives (Schminke et al., 2002), and the presence of wealth and perceptions of inequity (Gino and Pierce, 2009a, 2009b).

A large literature based on seminal work by Becker (1968) and Ehrlich (1973) relates both compliance with and enforcement of the law to economic costs and benefits. Much of this literature focuses on the relationship between enforcement mechanisms and crime. The application of this model to cheating within organizations has been dubbed the “rational cheater” model (Nagin *et al.*, 2002). In a fascinating field experiment, Nagin *et al.* (2002) find evidence that some employees of a telephone solicitation company respond to a reduction in monitoring with an increase in cheating, while others, perhaps motivated by conscience or guilt, do not. Rickman and Witt (2007) reach similar conclusions in their study of employee theft in the UK.

A considerable theoretical and empirical literature on tax evasion applies the “rational cheater” model to examine the relationship between the decision to evade and such enforcement mechanisms as the audit rate, audit selection methods, and the penalty if caught (e.g., Alm *et al.*, 1990; 1992a; 1992b; 1992c; 1993a; 1993b; 1995; 1999; Beron *et al.*, 1992; Boylan and Sprinkle, 2001; Cadsby *et al.*, 2006; Feld and Tyran, 2002; Moser *et al.*, 1995). At the same time, this literature also discusses other factors affecting incentives to evade or comply with one’s tax obligations such as the tax rate or the use to which tax revenues are put. Alm and McKee (1998) provide an excellent review of this literature, and argue that many of the results from laboratory experiments on tax compliance are directly applicable to organizational compliance with regulations and compliance with regulations within organizations. For example, experimental work on the effects of different enforcement mechanisms can be applied to the use of analogous schemes by regulatory authorities and within organizations.

However, compliance within organizations does not depend solely on enforcement mechanisms. It also depends on the incentives created by an organization’s compensation

practices to act in accordance with or to disregard company regulations. The literature on tax evasion has considered how incentives created by different tax rates or tax systems may affect compliance, but this is not directly applicable to the analogous issue of how compensation systems may create incentives that tempt employees to cheat.

Production is not always easy to observe and pay is often based upon employee reports of hours worked or tasks accomplished. For example, lawyers, accountants and business consultants are often paid based on self-reported billable hours. Automobile and appliance service technicians charge customers based on their own diagnosis of the problem and of the resultant repairs. Similarly, physicians in many countries are paid based upon their own diagnosis of illness and the resultant treatment. Many executives are paid based on the financial performance of their organizations, which in turn can be manipulated by false or misleading reports. Nagin et al. (2002), as mentioned above, discuss a case in which telephone canvassers soliciting money for non-profit organizations receive commissions based on their self-reports of contribution pledges.

Recently, in response to several business scandals associated with false sales reports to obtain rewards under goal-setting compensation systems (DeGeorge, *et al.*, 1999; Jensen, 2001), Michael Jensen (2003) has argued controversially that the use of production or sales targets in compensation formulas encourages people to lie or misrepresent their performance with serious consequences for firm productivity and profitability. Urging that such targets be replaced by linear pay-for-performance compensation systems in which people are rewarded in direct proportion to their productivity, he asserts:

“Everyone can benefit by bringing this game to an end, and I believe it starts by eliminating the use of targets in compensation systems, and in particular by eliminating the use of budgets as targets in compensation systems. Simply put this means creating linear pay-for-performance compensation systems” (Jensen, 2003, p. 405).

However, it is not obvious that adopting a linear pay-for-performance (henceforth PFP) compensation system would really give people incentives to report their performance truthfully. As long as people are paid on the basis of performance, linearly or otherwise, they may still have an incentive to exaggerate their performance. Indeed, it is possible that a linear PFP system would encourage bigger lies about the number of items produced or sold within a budgetary period. If one is close to a target under a target-based system, one need claim only to have produced or sold a few more items to reach the target, thereby obtaining a large financial bonus. To obtain a similarly increased payoff under a linear piece-rate system, one might have to make far more exaggerated claims relative to actual performance. Such exaggerated claims could damage the sales and production planning processes, perhaps even more seriously than under a target-based

system.

A recent study by Schweitzer *et al.* (2004) examined target-based production and compensation systems in the laboratory. They showed that a target-based system produces more lies about performance than simply paying people a lump sum and asking them to do their best. Grover and Hui (2005) also reported that people are more likely to over-report their performance level when performance pressure is created by linking financial rewards to the achievement of performance goals. More generally, several scholars in management and economics have warned that goal-setting as a management strategy may induce unintended, undesirable, and sometimes dire consequences including unethical behavior (e.g., Barsky, 2008; Ordóñez, Schweitzer, Galinsky, and Bazerman, 2009a, 2009b). These claims have proven highly controversial as evidenced by a recent vigorous exchange of views in the pages of *Academy of Management Perspectives* (Ordóñez *et al.*, 2009a, 2009b; Locke and Latham, 2009). The one thing that all of these authors agree on is that systematic scholarly research is the best way to shed further light on such issues.

To our knowledge no empirical study has compared cheating under a target-based compensation system with cheating under the linear piece-rate system favored by Jensen (2003). Nor has cheating under a target-based system been compared to cheating under another popular alternative, a tournament-based system. The purpose of this study is to compare the exaggerations and misrepresentations that occur under a target-based compensation system with those that occur under both a linear piece-rate and a tournament-based bonus setting by means of a controlled laboratory experiment with salient financial incentives.

Before eliminating target-based in favor of alternative PFP compensation systems, it is important to examine whether doing so will actually reduce misrepresentation. This is difficult to do in an actual business setting due to the hidden nature of misrepresentation and the many uncontrollable factors that might affect misrepresentation in the field. In contrast, a well-designed laboratory experiment allows us to observe directly the degree of misrepresentation under the three compensation systems— target-based, linear piece-rate, and tournament—while controlling for other confounding factors.

The next section outlines the theoretical motivation for the study, utilizing a simple illustrative model of the benefits and costs of cheating. This is followed by a section outlining the experimental methodology and another discussing the experimental results. A conclusion follows.

1. Theory

A Simple Illustrative Model

The Jensen hypothesis that target-based compensation encourages cheating and

misrepresentation relative to linear PFP is based on implicit assumptions about the cost of cheating to individuals and its relationship to the amount by which they cheat.¹ We illustrate this by constructing a simple model of cheating behavior. This model is not intended to encompass all possibilities, but rather has the more modest objective of illustrating some circumstances under which Jensen's arguments are correct and others in which they are not. Like the models of Becker (1968) and Ehrlich (1973), our model compares the benefits of cheating with the costs. In contrast to Becker and Ehrlich, the costs in our model are psychological costs such as the guilt experienced as a result of cheating rather than the expected costs of being caught and punished. Our model bears some resemblance to the one presented by Nagin et al. (2002). However, we focus more explicitly on the precise relationship between guilt and the amount of cheating and its interaction with the compensation system. In order to focus on this relationship, we do not include any system of monitoring, enforcement, or punishment in either our model or our experiment.

Suppose that an individual is working at a job that rewards each employee based on the number of self-reported units produced within a given time period. This may be thought of as a three-stage game. In stage one, the employee decides how much effort to exert. Individual output, q , is determined by a production function $q = f(e, \varepsilon)$, where e is effort and ε is a random shock. The random shock represents the possibility of being tired or alert, distracted or focused, or any other random factor that could have an impact on the transformation of effort into performance during a particular time period. In stage two, the person finds out q , the amount s/he has produced. In stage three, the person decides whether and by how much to misrepresent his/her performance. This paper focuses on the stage-three misrepresentation decision conditional on the realized level of output, q .²

Let c represent the number of over-claims (c stands for cheating) made by the individual in question. Over-claims may be beneficial in that under a PFP system, higher output leads to higher pay. Higher pay in turn leads to higher utility. In particular, the utility of the financial payoff is given by $U[P(q+c)]$, where P is the monetary payoff contingent on the reported performance level. The precise form of $P(q+c)$ is determined exogenously by the payment scheme. This will be the treatment variable in our experimental design. In all cases, $P(0) = 0$. For simplicity, we normalize $U[P(0)] = 0$. $U'[P(q+c)] > 0$ and $U''[P(q+c)] < 0$ by assumption.

$G(c)$ represents the disutility resulting from any psychological costs that may be

¹ Although costs of cheating are not discussed explicitly in Jensen (2001, 2003), his argument is not consistent with costless cheating. If cheating were costless, linear piece-rate schemes would produce unlimited amounts of cheating rather than a reduction in cheating as Jensen argues.

² The experimental results find no significant differences between output levels under the three schemes.

associated with cheating. For simplicity of exposition, we refer to such costs collectively as guilt. Guilt is modeled as a function of the number of over-claimed output units, c . $G(0) = 0$ since guilt arises only if cheating occurs. $G(c)$ is allowed to be discontinuous at 0. This allows for the possibility that for some people even a tiny amount of cheating results in a large amount of guilt. However, it is assumed to be continuous elsewhere. $G'(c) \geq 0$ for $c > 0$, indicating that guilt does not decrease as the amount of cheating rises.³ $U[P(q+c)]$ and $G(c)$ are assumed separable.

In the target-based setting, $P(q+c)$ is discontinuous. Suppose a person produces q_t , and q_t is less than the preannounced target, t . Then $P(q_t+c) = 0$ if $c+q_t < t$ and $P(q_t+c) = B$ if $q_t+c \geq t$, where B is a bonus received contingent on achieving the target. When faced with a decision about whether or not to cheat, an individual compares the benefits of the bonus with the psychological cost of the guilt. Define $c^* = t - q_t$. Then if $U(B) > G(c^*)$, the benefits exceed the costs and the individual will over-claim c^* units. In contrast, if $U(B) < G(c^*)$, the costs exceed the benefits and the individual will not cheat. When $U(B) = G(c^*)$, the person is indifferent, and the decision may go either way. Notice that if a person produces an amount greater than or equal to the target, $c^* \leq 0$, and B is received even in the absence of cheating. Hence, there is no opportunity to cheat for financial benefit in this instance.

In any group of people, $G(c^*)$ is likely to differ between individuals since different people will generally have different guilt responses to a given number of over-claims. However, since $G'(c) \geq 0$ for each individual, a given person with an unknown guilt function is more likely to cheat by making c^* over-claims, the closer s/he is to the target, i.e. the smaller is c^* . This is because a smaller c^* implies less guilt. This prediction has already received empirical support in the work of Schweitzer et al. (2004). We reexamine this issue in our setting.

Hypothesis 1: Under a target-based compensation scheme, cheating is more likely to occur the closer one is to the predetermined target.

In the linear piece-rate setting, $P(q_1+c) = k \cdot (q_1+c)$ is the monetary payoff resulting from reported performance, where k is the amount paid per unit of reported output. If $U[k \cdot (q_1+c)] - U(k \cdot q_1) < G(c)$ for all $c > 0$, there will be no cheating. In contrast, if $U[k \cdot (q_1+c)] - U(k \cdot q_1) > G(c)$ for some $c > 0$, cheating will occur at the level of $c > 0$, \hat{c} , that maximizes $U[k \cdot (q_1+c)] - U(k \cdot q_1) - G(c)$, thus providing the greatest possible net gain. In particular, there will be an interior maximum at \hat{c} where $kU'[k \cdot (q_1+\hat{c})] = G'(\hat{c})$, i.e. where the marginal benefit from cheating just equals the marginal psychological cost of guilt if $k^2 \cdot U''[k \cdot (q_1+\hat{c})] - G''(\hat{c}) < 0$, and it will be a

³ No restrictions are placed upon $G''(c)$ because it seems plausible for the marginal disutility of cheating to either rise or fall with the amount of cheating.

global maximum if $k^2 \cdot U''[k \cdot (q_1 + c)] - G''(c) < 0$ for all c associated with gains from cheating. If $U[k \cdot (q_1 + c)] - U(k \cdot q_1) > G(c)$ for some c and $k^2 \cdot U''[k \cdot (q_1 + \hat{c})] - G''(\hat{c}) > 0$ for all such c , gains from cheating will rise without limit as cheating increases. Notice that, in contrast to the target-based setting, cheating in the linear piece-rate setting always leads to a higher payoff regardless of the amount actually produced.

Comparing Cheating under the Target-Based and Linear Piece-Rate Settings

When comparing target-based and linear piece-rate compensation schemes, we assume that $B = k \cdot t$ and therefore that $B \geq k \cdot c^*$ for all possible c^* since, by definition, $c^* = t - q_1$. This simply means that the bonus under the target-based scheme is equal to the amount one would earn if one were to report exactly the targeted amount under the piece-rate scheme. The purpose of this assumption is to make the size of the compensation package under the two schemes equivalent for purposes of comparison. Nothing in the analysis that follows changes if $B > k \cdot t$.⁴

Proposition 1: If it is more beneficial for a person to over-claim c^* units than to over-claim zero units in the linear piece-rate setting, it will also be beneficial to over-claim c^* units in the target-based setting. However, the converse is not necessarily true.

Proof: Over-claiming c^* units in the linear setting in preference to over-claiming zero units implies $U[k \cdot (q_1 + c^*)] - U(k \cdot q_1) > G(c^*)$. As noted above, $B \geq k \cdot c^*$. Thus, $U(k \cdot q_1 + B) - U(k \cdot q_1) \geq U[k \cdot (q_1 + c^*)] - U(k \cdot q_1) > G(c^*)$. However, $U(B) = U(k \cdot q_1 + B) - U(k \cdot q_1)$ if $q_1 = 0$ and $U(B) > U(k \cdot q_1 + B) - U(k \cdot q_1)$ if $q_1 > 0$ since $U'' < 0$ by assumption.⁵ Thus, $U(B) > G(c^*)$ and c^* units will be over-claimed in the target setting.

The converse need not be true. Over-claiming c^* units in the target-based setting implies $U(B) > G(c^*)$. However, since $U(B) \geq U[k \cdot (q_1 + c^*)] - U(k \cdot q_1)$ as demonstrated above, this does not necessarily mean that $U[k \cdot (q_1 + c^*)] - U(k \cdot q_1) > G(c^*)$. Hence over-claiming c^* units in the target-based setting does not imply that an individual would over-claim c^* units in the linear piece-rate setting.

Intuitively, the financial incentives to over-claim c^* units of output are at least as high and generally higher in the target-based than in the linear case. This is perhaps the basis for Jensen's prediction that targets lead to substantially more cheating than linear pay systems. However, it is important to note that this general prediction is not implied by the theoretical model. Although a preference for over-claiming c^* rather than zero units in the linear case, but

⁴ A firm with a given amount of money available for compensation would in fact set $B > k \cdot t$ when moving between schemes. We set $B = k \cdot t$ so as to present as tough an empirical challenge as possible to the Jensen hypothesis discussed below. Setting $B > k \cdot t$ would only increase the temptation to cheat under the target-based scheme.

⁵ This is true because $U(k \cdot q_1 + B) - U(k \cdot q_1)$ falls as q_1 rises. In particular, the derivative of this expression with respect to q_1 is $k[U'(k \cdot q_1 + B) - U'(k \cdot q_1)] < 0$ when $U'' < 0$.

not in the target-based case, is ruled out by our simple model, every other conceivable outcome is possible, depending upon the shapes and positions of the $U[P(q+c)]$ and $G(c)$ functions, which in turn depend on individual attitudes towards earnings and guilt. Thus, it is impossible based on the theoretical model alone to predict whether target-based or linear financial incentives will lead to more cheating.

Table 1 indicates nine possible cases. In case *a*, there is no cheating in either the target-based or linear cases. This occurs when $U(B) < G(c^*)$ and $U[k \cdot (q_1+c)] - U(k \cdot q_1) < G(c)$ for all c . In case *b*, there is no cheating in the linear case, but cheating in the target-based case. This occurs when $U(B) > G(c^*)$, while $U[k \cdot (q_1+c)] - U(k \cdot q_1) < G(c)$ for all c . In case *c*, there are positive but fewer over-claims in the linear case than in the target-based case. This case occurs when: (i) $U(B) > G(c^*)$, (ii) $U[k \cdot (q_1+c)] - U(k \cdot q_1) > G(c)$ for some c , (iii) $k \cdot U'[k \cdot (q_1+\hat{c})] = G'(\hat{c})$ at $\hat{c} < c^*$, and (iv) $k^2 \cdot U''[k \cdot (q_1+\hat{c})] - G''(\hat{c}) < 0$ at $\hat{c} < c^*$. In case *d*, there are c^* over-claims in both the linear and target-based cases. It occurs under identical conditions to case *c* except that $\hat{c} = c^*$. In case *e*, there are more over-claims in the linear case than in the target-based case. It occurs under identical conditions to case *c* except that $\hat{c} > c^*$.

In case *f*, there are unbounded over-claims in the linear case, and c^* over-claims in the target-based case. This requires $U(B) > G(c^*)$ and $U[k \cdot (q_1+c)] - U(k \cdot q_1) > G(c)$ for some c . In addition, either $k \cdot U'[k \cdot (q_1+c)] > G'(c)$ for all c , or if $k \cdot U'[k \cdot (q_1+c)] = G'(c)$ for any c , $k^2 \cdot U''[k \cdot (q_1+c)] - G''(c) > 0$ at that point. In case *g*, there is no cheating in the target-based case, but cheating at $\hat{c} < c^*$ in the linear case. This occurs when: (i) $U(B) < G(c^*)$, (ii) $U[k \cdot (q_1+c)] - U(k \cdot q_1) > G(c)$ for some c , (iii) $k \cdot U'[k \cdot (q_1+\hat{c})] = G'(\hat{c})$ at $\hat{c} < c^*$, and (iv) $k^2 \cdot U''[k \cdot (q_1+\hat{c})] - G''(\hat{c}) < 0$ at $\hat{c} < c^*$. In case *h*, there is no cheating in the target-based case, but cheating at $\hat{c} > c^*$ in the linear case. It occurs under identical conditions to case *g* except that $\hat{c} > c^*$. Finally, in case *i*, there is no cheating in the target-based case, but unbounded over-claims in the linear case. This requires $U(B) < G(c^*)$ and $U[k \cdot (q_1+c)] - U(k \cdot q_1) > G(c)$ for some c . In addition, either $k \cdot U'[k \cdot (q_1+c)] > G'(c)$ for all c , or if $k \cdot U'[k \cdot (q_1+c)] = G'(c)$ for any c , $k^2 \cdot U''[k \cdot (q_1+c)] - G''(c) > 0$ at that point.⁶

Different people will generally possess different $U[P(c)]$ and $G(c)$ functions, representing different psychological reactions to changes in earnings and to different levels of cheating. At different times, $0 < c^* \leq t$ will also differ even for the same individual depending on his/her performance in a given pay period. If most people belong to cases *b* and/or *c*, Jensen's hypothesis

⁶ Each of these cases is consistent with some possible individual utility and guilt functions regardless of whether or not actual output under the target-based scheme, q_t , is equal to output under the linear scheme, q_l . In fact, the experimental results showed no significant productivity differences under the two schemes.

will hold and cheating will be greater under target-based than under linear piece-rate compensation schemes. Figure 1 illustrates case *b*, which is consistent with Jensen's prediction. If most people belong to cases *a* and/or *d*, there will be little difference between the two. Finally, if most people belong to cases *e*, *f*, *g*, *h*, and/or *i*, Jensen's hypothesis will not hold and there will be more cheating under the linear than under the target-based scheme. Figure 2 illustrates case *h*, which is inconsistent with Jensen's prediction. Generally speaking, the Jensen cases involve relatively low amounts of guilt for small levels of cheating that increase rapidly at an increasing rate as the level of cheating rises. The opposite case, in which more cheating occurs under the linear scheme, can involve a relatively large amount of guilt for low levels of cheating and/or relatively slow increases in guilt as the level of cheating rises.⁷

Furthermore, the analysis above assumes that $G(c)$ is independent of the compensation scheme. However, this may not be the case. For example, it is possible that a scheme perceived as being unfair would mitigate guilt. For example, if a person works diligently and comes close to, but does not reach the target, s/he may react with anger at the perceived unfairness, and use this perceived unfairness as a justification to cheat. Clearly, the validity of the Jensen prediction is ultimately an empirical issue, and an empirical examination of this issue is the primary focus of this paper.

Hypothesis 2 (the Jensen Hypothesis): According to Jensen (2001, 2003), more cheating will occur under a target-based than under a linear piece-rate compensation scheme.

Comparing Cheating under the Target-Based and Tournament Settings

The target-based compensation scheme examined in this study is binary in nature, rewarding reported performance if and only if the prespecified target is reported to have been reached. An analogous binary tournament rewards reported performance if and only if it is at or above a prespecified percentile relative to other participants. Suppose for comparability that this percentile is chosen so that if performance is honestly reported approximately the same percentage of participants will obtain the bonus as in the target-based scheme. Then the two schemes differ in two major respects.

First, the target-based scheme presents participants with a certain production target. Once a participant learns how much s/he has produced in a given pay period, s/he knows whether or not s/he has actually achieved the target, and if not, how many over-claims s/he would have to make to receive the bonus nonetheless. In contrast, in the tournament, a participant does not know either whether s/he has produced enough to obtain the bonus without cheating, or how many

⁷ Note that these general descriptions are not exhaustive of the possibilities discussed more formally above.

over-claims if any would be necessary to obtain it. This is because s/he knows neither how much other participants will actually produce, nor how much they will cheat. Thus, while the payoff from a given amount of cheating is certain under the target-based scheme, it is uncertain under the tournament scheme. It is difficult to predict how this will affect individual cheating behavior since the expected returns to such behavior depend upon expectations about the uncertain behavior of others as well as individual attitudes toward risk.

Second, under the target-based scheme, cheating is at the expense of the firm's owners who must pay out more in bonuses. In a laboratory experiment, it is analogously at the expense of the experimenter's research grant. In contrast, in a tournament, cheating is at the expense of fellow participants since it causes bonuses to be given to one participant rather than to another. To the extent that participants feel more solidarity and less social distance between themselves and the other participants than between themselves and the owners of their firm or between themselves and the experimenter, they may be more willing to cheat under a target-based scheme than in a tournament. This latter point gives rise to a somewhat speculative third hypothesis.

Hypothesis 3: There will be more cheating under the target-based compensation scheme than under the tournament compensation scheme.

2. Experimental Design

Participants were recruited at a medium-sized Canadian university by means of email solicitation through the Bachelor of Commerce program listserv. Potential participants were told that they were participating in a study about workplace issues, and that they would be paid. They were not given any other details prior to the experiment. All 158 participants were undergraduates and majors in economics or other business subjects. There were 72 males and 86 females with an average age of 20.25 years and a standard deviation of 1.73 years. A widely used anagram word-creation game (e.g., Cadsby *et al.*, 2007; Schweitzer *et al.*, 2004; Vance and Colella, 1990) was employed as the experimental task.

Upon arrival, participants were randomly divided into two groups of equal size. One group was instructed to stay in the room while the other group was taken to an adjacent room. The two groups followed exactly the same experimental instructions and went through the session in precisely the same way. The experimental instructions were read to the participants while they followed along on their own copies. The instructions informed the participants that they would begin by playing two practice rounds and subsequently play seven experimental rounds followed by a final trial round. The practice rounds were designed to familiarize participants with the experimental procedure, task and incentives. The purpose of the final trial round will be discussed below. Participants were paid based only on the data from the experimental rounds, and only

these data are used in our analysis.

The instructions informed the participants that they would be provided with a prepared workbook containing the anagrams and a separate performance record sheet on which to register their performance. Each anagram would be presented on a separate page of the workbook. Participants were told that they would be given exactly one minute to make as many words as possible from each anagram. A bell was used to signal the beginning and end of each round. There was a pause between the practice rounds and the experimental rounds for questions. Otherwise, each time the bell rang, participants would turn the page and move onto the next round. Participants were not permitted to look ahead to future pages or to go back to previous pages. Adherence to these instructions was carefully monitored by our team of research assistants. To ensure anonymity, participants wrote their assigned participant numbers, but not their names, on the covers of their workbooks.

The instructions also described the procedures that would be followed once all the rounds were complete. First, participants would be asked to fill out a short demographic questionnaire. The purpose of the questionnaire was to provide the experimenters with information regarding the gender, age, and first language of each participant. Second, participants would be asked to submit their workbooks to the experimenter, who would in turn hand them to another participant sitting in an adjacent room. Each participant would then be asked to check the work of another participant, whose workbook was given to him/her at random. To accomplish this task, each participant would be given a list of all known correct words and a Scrabble dictionary in case any correct words were missing from the list. The specific instructions were as follows: “At the end of the session, your performance will be checked by another person playing the same game. At the same time, you will be asked to check the words created by another participant. That participant will **not** be the same person checking your words.” Participants knew that a \$3 word-checking fee would be paid to compensate them for this cross-checking task.

Third, participants were informed that after the cross-checking was completed, each workbook would be returned to the participant to whom it belonged according to the participant number on the cover page. At that point, participants would be permitted to check over their own workbook to make sure that all of the correct words had been marked “correct” and that all of the incorrect words had been marked “incorrect.” They were given permission to correct any mistakes made by the rater. This system of cross-checking enabled us to distinguish between passive misrepresentation, in which a participant simply acceded to an error made by the rater, and active misrepresentation, in which a participant overruled the rater to make an incorrect claim. Fourth, participants would be asked to record the number of correct words they had created on the

performance record sheet in the space provided next to the corresponding round number. Fifth, participants were instructed that they should tear off the cover pages of their workbooks, deposit the cover pages and the workbooks into two separate sealed boxes, and bring their performance record sheets to another room one by one to get paid according to the performance record and compensation scheme plus the \$3 word-checking fee in cash.

Unbeknownst to the participants, the final trial round contained a unique set of seven letters for each participant that we used to match participants' workbooks with their self-reported performance records. Thus, we could compare actual productivity to claimed productivity for each participant by round after the completion of the experiment.⁸ However, since neither the performance records nor the coverless workbooks contained player names or id numbers, we could not match either to particular individuals.

The specific compensation scheme was the treatment variable of the study. The three experimental treatments were as follows. 1) *Linear Piece-rate*: Each participant was paid \$3 for word checking and \$0.40 for each word they created in the game. 2) *Target-Based Bonus*: Each participant was paid \$3 for word checking and \$3.60 for each of the seven experimental rounds in which s/he created nine words or more. Nine words was the 85th percentile performance level in our pre-test data, which were gathered in a piece-rate setting. 3) *Tournament-Based Bonus*: Each participant was paid \$3 for word checking and \$3.60 for each of the seven experimental rounds in which his/her performance was at or above the 85th percentile, relative to other participants in the same session.⁹ One session was devoted to each treatment with between 50 and 56 participants per session.

3. Results

We first examined whether demographic variables (age, gender, and first language) had any influence on our key dependent variables representing productivity and cheating behavior. We found no significant relationships between any of the demographic variables and the dependent variables. This is consistent with the finding reported in Schweitzer et al. (2004), who used a similar population and the same experimental task. As a result, we pooled the data from different demographic groups for the remainder of the analysis.

We then examined productivity under the three payment schemes. Productivity is defined as the number of correct words a participant created during the seven experimental rounds. These numbers are reported in the top row of Table 2. Productivity was similar under the

⁸ We followed Schweitzer et al. (2004) in this regard. However, in contrast to Schweitzer et al. (2004), we made no specific claim in the instructions that we could not match up the self-reported performance record to participants' workbooks after the experiment.

⁹ Please contact the authors for the complete instructions used in each of the three experimental treatments.

three PFP payment schemes, and the null hypothesis that it was identical could not be rejected at conventional levels of significance, using t-tests of differences in means.¹⁰ Thus, the remainder of our analysis focuses on over-reporting of productivity, i.e. cheating.

H1 predicts that under the target-based treatment, the probability of cheating behavior will be higher, the closer a person's actual output is to the prespecified target, i.e. the lower is c^* . We employed the following logit regression to test this hypothesis:

$$\ln [f/(1 - f)] = \beta_0 + \beta_1 \cdot c^* \quad (1)$$

where f is the probability of cheating for financial gain and c^* is the difference between the target and the actual number of words created. We used a random effects maximum-likelihood estimation procedure to deal with the lack of independence in the data due to the repeated measures from each participant over multiple rounds.¹¹ Since this hypothesis applies only to target-based compensation, we used only the data from the target-based treatment. We ran the regression twice, using two different definitions of cheating for financial gain. Under the first definition, a person was classified as having cheated for financial gain in a specific round if his/her actual productivity was less than the target level of nine words, but his/her reported productivity was greater than or equal to nine words. In this case, we utilized only the data from rounds in which a participant created fewer than nine words and hence had the opportunity to cheat for financial gain. The null hypothesis that $\beta_1 = 0$ was contrasted with the alternative specified by Hypothesis 1 that a smaller distance to the target, c^* , would be associated with a higher probability of cheating. The results of the logit regression reject the null hypothesis in the direction of the specified alternative ($\beta_1 = -1.486$, two-tailed $p = 0.01$).

Recall that each participant's work was graded by another participant. We refer to this person as the rater. In some instances, the rater marked incorrect words as correct, and sometimes these rater errors, if left in place by the original participant, resulted in cheating for financial gain under the first definition above. However, these instances of relatively *passive* cheating were possibly due to carelessness, and even if intentional, may have produced less guilt than the more active and clearly intentional cheating that took place when a rater's correct assessment was modified by the participant, resulting in financial gain. Thus, we also employed a second definition in which a person was classified as having *actively* cheated for financial gain in a specific round if his/her actual productivity was less than the target level of nine words, the rater correctly indicated that he had created fewer than nine words, and yet his/her reported

¹⁰ Non-parametric Mann-Whitney tests yield inferences that are virtually identical to those from the t-tests reported throughout this study. Therefore, to conserve space, we report only t-tests throughout the paper.

¹¹ We used the SAS NLMIXED Procedure and employed the Dual Quasi-Newton optimization technique.

productivity was greater than or equal to nine words. In this case, we used only data from rounds in which a participant both created fewer than nine words and was graded as having produced fewer than nine words by the rater, and thus had the opportunity to cheat actively. Under this definition, the null hypothesis is also rejected in the direction of H1 ($\beta_1 = -1.900$, two-tailed $p = 0.005$). These results are consistent both with our model of cheating behavior and with the evidence provided by Schweitzer et al. (2004).

H2 predicts more cheating under the target-based scheme than under the piece-rate scheme, while H3 predicts more cheating under the target-based scheme than under the tournament scheme. We compare cheating under all three schemes using several different measures of cheating. Table 2 reports measures that are based on numbers of over-claimed words over the seven experimental rounds. The first of these measures is the *gross number of over-claimed words*. This is defined as the sum over seven rounds of the per-round differences between the number of words reported as being correct and the number of correct words actually created by a participant, whenever such a difference is positive, and zero otherwise. The *gross number of over-claimed words* was more than two and a half times as high in the target-based treatment than in either of the other two treatments. These differences are both significant, supporting both H2 and H3.

It is possible that over-claims are the result of carelessness rather than intentional cheating. In that case, as pointed out by Schweitzer et al. (2004), one might expect there to be as many under-claimed as over-claimed words. In fact, there were far fewer rounds in which the number of reported words was less than, as opposed to more than, the actual number of correct words created. These rounds were ignored by our gross measure. If under-claims are entirely the result of carelessness, we might assume that a similar number of over-claims are also the result of carelessness. We can adjust for such carelessness by subtracting the number of under-claims in the under-claim rounds from the number of over-claims in the over-claim rounds. The *net number of over-claimed words* is thus the sum over seven rounds of the per-round differences between the number of words reported as being correct and the number of correct words actually created by a participant in each round, regardless of whether such differences are positive or negative. This reduces slightly the apparent cheating from each of the three treatments. Nonetheless, net over-claims are still significantly higher in the target-based treatment than in either of the other treatments.

Both of the measures reported above lump together instances in which over- or under-claims involve going along with the rater's assessment with instances in which a participant actively over-rules a rater's assessment. The former instances are more likely to be due to lack of

vigilance. Even if the participant realizes the rater made an error, it is possible s/he feels less guilt when deferring to another's erroneous judgment than when actively overruling another's accurate judgment to his/her own benefit. The third measure recognizes this distinction by considering only active cheating, in which a participant overrules the rater's correct assessment of his/her actual output with a higher self-report. The *gross number of active over-claimed words* is defined as the sum over seven rounds of the per-round differences between the number of words reported as being correct and the number of correct words actually created by a participant or incorrectly graded as correct by the rater, whenever this difference is positive, and zero otherwise. Once again, both Hypotheses 2 and 3 are strongly supported by the data, which show significantly more active over-claims in the target-based than in the other two treatments. This is also true for the fourth measure, the *net number of active over-claimed words*, which is the sum over seven rounds of the per-round differences between the number of words reported as being correct and the number of correct words actually created by a participant or incorrectly graded as correct by the rater, regardless of whether such differences are positive or negative.

Under the piece-rate scheme, an over-claim always has financial implications since pay increases directly with the number of reported words. Under the tournament scheme, over-claims do not always lead to increases in actual pay. However, they do always lead to a higher *ex-ante* probability of being among the tournament winners and thus receiving a higher payoff. In contrast, an over-claim under the target-based scheme only has financial implications insofar as it results in the achievement of the predetermined target. Suppose the target is nine words. A participant makes seven words. S/he reports nine words. The two over-claimed words result in increased pay. However, suppose this participant reports eight words. S/he overclaims by one word, but there are no financial implications. Similarly, if this participant reports 12 words, two of the over-claimed words have financial implications, but the other three do not. To our surprise, although the majority of over-claims in the target-based case did have financial implications, a considerable number did not. In particular, out of an average of 4.07 active overclaim words per participant in all seven rounds of the target-based treatment, 3.04 had financial implications. The other 1.03 words, 25.3% of the total, did not.

In retrospect, there could be a number of reasons for such non-financial over-claims. First, the participant who makes seven words and decides to cheat in order to get the bonus might be reluctant to report precisely nine words, thinking this might provoke suspicion. S/he may feel it would look less suspicious to report a number of words somewhat in excess of the target. Second, a participant may strongly believe that a word is correct, so much so that s/he feels no need to check it against the word list or in the dictionary. Nonetheless, the word may not be in the

dictionary, or may violate the rules of the game by, for example, being a proper noun. The participant erroneously counts it as a correct word, sometimes even in the face of a contrary decision by the rater regardless of whether or not there are financial consequences. Third, it is possible that a participant erroneously uses two forms of the same word, e.g. a singular and a plural form. Although this is clearly not permitted by the instructions, the participant may not have realized this, and thus counted both as correct words for reasons that have nothing to do with financial implications. The second and third reasons could also motivate over-claims in the piece-rate and tournament treatments; however, in those treatments they do not stand out as being different from any of the other over-claims since all over-claims may have financial implications.

Should over-claims from the target-based treatment that have no financial implications be counted as cheating when comparing cheating under the three treatments? On balance, we think so. Although the first motivation applies only to the target-based case, the other two apply to all three cases. Thus, similarly motivated over-claims may occur under all three treatments. It is impossible to remove them from the piece-rate and tournament data. Therefore, in our view, it is best to include them in the target-based data as we have done above. However, it is nonetheless prudent to undertake an additional comparison that removes such non-financial over-claims to determine whether doing so affects the support we have obtained for Hypotheses 2 and 3. To so do, we need to take account of the fact that financial over-claims are only possible in a given round of the target-based treatment if a participant has created fewer than the nine words specified as the target. There are seven rounds in the experiment. However, if a participant makes nine or more words in three of those rounds, he only has the opportunity to make a financial over-claim during the other four rounds.

We thus define the *gross number of active financial over-claimed words per eligible round* for the target-based treatment as follows: First, we calculate the number of active over-claim words that have financial implications for each participant. We then divide this number by the number of rounds in which that participant (i) made fewer than the target of nine words and (ii) was correctly deemed to have made fewer than nine words by the rater, i.e., the rounds in which the participant had the opportunity to actively cheat. The result is the gross number of active financial over-claimed words per eligible round for that particular participant. Table 2 reports the average of these numbers over all participants as 0.50 words per eligible round. For the piece-rate and tournament treatments, all active over-claims have financial implications, and all rounds are eligible rounds in which it is possible to cheat. Therefore, active financial over-claimed words are simply calculated as the gross number of active over-claimed words divided by the total number of seven rounds. We find some support for Hypotheses 2 and 3 with p - values that are on the

border of significance at $p = 0.057$ and $p = 0.056$ respectively. Since incidents of active financial under-claims were negligible in all three treatments, with just one participant making such under-claims in just one round in each case, we do not repeat the analysis for net active financial under-claims.

So far, we have examined cheating only in terms of the number of over-claimed words under each payment scheme over the course of the experiment. This is comparable to comparing over-claims on units of output or sales over several pay periods. The experiment consisted of seven rounds or work/pay periods. Another way of comparing cheating is to consider the number of rounds in which each participant cheated by making at least one over-claim under each scheme. Table 3 reports the relevant statistics using definitions for over-claim rounds that are analogous to the definitions used for over-claim words. Regardless of the definition employed, the number of rounds per person in which over-claims occurred is more than twice as high under the target-based as under the linear piece-rate or tournament compensation schemes. Furthermore, the t-tests all reject the null hypotheses in favor of Hypotheses 2 and 3. Thus, Hypotheses 2 and 3 both continue to receive strong support when cheating is compared based on the number of rounds in which it occurs rather than on the number of over-claims made over the entire session.

A third perspective can be gained by comparing the probability of a person cheating at least once between treatments. Table 4 reports the number and percentage of people who cheated under each payment scheme for the different definitions of cheating. For the target-based scheme, these numbers are between 41% and 54%, while for the piece-rate scheme, they are between 25% and 33% and for the tournament scheme, between 26% and 32%. We ran the following logit regression separately for each of the various definitions of cheating and for each pair of treatments:

$$\ln [q/(1 - q)] = \beta_0 + \beta_1 \cdot T \quad (2)$$

where q is the probability of a person cheating at least once and T is a dummy variable representing the treatment or compensation scheme. The results are reported in Table 4. They again provide support for both Hypotheses 2 and 3, significantly rejecting the null in the direction of these hypotheses for all but one of the cheating definitions. The exception is gross active financial over-claims for which the p -values were marginal, 0.079 for H2 and 0.104 for H3. Since gross active financial over-claims and gross active over-claims are identical by definition under the piece-rate and tournament schemes, these marginal p -values are due to the fact that three fewer people made gross active *financial* over-claims than made gross active over-claims under the target-based scheme. It should be remembered that in the target-based treatment, there were fewer opportunities to make active financial over-claims than in the other treatments since such

cheating was only possible in rounds in which a participant made fewer than nine words.

Although we did not put forward any hypothesis concerning the relative frequency of cheating between the piece-rate and tournament schemes, we nonetheless report comparisons between these two treatments for each definition of cheating in each of the relevant tables. In no case was there a significant difference.

Despite the strong support for the notion that the target-based scheme provides the greatest temptation to cheat for most participants, this is not necessarily the case for everyone. Since each participant was randomly assigned to just one treatment, we are unable to compare directly the cheating behavior of particular individuals under the different schemes. However, it is interesting to note that the highest number of over-claims by one person, 38, occurred under the piece-rate scheme. Since 79 over-claims were made by all 52 participants in the piece-rate treatment, that one person was responsible for fully 48% of the piece-rate cheating. Under the tournament scheme, the biggest cheater was responsible for 38% of the over-claims, while under the target-based scheme, the comparable percentage was just 13%. Thus, we must note that while the target-based scheme produced far more cheating overall than the other two schemes, some people were nonetheless willing to cheat very substantially under the other schemes.

4. Conclusions and Implications

Warren Buffett once remarked, “Managers that always promise to ‘make the numbers’ will at some point be tempted to *make up* the numbers” (Buffett, 2003). Our experiment provides strong support for this insight. Whether one considers the number of over-claimed words, the number of work/pay periods in which such over-claims occur, or the number of participants who make an over-claim at least once, financially-salient numerical targets produce more cheating than other compensation systems. In particular, as argued by Michael Jensen (2001, 2003), a linear piece-rate system produces significantly less cheating than one based on a prespecified target. Moreover, a tournament compensation system based on relative performance also results in significantly less cheating than a target-based scheme. In addition, as first demonstrated by Schweitzer *et al.* (2004), cheating is more likely under a target-based scheme the closer a participant is to the target.

In our experimental design, participants were instructed to report accurately the number of words that they had created. Furthermore, before making a decision about how much to report, each participant received anonymous feedback from another participant selected at random. Thus, in contrast to many experiments on tax evasion, participants were placed in a situation where cheating was really cheating rather than the acceptance of a perceived invitation to gamble (see Cadsby *et al.*, 2006 for a discussion of this issue). In addition, we were able to separate instances

of active cheating, in which a participant overruled the verdict of an anonymous rater, from instances where a participant simply went along with a beneficial error made by a rater. Both these active cheating data and the total cheating data strongly support the notion that a prespecified numerical target linked to a financial bonus leads to more cheating than the other compensation schemes.

In our design, there was no chance of being caught. The possibility of an audit was never mentioned throughout the experiment, and no such audit ever occurred. Of course, participants may have imagined that if they were to make outlandish claims, they might be challenged. Indeed if any participant under the piece-rate scheme had claimed to have made more words than our budget could tolerate, we would have been forced to depart from the experimental protocol and check his/her claim.¹² As mentioned above, one participant made 38 over-claims in the piece-rate treatment. The belief that a ridiculously high claim might be challenged may have kept him/her from going further.

We purposely avoided an explicit audit system in order to examine the role played by the interplay of the utility of the financial rewards that may accrue from cheating with the psychological costs that may arise from such dishonest anti-social actions. If an audit and punishment system were in place, the costs of this system both to the employer and the employee would have to be considered in addition to psychological costs. There is some evidence that the existence of such an audit system might replace in whole or in part the costs of guilt (Gneezy and Rustichini, 2000). The design of an audit system must be considered in conjunction with a particular compensation scheme. For example, an audit system that responds to large self-reports with a higher probability of audit, would likely constrain cheating under a piece-rate or tournament system much more effectively than under a target-based system, in which there is no reason to report selling much more than the prespecified target. Under a target-based system, a random audit of those claiming to have achieved the target would likely be more effective. This is a subject that warrants further study.

Michael Jensen's passionate appeal for an end to bonuses based on prespecified numerical targets and their replacement with linear compensation systems cannot be supported by economic theory alone. As we have demonstrated, economic theory is consistent with people cheating either more or less under a target-based than under a linear PFP system. However, Jensen's argument receives strong and unambiguous support from our experimental data. While it is important to take account of other characteristics of compensation systems in addition to their

¹² We budgeted for the possibility of large amounts of cheating. However, if somebody had brazenly claimed to have made a million words, we would have been unable to pay.

effects on cheating behavior, especially potential effects on effort and output, knowing that the choice of compensation scheme can have an important effect on cheating is clearly an important factor in the design of such systems. Jensen's appeal deserves serious attention.

References:

- Alm, J., M. Cronshaw, and M. McKee. 1993a. Tax Compliance with Endogenous Audit Selection Rules. *Kyklos*, 46, 27-45.
- Alm, J., B. Jackson, and M. McKee. 1992a. Institutional Uncertainty and Taxpayer Compliance. *American Economic Review*, 82, 1018-1026.
- Alm, J., B. Jackson, and M. McKee. 1992b. Estimating the Determinants of Taxpayer Compliance with Experimental Data. *National Tax Journal*, 45, 107-114.
- Alm, J., B. Jackson, and M. McKee. 1993b. Fiscal Exchange, Collective Decision Institutions and Tax Compliance. *Journal of Economic Behavior and Organization*, 22, 285-303.
- Alm, J., G. H. McClelland, and W. D. Schulze. 1992c. Why Do People Pay Taxes? *Journal of Public Economics*, 48, 21-38.
- Alm, J., G. H. McClelland, and W. D. Schulze. 1999. Changing the Social Norm of Tax Compliance by Voting. *Kyklos*, 52, 141-171.
- Alm, J. and M. McKee. 1998. Extending the Lessons of Laboratory Experiments on Tax Compliance to Managerial and Decision Economics, *Managerial and Decision Economics*, 19, 259-275.
- Alm, J., M. McKee, and W. Beck. 1990. Amazing Grace: Tax Amnesties and Compliance. *National Tax Journal*, 43, 23-37.
- Alm, J., I. Sanchez, and A. de Juan. 1995. Economic and Noneconomic Factors in Tax Compliance. *Kyklos*, 48, 3-18.
- Ambrose, M., and M. Schminke. 1999. Sex Differences in Business Ethics: The Importance of Perceptions. *Journal of Managerial Issues*, 11, 454-475.
- Barsky, A. (2008). Understanding the ethical cost of organizational goal-setting: A review and theory development. *Journal of Business Ethics*, 81(1), 63-81.
- Becker, G. S. 1968. Crime and Punishment: An Economic Approach. *Journal of Political Economy*, 76, 169-217.
- Beron, K. J., H. V. Tauchen, and A. D. Witte. 1992. The Effect of Audits and Socioeconomic Variables on Compliance. In Slemrod, J. (ed.), *Why People Pay Taxes: Tax Compliance and Enforcement*. Ann Arbor: The University of Michigan Press, pp. 67-89.
- Boylan, S. J., and G. B. Sprinkle. 2001. Experimental Evidence on the Relation between Tax Rates and Compliance: The Effect of Earned Versus Endowed Income. *Journal of the American Taxation Association*, 23, 75-90.
- Buffett, W. E. 2003. *Letter to the Shareholders of Berkshire Hathaway Inc. for 2002*, <http://www.berkshirehathaway.com/letters/2002pdf.pdf>.

- Cadsby, C. B., E. Maynes, and V.U. Trivedi. 2006. Tax Compliance and Obedience to Authority at Home and in the Lab: A New Experimental Approach. *Experimental Economics*, 9, 343-359.
- Cadsby, C. B., F. Song and F. Tapon. 2007. Sorting and Incentive Effects of Pay-for-Performance: An Experimental Investigation. *Academy of Management Journal*, 50, 387-405.
- Chen, A. Y., R. B. Sayers, and P. F. Williams. 1997. Reinforcing Ethical Decision Making Through Corporate Culture. *Journal of Business Ethics*, 16, 855-865.
- Covey, M. K., S. Saladin, and P. J. Killen. 1989. Self-monitoring, Surveillance, and Incentive Effects on Cheating. *Journal of Social Psychology*, 129, 673-680.
- Degeorge, F., J., Patel, and R. Zeckhauser. 1999. Earnings Management to Thresholds. *Journal of Business*, 72, 1-33.
- Delaney, J. T., and D. Sockell. 1992. Do Company Ethics Training Programs Make a Difference? An Empirical Analysis. *Journal of Business Ethics*, 11, 719-727.
- Donaldson, T., and T. Dunfee. 1994. Toward a Unified Conception of Business Ethics: Integrative Social Contracts Theory. *Academy of Management Review*, 19, 252-284.
- Ehrlich, I. 1973. Participation in Illegitimate Activities: A Theoretical and Empirical Investigation. *Journal of Political Economy*, 81, 521-565.
- Feld, L. P., and J. R. Tyran. 2002. Tax Evasion and Voting: An Experimental Analysis. *Kyklos*, 55, 197-221.
- Ford, R. C., and D., Richardson. 1994. Ethical Decision Making: A Review of the Empirical Literature. *Journal of Business Ethics*, 13, 205-221.
- Gino, F., Ayal, S., and Ariely, D. (2009). Contagion and differentiation in unethical behavior: The effect of one bad apple on the barrel. *Psychological Science*, 20(3), 393-398.
- Gino, F., and Bazerman, M. H. (2009). When misconduct goes unnoticed: The acceptability of gradual erosion in others' unethical behavior. *Journal of Experimental Social Psychology*, forthcoming.
- Gino, F., and Pierce, A. (2009a). The abundance effect: Unethical behaviour in the presence of wealth. *Organizational Behavior and Human Decision Processes*, forthcoming.
- Gino, F., and Pierce, L. (2009b). Dishonesty in the name of equity. *Psychological Science*, Forthcoming.
- Grover, S., and Hui, C. (2005). How job pressures and extrinsic rewards affect lying behavior. *International Journal of Conflict Management*, 16, 424-437.
- Glover, S., M. A. Bumpus, G. F. Sharp, and G. A. Munchus, 2002. Gender Difference in Ethical Decision Making. *Women in Management Review*, 17, 217-228.

- Gneezy, U., and A. Rustichini, 2000. A Fine is a Price. *Journal of Legal Studies*, 29, 1-18.
- Holt, C., and S. Laury. 2002. Risk Aversion and Incentive Effects in Lottery Choices. *American Economic Review*, 92, 1644-1655.
- Jensen, M. C. 2001. Corporate budgeting is broken – Let’s fix it. *Harvard Business Review*, 79, 94-101.
- Jensen, M. C. 2003. Paying People to Lie: The Truth about the Budgeting Process. *European Financial Management*, 9, 379–406.
- Loe, T. W., L. Farrell, and P. Mansfield, 2000. A Review of Empirical Studies Assessing Ethical Decision Making in Business. *Journal of Business Ethics*, 25, 185-204.
- Moser, D. V., J. H. Evans III, and C. K. Kim. 1995. The Effects of Horizontal and Exchange Inequity on Tax Reporting Decisions. *The Accounting Review*, 70, 619-634.
- Nagin, D. S., J. Rebitzer, S. Sanders, and L. Taylor. 2002. Monitoring, Motivation and Management: The Determinants of Opportunistic Behavior in a Field Experiment. *American Economic Review*, 92, 850-872.
- Ordóñez, L., Schweitzer, E., Galinsky, A., and Bazerman, M. (2009a). Goals Gone Wild: The Systematic Side Effects of Over-Prescribing Goal Setting. *Academy of Management Perspectives*, 23, 9-16.
- Ordóñez, L., Schweitzer, E., Galinsky, A., and Bazerman, M. (2009b). On Good Scholarship, Goal Setting, and Scholars Gone Wild, *Academy of Management Perspectives*, in press.
- Rickman, R., and R. Witt. 2007. The Determinants of Employee Crime in the UK. *Economica*, 74, 161-175.
- Schminke, M., M. Ambrose, and T. Noel. 1997. The Effect of Ethical Frameworks on Perceptions of Organizational Justice. *Academy of Management Journal*, 40, 1190-1208.
- Schminke, M., D. Wells, J. Peyrefitte, and T. C. Sebor. 2002. Leadership and Ethics in Work Groups: A Longitudinal Assessment. *Group and Organization Management*, 27, 272-294.
- Schweitzer, M., L. Ordóñez, and B. Douma. 2004. Goal-setting as a Motivator of Unethical Behavior. *Academy of Management Journal*, 47, 422–433.
- Trevino, L., and S. Youngblood, 1990. Bad Apples in Bad Barrels: A Causal Analysis of Ethical Decision-Making Behavior. *Journal of Applied Psychology*, 75, 378-385.
- Vance, R., and A. Colella. 1990. Effects of Two Types of Feedback on Goal Acceptance. *Journal of Applied Psychology*, 75, 68–77.
- Weaver, G., L. Trevino, and P. Cochran. 2000. Integrated and Decoupled Corporate Social Performance: Management Commitments, External Pressures, and Corporate Ethics Practices. *Academy of Management Journal*, 42, 539-553.

Table 1 Comparing Cheating under the Target-based and Linear Piece-rate Settings

		Linear Piece-Rate				
		No	$\hat{c} < c^*$	$\hat{c} = c^*$	$\hat{c} > c^*$	$\hat{c} \rightarrow \infty$
Target-based	Yes	b	c	d	e	f
	No	a	g	Not possible	h	i

Table 2: Treatment Effects Based on Over-Claimed Words per Person (two-tailed p-values in parentheses)

Measures	Treatments			T-tests of Differences in Means		
	1. Piece-rate N=52	2. Target-based N=56	3. Tournament N=50	1 vs. 2	2 vs. 3	1 vs. 3
Productivity	41.56	43.78	42.38	-0.910 (0.365)	0.539 (0.591)	-0.298 (0.766)
Gross over-claimed words	1.52	4.34	1.58	-2.29 (0.024)	2.29 (0.024)	-0.06 (0.953)
Net over-claimed words	1.42	4.02	1.48	-2.15 (0.034)	2.10 (0.038)	-0.06 (0.955)
Gross active over-claimed words	1.29	4.07	1.38	-2.27 (0.026)	2.32 (0.022)	-0.09 (0.925)
Net active over-claimed words	1.25	3.89	1.36	-2.20 (0.030)	2.19 (0.030)	-0.12 (0.908)
Gross active financial over-claimed words by eligible rounds	0.18	0.50	0.20	-1.92 (0.057)	1.94 (0.056)	-0.09 (0.925)


Table 3: Treatment Effects on Over-Claim Rounds per Person (two-tailed p-values in parentheses)

Measures	Treatments			T-tests of Differences in Means		
	1. Piece-rate N=52	2. Target-based N=56	3. Tournament N=50	1 vs. 2	2 vs. 3	1 vs. 3
Gross over-claim rounds	0.62	1.61	0.64	-3.09 (0.003)	2.94 (0.004)	-0.11 (0.912)
Net over-claim rounds	0.54	1.31	0.54	-2.32 (0.023)	2.20 (0.031)	-0.01 (0.995)
Gross active over-claim rounds	0.44	1.39	0.50	-3.12 (0.003)	2.88 (0.005)	-0.30 (0.766)
Net active over-claim rounds	0.42	1.29	0.48	-2.96 (0.004)	2.65 (0.010)	-0.30 (0.627)
Gross active financial over-claim rounds by eligible rounds	0.06	0.17	0.07	-2.91 (0.005)	2.62 (0.010)	-0.30 (0.766)

Table 4: Treatment Effects on Number of People who Over-Claimed (two-tailed p-values in parentheses)


Measures	Treatments			Logit Regression Treatment Coefficients		
	1. Piece-rate N=52	2. Target-based N=56	3. Tournament N=50	1 - 2	2 - 3	1 - 3
Gross over-claim	17 (0.33)	30 (0.54)	16 (0.32)	0.87 (0.030)	-0.90 (0.027)	-0.02 (0.940)
Net over-claim	13 (0.25)	28 (0.50)	16 (0.32)	1.10 (0.008)	-1.05 (0.013)	0.03 (0.908)
Gross active over-claim	13 (0.25)	26 (0.46)	13 (0.26)	0.96 (0.022)	-0.90 (0.031)	0.03 (0.908)
Net active over-claim	13 (0.25)	26 (0.46)	13 (0.26)	0.96 (0.022)	-0.90 (0.031)	0.03 (0.908)
Gross active financial over-claim	13 (0.25)	23 (0.41)	13 (0.26)	0.74 (0.079)	-0.69 (0.104)	0.03 (0.908)

Figure 1



Note: $U[k \cdot (q_1 + c)] - U(k \cdot q_1) < G(c)$ for all c . The discontinuity of $G(c)$ at zero is represented by the solid and open circles at $c=0$.

Figure 2



Note: $U[k \cdot (q_1 + \hat{c})] - U(k \cdot q_1) > G(\hat{c})$. The discontinuity of $G(c)$ at zero is represented by the solid and open circles at $c=0$.