Discussion Paper Series

2009 – 12

Department of Economics
Royal Holloway College
University of London
Egham TW20 0EX

Signaling in Deterministic and Stochastic Settings*

Thomas D. Jeitschko[†]
Department of Economics
Michigan State University
East Lansing, MI 48824
U.S.A.
jeitschk@msu.edu

Hans-Theo Normann
Department of Economics
Goethe University Frankfurt
60323 Frankfurt
Germany
normann@econon.uni-frankfurt.de

October 28, 2009

Abstract

We contrast a standard deterministic signaling game with one where the signal-generating mechanism is stochastic. With stochastic signals a unique equilibrium emerges that involves separation and has intuitive comparative-static properties as the degree of signaling depends on the prior type distribution. With deterministic signals both pooling and separating configurations occur. Laboratory data support the theory: In the stochastic variant, there is more signaling behavior than with deterministic signals, and less frequent types distort their signals relatively more. Moreover, the degree of congruence between equilibrium and subject behavior is greater in stochastic settings compared to deterministic treatments.

Keywords: experiments, noise, signalling, learning, stochastic environments.

JEL classification: C7, C9, D8

---

*We would like to thank Tony Creane, John Kagel, audience members at Seminars at the University of York 2005, Royal Holloway College (University of London) 2006, University of Frankfurt 2008, and Max-Planck-Institute Jena 2009; as well as participants at the Midwest Economic Theory Meetings at Michigan State University 2006, the Midwest Mathematical Economics Meetings at Ohio State University 2008, and the Economic Science Association Meetings in Innsbruck 2009, for comments. We are grateful to the Max-Planck Institute for Research on Collective Goods (Bonn) for financial support.

[†]Phone: 517 355 8302, Fax: 517 432 1068

1 Introduction

Since the seminal work by Spence (1973, 1974), signalling—that is, the costly undertaking of actions in order to either convey or hide private information from others—has become the focus of much research within and beyond economics. In addition to the original work by Spence focusing on education, applications and variations have been seen in industrial organization (e.g., entry deterrence through limit pricing (Milgrom and Roberts, 1982) or signalling of product quality (Milgrom and Roberts, 1986; Bagwell and Riordan, 1991)), monetary policy (Backus and Driffill, 1985) and the economics of litigation (Reinganum and Wilde, 1984), to name just a few; and outside of economics insights from signalling have found applications in biology (e.g., Zahavi, 1975) and anthropology (e.g., Sosis and Ruffle, 2003).

Although these models greatly differ in their approaches and applications, there is one thing they have in common. In the vast majority of signalling games, the signal-generating mechanism is deterministic. The sender is able to perfectly control the signal, and the receiver precisely observes the signal that is sent. The receiver has no trouble interpreting the signal and can therefore correctly infer its cost. that is, there are no inaccuracies in sending or receiving the signal.

Even though standard, the assumption of deterministic signalling is not always plausible. For illustrative purposes, consider a Spence-type education-signaling game in which students signal their (unobservable) ability to potential employers through their choice of an (observable) level of education attained, and suppose that the signal that employers observe is a student's grade-point-average. Problems at the signal-generating stage may occur if, for instance, a student has a "bad day" (or a "good day," for that matter) during an exam. In this case, the sender is only imperfectly able to control the signal. Problems at the receiving end may occur if the employer cannot assess whether the classes taken by the student were particularly easy or hard. Similarly, education will also be a noisy signal for the receiver in a scenario where education is measured by the observable number of years of school attendance but where the education choice is affected by an intrinsic (dis)utility for education. As the receiver will be unaware of the utility of education, perfect inference of the cost of the signal is no longer possible. In these examples, the signal-generating process is, in effect, stochastic.

Matthews and Mirman (1983) were the first to study such a setting. They consider a variation of the limit-pricing game introduced by Milgrom and Roberts (1982). In particular, they suppose that the incumbent monopolist chooses an unobservable output level before stochastic demand for the product is realized. This results in an observable but stochastic price which only imperfectly reveals the underlying output choice and hence the type of the incumbent. Another study, Hertzendorf (1993), adds noise to the Milgrom and Roberts (1986) advertising model. Hertzendorf argues that the recipients of advertising signals will only rarely be informed about the exact advertising budget of a company. Instead, people receive a noisy signal of the budget when observing adverts.

Technically, what happens in stochastic signaling games is that any signal realization is consistent with any action taken by any type whenever the noise perturbing the signal has full support. Thus, signals are no longer invertible and therefore do not allow complete information about the underlying actions of the sender, even when agents of different types undertake different actions in equilibrium (i.e., a separating equilibrium). The observable signals only allow incomplete inferences about the sender's true (unobservable) type. In other words, Bayesian updating leads to incremental information dissemination when agents undertake distinct actions, rather than immediate and complete learning.

A main result in noisy signaling games is that often a unique separating equilibrium emerges (Matthews and Mirman, 1983), instead of the large number of possible equilibrium configurations that emerge without noise and that differ quantitatively and qualitatively in deterministic games. Carlsson and Dasgupta (1997) make use of the uniqueness result of the stochastic game to suggest an equilibrium-selection criterion for deterministic signalling games when noise vanishes—demonstrating the conditions for a unique noise-proof equilibrium to exist in deterministic games.

A second significant deviation of the equilibrium in noisy signalling games compared to deterministic versions is that the former admit a much richer comparative-statics analysis. In particular, in deterministic games actions are generally independent of the underlying distribution of types within a class of equilibrium configurations so that there are no meaningful comparative statics with respect to prior beliefs. In contrast, in a noisy signalling model the unique equilibrium is sensitive to variations in the underlying distributions, which yields smooth comparative statics with respect to prior beliefs.

In the present paper, we consider a generic sender-receiver signaling game in which we compare the deterministic and the stochastic signal-generating mechanism. In our setting, a sender chooses from a continuum of actions while the receiver only has two actions. We provide theoretical analyses and experimental data for this setup.

Our theoretical analysis shows that there are many perfect Bayesian equilibrium constellations in the variant without noise. After the application of equilibrium refinements, we obtain a unique perfect Bayesian equilibrium. Depending on the prior, this unique equilibrium is either pooling or separating. For the stochastic case, even without resorting to refinements, we obtain a unique equilibrium which is separating. Thus, a first hypothesis (testable in the experiments) is that for certain priors pooling behavior should occur without noise as opposed to separating behavior with noise. A second implication of the noisy signalling framework is that players should always signal, that is, always choose actions that differ from their myopically optimal actions. This is in contrast to the deterministic case in which one type always chooses the myopic best action and does not engage in signaling.[1] Often, the impact of noise on players' decisions is ambiguous and depends on the prior beliefs and players' types. This sometimes leads to intriguing comparative-statics predictions that can be tested experimentally. For example, it is the less frequent type who chooses a message that is more strongly distorted away from the sender's myopic best action.

We complement the theoretical analysis with experimental data. Experimental research on signalling games has proven useful is assessing the relevance of the theory. For early contributions see for example Miller and Plott (1985), Brandts and Holt (1992), Potters and van Winden (1996) and Cooper, Garvin and Kagel (1997a, 1997b). More recent studies include Cooper and Kagel (2005) and Kübler *et al.* (2008). The case for studying a noisy signalling game seems particularly strong given the various propositions that differ markedly from the deterministic environment.

[1] In a in a separating equilibrium the "low" type chooses his myopic first-best action, and in a (refined) pooling equilibrium it is the "high" type who chooses his myopic first-best action.

In our experiments, we study treatments with two different priors (a "high" and "low" prior belief for the sender's type). For each of the two priors, we implement a deterministic and a noisy variant. We find that, as in previous experiments, our sessions do not completely converge to equilibrium.[2] Nevertheless, our experimental results provide some clear confirmation of the theory. Regarding the key variables of our experiment, the theory has predictive power. In addition, the hypotheses mentioned in the previous paragraph are supported by the data. Thus, for the high prior, there is more pooling behavior in the deterministic variant; we find indeed that there is significantly more signalling with noise. While there is no support for the hypothesis that the less frequent type signals more, in relative terms, this prediction is confirmed. Overall, the empirical data are closer to their equilibrium counterparts in the stochastic variant compared to the deterministic setting. We attribute this to the fact that noise in the model is similar to imperfect play by subjects leading to a greater congruence between equilibrium observations and subject behavior.

Our study is among the first experiments to analyze a noisy signalling game. In independent research, de Haan, Offerman and Sloof (2008) also construct a model with noisy signals and run experiments. Their model differs from ours in that a pooling equilibrium may exist with noise, because the two sender types have the same first-best preferred action and the marginal cost of signaling is strictly positive. In addition to their altered model, their main focus also differs from ours in that they look at varying levels of noise, whereas we examine how prior beliefs affect play in noisy and deterministic games.

2 The Model

There are two players who act in sequence. The first player to move is referred to as the sender (of male gender), and the second player is referred to as the receiver (of female gender). Before play begins, nature draws the sender's type. With probability $\rho_0 \in (0, 1)$ the sender is the "high" type, denoted by $\bar{t}$. With complementary probability of $1 - \rho_0$ the sender is

[2]See Cooper and Kagel (2005) on this issue. They convincingly argue that previous experimental signalling games do not immediately converge. Without repetitions or other mechanisms facilitating learning, equilibrium play emerges only gradually, if at all. Cooper and Kagel (2005) show that teams play dramatically more strategically than individuals.

the "low" type denoted by $\underline{t}$ $(< \overline{t})$.

The sender observes his type and then chooses a hidden/unobservable action a that affects his payoffs both directly and indirectly. The indirect effect comes about because the unobservable action a generates a (possibly noisy) signal s that triggers a payoff-relevant reaction r by the other player, the receiver. Specifically, the sender's (type-dependent) payoff is given by

$$u(a,r) = \begin{cases} U - c\left(a - \overline{t}\right)^2 + W\, r(s), & \text{if } t = \overline{t} \\ U - c\left(a - \underline{t}\right)^2 + W\, r(s), & \text{if } t = \underline{t}, \end{cases} \tag{1}$$

where U is a normalization parameter, $r \in \{0, 1\}$ is the receiver's response (based on the observed signal s), $c > 0$ is a scaling parameter, and $W > 0$ is a windfall that the sender obtains when $r = 1$.

The agent's most preferred (myopic best) action is thus $a = t$, and deviations from this (i.e., $a \neq t$) entail signalling. Signalling may be undertaken in order to induce the receiver to take a response of $r = 1$, rather than a response of $r = 0$, as $r = 1$ results in the agent obtaining the added windfall payment of W. The sender's type-dependent payoff as a function of the action is depicted in Figure 1, where, in order to observe signaling behavior in the equilibrium of the deterministic setting, we have restricted parameters such that $\overline{t} - \underline{t} < 2\sqrt{W/c}$.

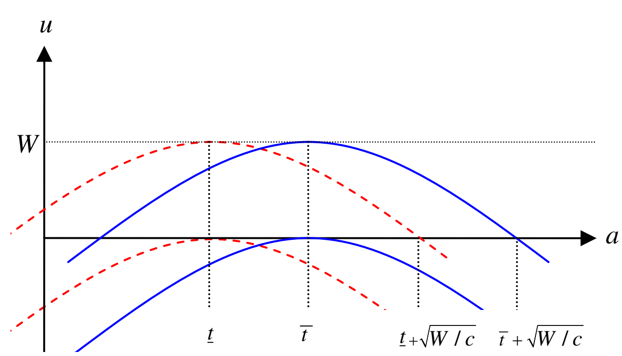

Figure 1: *Sender's Payoff* (with W (top) and without W (bottom)—solid lines: high type; dashed lines: low type)

The receiver does not know the type of the sender. Her prior beliefs are captured by ρ_0. These prior beliefs are updated to ρ_1 upon observing the signal s on the basis of the relationship between the sender's actions a and the resulting signal s, given beliefs about how a sender's type t determines his action.

The receiver's payoffs are affected by her response $r \in \{0, 1\}$ and are given by

$$v(r,t) = \begin{cases} V + B\,r, & \text{if} \quad t = \overline{t} \\ V + B\,(1-r), & \text{if} \quad t = \underline{t}, \end{cases} \tag{2}$$

where V is some base-utility and $B > 0$ is a bonus that increases the weight of the decision variable r on the inference that the receiver has drawn about the agent's type. Given posterior belief $\rho_1(s)$, the receiver chooses $r \in \{0, 1\}$ in order to maximize

$$Ev(r) = V + B\left[\rho_1 r + (1-\rho_1)(1-r)\right]. \tag{3}$$

Hence, the receiver responds with $r = 0$ whenever $\rho_1 \leq 1/2$, and chooses the response $r = 1$ otherwise.

We analyze the game by postulating a switching strategy. That is, the receiver's response is determined by a critical threshold value of the signal $\tilde{s}$ for which

$$r = \begin{cases} 1 & \text{if} \quad s \geq \tilde{s} \\ 0 & \text{if} \quad s < \tilde{s}. \end{cases} \tag{4}$$

Using a switching point makes sense if the receiver thinks facing the higher type is more likely as the signal increases. This indeed occurs in the unique equilibrium of the noisy game, when noise has the monotone-likelihood-ratio property. For the deterministic game, a switching strategy is a plausible assumption although the game can be studied without switching strategies, which changes the results only marginally (as discussed below).

It is worth mentioning at this point that the receiver's payoff is strictly increasing in properly identifying the sender's type. That is, the receiver obtains B whenever she correctly identifies the sender's type. This is in contrast to many signalling games in which one response gives a constant payoff independent of the sender's type (e.g., not hiring the worker in Spence's model leads to a reservation payoff that is independent of the true type of the job applicant). Our setup corresponds to a scenario where a manager has to assign a worker to specific tasks within the firm (one requiring higher skills and therefore yielding greater compensation), and where the worker's subsequent performance correctly reveals his or her type in either case and thus may serve as a basis for the payment of the manager.

In summary, the sequence of events is:

1. after nature chooses the sender's hidden type $t \in \{\bar{t}, \underline{t}\}$ with $\Pr\left(t = \bar{t}\right) = \rho_0$, the sender chooses an unobservable action, denoted by $\underline{a}$ for $\underline{t}$, and $\bar{a}$ for $\bar{t}$;

2. the unobservable action a generates the observable signal s, which the receiver uses to update her beliefs about the sender's type upon which she chooses a response $r \in \{0, 1\}$;

3. both players' payoffs are realized according to (1) and (2).

The relationship between the sender's action a and the observed signal s depends on whether the signal-generating technology is deterministic or noisy. We analyze the two distinct environments in turn, concentrating on perfect Bayesian equilibrium solutions. In the deterministic setting the set of solutions is further refined, whereas in the noisy setting the PBE is unique and therefore does not require further refinement arguments.

2.1 Equilibrium with Deterministic Signals

In deterministic models the signal allows a perfect inference about the actions that were taken (i.e., the signal-generating mechanism is invertible). Indeed, this is informationally equivalent to a setting in which the action itself is observable. Thus, we consider them to be identical,

$$s \equiv a. \tag{5}$$

Restricting attention to pure-strategy equilibrium configurations, the action taken by the sender is either type-dependent and distinct for the two types (a separating equilibrium), or is independent of his type (a pooling equilibrium).

Separating Configurations. In a separating equilibrium the receiver infers the sender's type: the low type faces the unfavorable response of $r = 0$, whereas the high type achieves the favorable response of $r = 1$. As $\underline{t}$ faces the unfavorable response he chooses his myopic first best action, i.e., $\underline{a}^* = \underline{t}$, since any other action yields a lower payoff when his type is revealed.

(See Figure 1 for illustration.) Incentive compatibility for the low type requires that he does not find it advantageous to trigger the favorable response of $r = 1$ by choosing the high type's equilibrium action. This implies that $\overline{a}^* \notin \left[\underline{t} - \sqrt{W/c},\ \underline{t} + \sqrt{W/c}\right]$. Moreover, individual rationality for the high type dictates that he does not prefer to accept the unfavorable response $r = 0$ over taking the equilibrium action in order to obtain the favorable response of $r = 1$. That is, $\overline{a}^* \in \left[\overline{t} - \sqrt{W/c},\ \overline{t} + \sqrt{W/c}\right]$. Taking these considerations together yields a continuum of separating equilibrium constellations $\underline{a}^* = \underline{t}$ and $\overline{a}^* \in \left[\underline{t} + \sqrt{W/c},\ \overline{t} + \sqrt{W/c}\right]$, with switching point $\tilde{s}^* = \overline{a}^*$.[3]

Using Cho and Kreps' (1987) intuitive criterion, any separating equilibrium with $\tilde{s}^* > \underline{t} + \sqrt{W/c}$ can be upset. Specifically, suppose $\tilde{s}^* = \overline{a}^* > \underline{t} + \sqrt{W/c}$ and consider an out-of-equilibrium action $a \in \left[\underline{t} + \sqrt{W/c}, \tilde{s}^*\right)$. Such an action is dominated for $\underline{t}$. Even if the receiver responded with $r = 1$ to such an action, $\underline{t}$ would still be strictly better off choosing $a = \underline{t}$ and getting the $r = 0$ response. Thus, the receiver should believe $\rho_1 = 1$ after such a deviation and then $\overline{t}$ can profitably deviate to $\overline{a}' = \underline{t} + \sqrt{W/c}$. This leaves a unique separating equilibrium, the least-cost separating equilibrium, with $\underline{a}^* = \underline{t}$, $\overline{a}^* = \underline{t} + \sqrt{W/c}$ and $\tilde{s}^* = \overline{a}^*$.

Pooling Configurations. Note first that when $\rho_0 < 1/2$ there cannot be a pooling equilibrium. If types pool their actions with $\rho_0 < 1/2$, the receiver chooses the unfavorable response of $r = 0$. But then both types are no worse off by choosing their myopic best actions, which precludes them taking the same action so that a pooling equilibrium does not exist. With $\rho_0 \geq 1/2$ and pooling, the receiver chooses the favorable response of $r = 1$, that is, both types get W in equilibrium. Because the receiver employs a switching point, there does not exist a pooling equilibrium with $\tilde{s}^* = \underline{a}^* = \overline{a}^* < \overline{t}$ as $\overline{t}$ could profitably deviate to the myopically optimal $a = \overline{t}$ and still trigger $r = 1$. There is a continuum of equilibrium pooling configurations with $\underline{a}^* = \overline{a}^* \in \left[\overline{t},\ \underline{t} + \sqrt{W/c}\right]$ (having assumed that $\overline{t} - \underline{t} < 2\sqrt{W/c}$).[4]

[3]Note that, because we assume a switching point, all actions $a > \tilde{s}^* = \overline{a}^*$ trigger $r = 1$. Without switching-point strategies, these signals may induce the response $r = 0$. However, neither type has an incentive to choose $a > \tilde{s}^*$, because this reduces the sender's payoff without affecting the receiver. Hence, the set of separating equilibrium actions is the same with and without a switching-point strategy.

[4]Analyzing the game without switching strategies, the set of equilibrium pooling configurations is larger, namely $\underline{a}^* = \overline{a}^* \in \left[\overline{t} - \sqrt{W/c},\ \underline{t} + \sqrt{W/c}\right]$, which includes actions below the low-type's most preferred action. However, these additional configurations do not pass the intuitive criterion. We also note already

The equilibrium pooling configurations are strictly Pareto-rankable with lower actions strictly preferred by both types of sender. For that reason, the intuitive criterion does not refine the set of equilibrium configurations. (Whenever $\underline{a}^* = \overline{a}^* > \overline{t}$, both types would be better off choosing an out-of-equilibrium action of $a^* = \overline{t}$ if this triggered $r = 1$, thus, the out-of-equilibrium action is not equilibrium dominated and hence configurations with $a^* > \overline{t}$ survive.) One can select the efficient pooling equilibrium based on the Pareto criterion only, and, even if one rejects the Pareto criterion as *ad hoc*, the application of Grossman and Perry's (1986) perfect sequential equilibrium or Mailath *et al.*'s (1993) undefeated equilibrium still yields a unique pooling equilibrium in which $\underline{a}^* = \overline{a}^* = \overline{t}$.[5]

Note finally that, when $\rho \geq 1/2$, the efficient pooling equilibrium with $\underline{a}^* = \overline{a}^* = \overline{t}$ also Pareto dominates the least-cost separating equilibrium from the sender's point of view. Specifically, in this pooling equilibrium, $\overline{t}$ has no incentive to separate himself by choosing some action $a > \overline{t}$ since $\overline{t}$ already gets the maximum payoff in the pooling equilibrium. Applying the same equilibrium selection arguments as in footnote 5 leaves the efficient pooling equilibrium as the unique equilibrium if $\rho_0 \geq 1/2$. If $p < 1/2$, the least-cost separating equilibrium does survive the application of these additional refinements.

We summarize the refined equilibrium constellation for the deterministic case in Proposition 1 and Table 1.

Proposition 1 (Equilibrium in Deterministic Settings) *If $\rho_0 < \frac{1}{2}$, the unique perfect Bayesian equilibrium surviving equilibrium refinements is the least-cost separating equilibrium with $\underline{a}^* = \underline{t}$ and $\overline{a}^* = \underline{t} + \sqrt{W/c}$. If $\rho_0 \geq \frac{1}{2}$, the unique perfect Bayesian equilibrium surviving equilibrium refinements is the pooling equilibrium with $a^* = \overline{t}$.*

---

at this point that subjects of either type in the experiment only rarely choose actions $a \in (\underline{t}, \overline{t})$ in the deterministic framework.

[5]Perfect sequential equilibrium (Grossman and Perry, 1986) requires that, for each out-of-equilibrium message, the receiver hypothesizes that the message was sent by some set of types of the sender and revises her beliefs conditional accordingly. If precisely the set of hypothesized players best responds by choosing the out-of-equilibrium message, the original equilibrium is upset. A similar rationale underlies Mailath *et al.*'s (1993) undefeated equilibrium except that the out-of-equilibrium message must be chosen with positive weight by some sender type in another perfect Bayesian equilibrium. To see the application of these refinements to our game, consider a pooling equilibrium with $\underline{a}^* = \overline{a}^* > \overline{t}$ (given $\rho_0 > 1/2$). Now assume an out-of-equilibrium action $a \in [\overline{t}, a^*)$. This is profitable for both $\underline{t}$ and $\overline{t}$ whenever the response is $r = 1$. While the equilibrium requires $\rho_1 = 0$ after the deviation, the refinements requires $\rho_1 = \rho_0 > 1/2$ since the deviation is profitable for both types provided $r = 1$ and both types have an incentive to deviate.

		Prior Beliefs	
		$\rho_0 < 1/2$	$\rho_0 \geq 1/2$
Equilibrium Type	Pooling	—	$\underline{a}^* = \overline{a}^* = \overline{t}$
	Separating	$\underline{a}^* = \underline{t}$ $\overline{a}^* = \underline{t} + \sqrt{W/c}$	—

Table 1: *Refined Equilibrium Configurations with Deterministic Signals*

The predictions in Proposition 1 are based on the application of equilibrium refinements. The literature on the relevance of refinements in experiments (starting with Brandts and Holt, 1992, 1993) has not been conclusive and has not always found support for refinements. Thus, *ex ante*, it appears demanding to consider the above theoretical results as benchmarks for an experiment. Note, however, that the implications of the refinements are rather modest. They merely give preference to the least-cost separating equilibrium over Pareto inferior separating equilibria, and they select the Pareto efficient pooling equilibrium. The experimental results will show that there is support for Proposition 1.

2.2 Equilibrium with Noisy Signals

Consider now the case where the signal that the receiver observes is not invertible and therefore does not reveal the agent's action perfectly. As indicated in the introduction such noise may result because the sender does not have perfect control over the signal, or it may be that the receiver cannot clearly observe the action. In either case, the receiver must use statistical inference in order to update her beliefs about the action taken, and thus learn about the agent's type.

After we formalize the signal-generating mechanism, we consider the receiver's best response conditioned on her conjectures about the actions taken by the two types of agent. In anticipation of this response the optimal actions of the agent is derived. The equilibrium is found by imposing beliefs of the receiver that are consistent with the actions taken by the sender.

The signal-generating mechanism is given by

$$s \equiv a + \epsilon, \tag{6}$$

where ϵ is an unobservable noise term that is distributed independently of a (i.e., homoskedastically) according to a normal distribution with zero-mean and standard deviation of σ, i.e., $\epsilon \sim N(0, \sigma^2), \forall a$. We assume that the noise term is realized only after the sender has taken the action a. Thus, the sender is unable to adjust his actions in light of the realization of noise and, hence,

$$s \sim N(a, \sigma^2).$$

We consider first the receiver's inference problem and best response. Let $\underline{a}^c$ and $\overline{a}^c$, with $\underline{a}^c \neq \overline{a}^c$, denote the receiver's conjectures about which (unobservable) type-dependent actions are taken. As before ρ_1 denotes updated (posterior) beliefs. That is, ρ_1 is the receiver's subjective probability-assessment that the sender is a $\overline{t}$-type sender, conditional upon having observed the signal s, given the conjectures $\overline{a}^c$ and $\underline{a}^c$.

Then, with $f(s|a) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(s-a)^2}{2\sigma^2}\right)$ denoting the normal density of the distribution of s with mean a, Bayes' Rule yields:

$$\begin{aligned} \rho_1(s|\underline{a}^c, \overline{a}^c) &= \frac{\rho_0 f(s|a = \overline{a}^c)}{\rho_0 f(s|a = \overline{a}^c) + (1-\rho_0) f(s|a = \underline{a}^c)} \\ &= \frac{LR_0}{LR_0 + \exp\left(\frac{(\overline{a}^c - \underline{a}^c)(\overline{a}^c + \underline{a}^c - 2s)}{2\sigma^2}\right)}, \end{aligned} \tag{7}$$

where LR_0 denotes the likelihood ratio of prior beliefs, $\rho_0/(1-\rho_0)$.

Combining (7) with (3) yields the receiver's best response.

Lemma 1 (Best Response) *Given the conjecture $\overline{a}^c$ and $\underline{a}^c$, the receiver's best response is determined by a critical threshold value of s denoted by $\tilde{s}^c$. That is, $r^* = 1$ if and only if $s \geq \tilde{s}^c$ with*

$$\tilde{s}^c = \frac{\overline{a}^c + \underline{a}^c}{2} - \ln(LR_0) \frac{\sigma^2}{\overline{a}^c - \underline{a}^c}.$$

Notice that $\tilde{s}^c$ has several intuitive properties. If $\rho_0 = 1/2$, then the critical threshold is simply the average of the actions of the two types of agent. As prior beliefs become strongly biased in favor of one or the other type of sender (i.e., $LR_0 \to 0, \infty$), only extreme signals will lead to updating sufficiently strong to revise prior beliefs to change a response. Similarly, as the sender chooses a similar action regardless of type (i.e., we approach a pooling equilibrium, so to speak, and $\overline{a} \approx \underline{a}$), again only extreme signals trigger a response by the receiver that differs from what prior beliefs indicate. Finally, the same holds true for increases in the variance of the noise σ^2 so that, for given beliefs about the senders' actions, a noisier environment leads to less updating.

Having characterized the receiver's learning and response, consider now the sender's optimal actions. Recall that the receiver's choice of r affects the sender's payoff (see (1)). Since the choice of r is governed by ρ_1, which is a function of the sender's action a (see (7)), it is clear that the sender accounts for how a affects r. Since r is increasing in s, given $\overline{a}^c$ and $\underline{a}^c$, both types of sender have an incentive to increase s. That is, both types would like to be identified as being a high type: the high type wants to set himself apart, and the low type wants to deceive.

Thus, given $\overline{a}^c$ and $\underline{a}^c$ (the sender's action affects only the signal s, but cannot affect the receiver's conjectures) and given Lemma 1, the sender's objective is to choose a in order to maximize

$$\begin{aligned} & U - c(a-t)^2 + W \Pr\left(s \geq \tilde{s}^c | a\right) = \\ & U - c(a-t)^2 + W \int_{\tilde{s}^c}^{\infty} \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(s-a)^2}{2\sigma^2}\right) ds, \end{aligned} \tag{8}$$

where $\tilde{s}^c$ is given in Lemma 1. The (type-dependent) first-order condition is given by

$$2c(a-t) + \frac{W}{\sigma\sqrt{2\pi}} \int_{\tilde{s}^c}^{\infty} \frac{s-a}{\sigma^2} \exp\left(-\frac{(s-a)^2}{2\sigma^2}\right) ds = 0. \tag{9}$$

Hence,

Lemma 2 (Best Action) *Given the receiver's response conditioned on her conjecture $\overline{a}^c$ and $\underline{a}^c$, the sender's best action is implied by*

$$2c(a-t) = \frac{W}{\sigma\sqrt{2\pi}} \exp\left(-\frac{\left(\tilde{s}^c - a\right)^2}{2\sigma^2}\right),$$

where $\tilde{s}^c$ is as before, given in Lemma 1.

Proof. Let $g(s,a) = -\frac{(s-a)^2}{2\sigma^2}$. Then $g_a = -g_s = \frac{s-a}{\sigma}$. Hence the term under the integral in the FOC (9) is $-g_s e^{g(s,a)}$ and therefore the integral itself is $e^{g(\tilde{s}^c,a)}$, since $\lim_{s\to\infty} e^{g(s,a)} = 0$.□

Notice that Lemma 2 implies that both types of agent engage in signaling (i.e., $a^* > t$ for both types), independent of the receiver's conjectures about the actions taken, provided $\overline{a}^c \neq \underline{a}^c$. This is a reflection of the fact that the marginal gain from signaling is positive, and hence the sender is willing to trade-off deviations of a from t in order to obtain the positive marginal signalling gains. Specifically, the marginal cost of signalling is zero at t, whereas the marginal gains are strictly positive. Thus, in the noisy environment, players always signal.

In equilibrium, the receiver is aware of the sender's desire to manipulate the flow of information. That is, she is aware that the high type will choose an action in the hopes of distinguishing himself from the low type and, similarly, that the low type will attempt to mimic the high type. As a consequence, she is aware of Lemma 2. This leads to consistent beliefs in which $\overline{a}^c = \overline{a}^*$ and $\underline{a}^c = \underline{a}^*$. Thus,

Proposition 2 (Equilibrium in Stochastic Settings) *The equilibrium actions, $\overline{a}^*$ and $\underline{a}^*$, are implied by the equations*

$$2c\left(\overline{a}^* - \overline{t}\right) = \frac{W}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2\sigma^2}\left(\frac{\Delta a^*}{2} + \ln(LR_0)\frac{\sigma^2}{\Delta a^*}\right)^2\right), \tag{10}$$

$$2c\left(\underline{a}^* - \underline{t}\right) = \frac{W}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2\sigma^2}\left(\frac{\Delta a^*}{2} - \ln(LR_0)\frac{\sigma^2}{\Delta a^*}\right)^2\right); \tag{11}$$

where $\Delta a^* := \overline{a}^* - \underline{a}^*$. The equilibrium response is given in Lemma 1, with the $\overline{a}^*$ and $\underline{a}^*$ replacing $\overline{a}^c$ and $\underline{a}^c$.

Notice that, if $\rho_0 = 1/2$, then both types will deviate from their myopic best actions by exactly the same amount, otherwise the relatively less likely type deviates (i.e., signals) more.

3 Experimental design and procedures

The experiments were framed as an interaction of a worker and a personnel manager. In the instructions (see appendix), subjects were informed about the game as described above.

The worker's decision was framed as an effort in a "test" which preceded the employment decision. We made it clear that no real effort had to be invested, and we explained how the effort level chosen affected the worker's profit. The payoff information regarding the effort choice was given in a table.

As in the model, personnel managers then received the (deterministic or stochastic) signal derived from workers' effort levels in the test. Next, they had to decide whether or not to employ the worker for some (not specified) task. The descriptions of the payoffs explained that the manager is paid B if he or she employs a worker "suitable for the task" or if he or she does not employ a worker who is not suitable for the task. The suitability of the workers was randomly determined by the computer individually and in every period. Workers were paid W only if they were employed, in addition to the payoff from the effort choices.

The parameters we used for the experiment were $U = 100$, $V = 0$, $W = B = 100$, $c = 1/2$, $\bar{t} = 50$ and $\underline{t} = 40$. Effort levels had to be chosen from the interval $[25, 65]$. These parameters yield the payoff table in the instructions (see appendix).

Our treatment variables are the prior belief and the noise parameter. Specifically, we compare games without noise to those with noise. In the sessions with noise, the noise was normally distributed with $\varepsilon \sim N(0,\ 5^2)$. Following Ashenfelter *et al.* (1992), subjects were not given the specific formal details of the normal distribution. Instead, they were given 100 "past realizations" of the noise term and were told that they should expect "similar distortions today" (see appendix). As for the prior belief, we use priors of $\rho_0 = 1/3$ and $\rho_0 = 2/3$. We use the labels *NoNoise.33*, *NoNoise.67*, *Noise.33* and *Noise.67* for the corresponding treatments. Table 2 summarizes the treatment design.

At the end of each period, subjects were given the following feedback: They were informed about the worker's type and the actual effort decision. In sessions with noise, they were also told the noisy signal that the personnel manager received. Further, they were reminded of the personnel manager's decision and were given the resulting payoffs of both players.

We decided to allow for many repetitions because learning is necessary in such complex situations. Our experiments had a length of 40 periods.

Subjects were randomly rematched in every period in order to create an environment as close as possible to a single-period interaction between subjects. In each session, 20 subjects

		Noise	
		$\sigma = 0$	$\sigma = 5$
Prior	$\rho_0 = 1/3$	*NoNoise.33*	*Noise.33*
	$\rho_0 = 2/3$	*NoNoise.67*	*Noise.67*

Table 2: *Treatments*

participated. The matching scheme was such that subjects interacted within a group of ten subjects. Thus each session consists of two entirely independent groups.

We applied role switching in this experiment. That is, participants acted both in the role of the worker and in the role of the personnel manager. Roles were switched every five periods, so all participants played either role four times for five periods. Role switching is often believed to enhance learning. Although this has not been tested experimentally, it seems intuitive that subjects better understand the decision problem of the other players and therefore the overall game if they play in both roles. Role switching also emphasizes the one-shot nature of the interaction and therefore strengthens the effects of the random-matching scheme. Many signaling experiments employ role switching; see Brandts and Holt (1992, 1993), Cooper *et al.* (1997a,b), Potters and van Winden (1996), and Kübler *et al.* (2008).

Experiments were computerized. We used *z-Tree*, developed by Fischbacher (2007). Sessions were conducted at *BonnEconLab*, the University of Bonn's Experimental Economics Laboratory. In total, 160 subjects participated in eight sessions. We have two sessions (40 participants) for each treatment.

Sessions lasted between 60 and 75 minutes, including the time for reading the instructions and paying the subjects. Earnings were denoted in "points." The exchange rate of one euro for 500 points was commonly known. Subjects also received a show-up fee of four euros. Average earnings were about 13 euros, including the show-up fee.

4 Hypotheses

Given the experimental parameters ($U = 100$, $V = 0$, $W = B = 100$, $c = 1/2$, $\bar{t} = 50$ and $\underline{t} = 40$), the equilibrium benchmarks against which we compare the experimental data are given in Table 3. Recall that $\bar{a}$ and $\underline{a}$ refer to $\bar{t}$'s and $\underline{t}$'s effort choices, respectively, and $\Delta a := \bar{a} - \underline{a}$ is the effort difference. Employment rates are denoted by $\bar{e}$ and $\underline{e}$, and the average employment rate is $e := \rho_0 \bar{e} + (1 - \rho_0)\underline{e}$. Finally, $\tilde{s}$ denotes the switching point, that is, the signal above which employers choose $r = 1$. Equilibrium values are indicated by an asterisk (*) throughout.

	NoNoise	*Noise*
$\rho_0 = 1/3$	$\underline{a}^* = 40.0$, $\bar{a}^* = 54.1$, $\Delta a^* = 14.1$ $\underline{e}^* = 0$, $\bar{e}^* = 1$, $e^* = 0.33$ $\tilde{s}^* = 54.1$	$\underline{a}^* = 42.5$, $\bar{a}^* = 55.0$, $\Delta a^* = 12.5$ $\underline{e}^* \approx 0.06$, $\bar{e}^* \approx 0.83$, $e^* \approx 0.32$ $\tilde{s}^* \approx 50.2$
$\rho_0 = 2/3$	$\underline{a}^* = 50.0$, $\bar{a}^* = 50.0$, $\Delta a^* = 0.0$ $\underline{e}^* = 1$, $\bar{e}^* = 1$, $e^* = 1$ $\tilde{s}^* = 50.0$	$\underline{a}^* = 48.0$, $\bar{a}^* = 54.0$, $\Delta a^* = 6.0$ $\underline{e}^* \approx 0.49$, $\bar{e}^* \approx 0.88$, $e^* \approx 0.75$ $\tilde{s}^* \approx 48.1$

Table 3: *Equilibrium Constellations Given the Experiment Parameters*

Table 3 reveals that the predictions about the effects of our treatments are not always unambiguous. Consider, for example, the impact of noise on $\underline{a}$. In *Noise.33*, $\underline{a}$ should be higher than in *NoNoise.33*, but it is exactly the other way round in *Noise.67* and *NoNoise.67*. For $\bar{a}$, the impact of noise is different again. Instead of unambiguous hypotheses, the impact of noise often depends on the prior (or *vice versa*) in theses cases. In what follows, we accordingly present hypotheses in the form of ordinal rankings of the relevant variable across all four treatments. We then use the non-parametric Jonckheere-Terpstra test,[6] testing the null hypothesis that all treatments come from the same distribution against the predicted ranking of treatments.

[6]The Jonckheere-Terpstra test is a non-parametric test for more than two independent samples, like the Kruskal-Wallis test. Unlike Kruskal-Wallis, Jonckheere-Terpstra tests for ordered differences between treatments and thus requires an ordinal ranking of the test variable. See, e.g., Hollander and Wolfe (1999).

We start with sender behavior. The first hypotheses are on the effort levels chosen by the low type $\underline{a}$, by the high type $\overline{a}$, and on the effort difference $\Delta a = \overline{a} - \underline{a}$. All of these can be obtained from Table 3.

Hypothesis 1 (Sender's Effort Choice) *Concerning the senders' type-dependent equilibrium effort choices, the following rankings hold:*

(a) for the low-type's action $\underline{a}^*$: NoNoise.33 < Noise.33 < Noise.67 < NoNoise.67;

(b) for the high-type's action $\overline{a}^*$: NoNoise.67 < Noise.67 < NoNoise.33 < Noise.33;

(c) for the difference in actions Δa^*: NoNoise.67 < Noise.67 < Noise.33 < NoNoise.33.

A general implication of noise is as follows.

Hypothesis 2 (Signalling with Noise) *In the noisy treatments, senders should always signal, that is, they should always choose $a > t$.*

For the treatments with noisy signalling, we have an intriguing hypothesis which we already noted following Proposition 2.

Hypothesis 3 (Signalling Distortions with Noise) *The sender whose type is less likely under the prior beliefs engages in more costly signalling efforts, i.e., $\underline{a} - \underline{t} < \overline{a} - \overline{t}$ in Noise.33, and $\underline{a} - \underline{t} > \overline{a} - \overline{t}$ in Noise.67.*

We now turn to the receiver's behavior. Table 3 contains the data for the type-dependent employment rates and also for the average employment rates per treatment.

Hypothesis 4 (Employment Rates) *Concerning the senders' type-dependent equilibrium employment probabilities, the following rankings hold:*

(a) for the low type's employment rate $\underline{e}^*$: NoNoise.33 < Noise.33 < Noise.67 < NoNoise.67;

(b) for the high type's employment rate $\overline{e}^*$: Noise.33 < Noise.67 < NoNoise.67 = NoNoise.33;

(c) for the average employment rate e^*: Noise.33 < NoNoise.33 < Noise.67 < NoNoise.67.

We finally turn to the receiver's behavior, in particular, the equilibrium choice of switching points.

Hypothesis 5 (Switching Points) *Regarding the receiver's switching points, $\tilde{s}^*$, the following ranking holds: Noise.67 $<$ NoNoise.67 $<$ Noise.33 $<$ NoNoise.33.*

5 Results of the Experiments

The results section is structured as follows. We begin with an analysis of the worker (sender) behavior. Then we move on to the managers (receivers), before analyzing workers and managers jointly to see how they respond to one another's actual behavior.

We usually employ non-parametric tests where we (conservatively) count one group of randomly matched participants as one observation. Whenever we depart from this, we indicate how we deal with the possible non-independence of observations. With the help of the non-parametric tests, we test directed hypotheses throughout and report one-sided p-values, accordingly.

The results reported here are based on the analysis of periods 21 to 40, that is, the second half of the experiment. This is to take learning effects into account. We will not report on learning effects in detail. However, noticeable time trends do not occur after period 15 and the variance of choices does not decrease much further in the second half of the experiment. Hence, we restrict the data analysis to periods 21 to 40.

5.1 Worker (sender) behavior

Table 4 and Figure 2 summarize the effort choices across the four treatments. Table 4 reports average effort choices and their standard deviation. It also states the equilibrium benchmarks. Figure 2 displays the CDFs of choices by types and treatment.

As in many previous signaling games, choices do not perfectly settle on the equilibrium benchmarks, as can be seen in Table 4. Figure 2 also indicates that there is no complete separating or pooling behavior in any treatment. Frequently, workers choose their myopic best action. While this is consistent with equilibrium behavior for one type in the *NoNoise*

Prior	*NoNoise*		*Noise*	
	$\underline{t}$	$\bar{t}$	$\underline{t}$	$\bar{t}$
	40.00	*54.14*	*42.50*	*55.00*
1/3	44.34	51.57	45.19	52.77
	(5.40)	(2.61)	(4.59)	(3.67)
	50.00	*50.00*	*48.00*	*54.00*
2/3	46.93	50.96	45.67	51.50
	(5.13)	(1.79)	(4.61)	(2.71)

Table 4: *Effort Levels* (equilibrium values in italics, standard deviation in parenthesis.)

treatments, the frequency is nowhere near 100% for any type in any treatment. Another general observation is that the average signalling distortions (that is, the $a - t$ margins) of the $\underline{t}$ types are larger than those of the $\bar{t}$ types in all treatments.[7] Whereas $\bar{t}$ generally provide too little effort compared to the equilibrium benchmark (except in *NoNoise.67*), $\underline{t}$ provide too much effort with the low prior and too little effort with the high prior.

However, there are a number of observations that are consistent with the equilibrium benchmarks. Table 4 shows that the ranking of effort averages across the four treatments is consistent with the theory. For $\underline{t}$, recall from Hypothesis 1(a) that the lowest effort benchmark (namely 40.0) should occur in *NoNoise.33*, next comes *Noise.33* with a level of 42.5, followed by *Noise.67* (48.0) and the highest effort levels for $\underline{t}$ (50.0) should occur in *NoNoise.67*. The actual averages are ranked precisely in this way. A Jonckheere-Terpstra test rejects the null hypothesis (that effort averages of the $\underline{t}$ type workers are drawn from the same distribution) at p=0.016. This supports Hypothesis 1(a). Conducting the Jonckheere-Terpstra test for the ranking of effort choices by the $\bar{t}$ types yields a similar result (p=0.031), rejecting the null in favor of Hypothesis 1(b).

[7]Indeed, this is true right from the beginning. In period one (where all effort choices are still completely independent), 30 of 44 high types choose $a = \bar{t} = 50$ but only 14 of 36 low types $a = \underline{t} = 40$. A chi-square test indicates that the difference in proportions is significant ($d.f. = 2$, p=0.009). Similarly, the 95% confidence interval of period-one actions for the high type, $[49.01, 50.94]$, includes the myopic best action, 50, whereas the confidence interval for $\underline{t}$, $[41.54, 44.85]$, does not include 40.

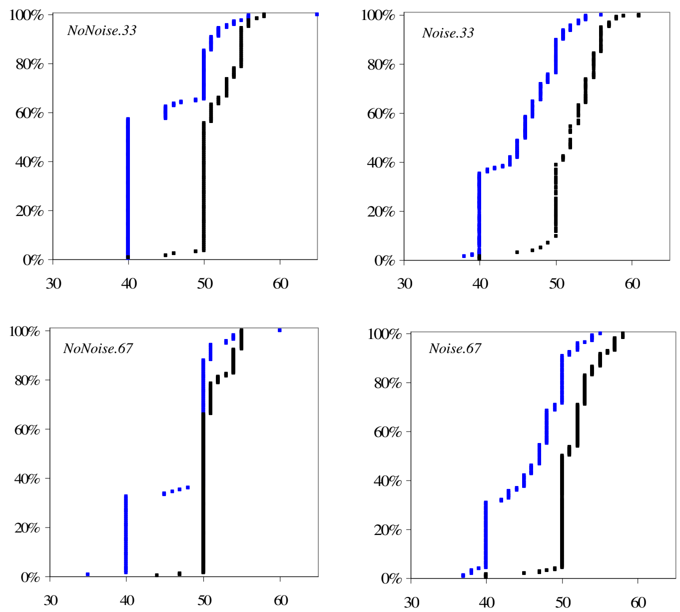

Figure 2: *Cumulative Distribution Functions (CDFs) of Effort Choices*

Now consider the effort differential, $\Delta a = \overline{a} - \underline{a}$. The effort differential is the amount of separation between the types and therefore tells us something about the important pooling *vs.* separating issue. The impact of noise on this variable is ambiguous in theory. There is less separation with noise for the low prior but more separation with noise for the high prior. From Hypothesis 1(c), the theoretical benchmark for Δa^* (ranked in ascending order) is $\Delta a^* = 0.0$ in *NoNoise.67* (the pooling case), 6.0 in *Noise.67*, 12.5 in *Noise.33* and finally 14.1 in *NoNoise.33*. Figure 3 shows the average amount of separation between types for each group and the theoretical benchmark. The picture shows that the theory works well in organizing the data. The ranking of the group averages by treatment is the one predicted except that *Noise.33* has a marginally higher average than *NoNoise.33*. Applying the Jonckheere-Terpstra test on the average effort differentials yields a highly significant rejection of the null hypothesis (p=0.006), therefore providing support for Hypothesis 1(c).

Part of Hypothesis 1(c) is the proposition that, with $\rho_0 = 2/3$, pooling ($\Delta a^* = 0.0$) should occur without noise but separating ($\Delta a^* = 6.0$) with noise. This is an intriguing hypothesis which can be tested directly in a pairwise comparison of these treatments. A ranksum test confirms that the effort differential is smaller in *NoNoise.67* than in *Noise.67* (p=0.042).

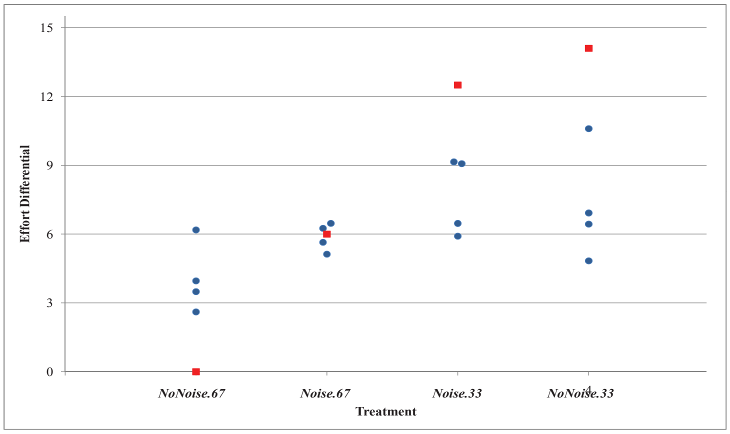

Figure 3: *Effort Differentials* $\Delta a = \overline{a} - \underline{a}$ (equilibrium: squares; group averages: bullets).

The CDFs in Figure 2 also provide evidence in this direction. For *NoNoise.67*, Figure 2 reveals that the pooling equilibrium effort level of 50 is the mode (60.25%). This is true for both types as $\underline{t}$ workers choose $a = 50$ in 51.5% and $\overline{t}$ types in 64.8% of the cases. The frequency of $a = 50$ effort choices is significantly smaller in *Noise.67* and indeed in all of the other three treatments where $a^* = 50$ should not occur in equilibrium (pairwise comparisons with rank-sum tests, all p=0.021 or smaller). This is support for the pooling *vs.* separating hypothesis with $\rho_0 = 2/3$. By contrast, if $\rho_0 = 1/3$ there should be separation both with and without noise and the prediction regarding the effort differential is 12.5 and 14.1 for *Noise.33* and *NoNoise.33*, respectively. As there is separation either way and since the predicted effort differentials do not differ much, unsurprisingly, *Noise.33* and *NoNoise.33* do not differ significantly.[8]

Hypothesis 2 suggests that workers should always signal (that is, choose $a > t$) in treatments with noise. Specifically, $\underline{t}$ should choose $a > 40$, and $\overline{t}$ should choose $a > 50$ with noise. Whereas the CDFs show that in many cases workers do actually choose their myopic best actions, they also show that $\underline{a} = 40$ and $\overline{a} = 50$, respectively, are selected less frequently in the stochastic-signal treatments. Statistical support for this can be obtained by collecting the share of effort choices strictly larger than the (type-specific) myopic best action for each

[8]For the sake of completeness, the third pairwise comparison, *NoNoise.67 vs. NoNoise.33*, confirms the theory (p=0.042).

group. In the eight groups of the noisy treatments, 62.8% of the workers' actions have $a > t$ but this is only the case in 47.8% of the deterministic treatments. A ranksum test reveals that this difference is significant (p=0.023). Note that we obtain this significant result even though one type is predicted to choose $a > t$ also in the *NoNoise* treatments.

We can also check Hypothesis 2 for $\underline{t}$ only. In the *Noise* treatments, $\underline{t}$'s equilibrium actions lie strictly between 40 $(= \underline{t})$ and 50 $(= \overline{t})$. By contrast, in the *NoNoise* treatments, the equilibrium actions for $\underline{t}$ are 40 and 50, and moreover the worker is (at least theoretically) revealed as being $\underline{t}$ when choosing $a \in (40, 50)$. With noise, effort choices between 40 and 50 generate signals that only imperfectly reveal the worker as $\underline{t}$. These considerations are confirmed in the data. The CDFs show that 39.04% of the $\underline{t}$ observations are in the $a \in (40, 50)$ interval with noise but, without noise, only 5.58% are. This difference is significant (rank-sum test, $p < 0.001$). As an aside, the result shows that subjects clearly understood the noisy signal-generating mechanism.

Finally, for the noise treatments, Hypothesis 3 states that signalling distortions $(a - t)$ are larger for the less frequent type. At face value, this hypothesis is clearly rejected. As mentioned above (and as in other signalling experiments), the $\underline{t}$ types distort more in all treatments and we find $\underline{a} - \underline{t} > \overline{a} - \overline{t}$ in all groups of all treatments. However, in *relative* terms, the prediction is supported. From Table 4, note that the $\underline{t}$ types distort more with the high prior whereas the $\overline{t}$ types distort more for the low prior. Testing this formally, the ratio $(\underline{a} - \underline{t})/(\overline{a} - \overline{t})$ is significantly smaller in *Noise.33* than in *Noise.67* (ranksum test, p=0.042). Interpreting the predictions in relative rather than absolute terms (which seems warranted as the $\underline{t}$ types signal too much anyway, right from the first period on), we find support Hypothesis 3.

5.2 Manager (receiver) behavior

Table 5 shows how frequently managers employ the workers. Compared to the equilibrium benchmark, $\overline{t}$ is employed too rarely, and $\underline{t}$ is employed to often. This finding is not particularly surprising given the above result that high types usually signal too little and low types sometimes too much. Regarding the ranking of $\overline{e}$ (the employment rates of $\overline{t}$), we cannot reject the null hypothsis that all treatments are drawn from the same distribution (Jonckheere-Terpstra, p=0.245), that is, we find no support for Hypothesis 4(a). The hypothesis regarding $\underline{t}$, Hypothesis 4(b), turns out to be supported by the data though (Jonckheere-Terpstra, p=0.003). That is, even though quantitatively the predictions fail, the theory still yields a useful qualitative prediction regarding $\underline{e}$.

Prior	*NoNoise*		*Noise*	
	$\underline{t}$	$\bar{t}$	$\underline{t}$	$\bar{t}$
	0.000	*1.000*	*0.060*	*0.833*
1/3	0.232	0.609	0.272	0.609
	(0.42)	(0.50)	(0.45)	(0.49)
	1.000	*1.000*	*0.490*	*0.880*
2/3	0.470	0.864	0.565	0.834
	(0.50)	(0.34)	(0.50)	(0.37)

Table 5: *Employment Rates* (equilibrium values in italics, standard deviation in parenthesis).

Figure 4 shows that average employment rates across the two types, e, meet the equilibrium benchmarks rather accurately in three of the four treatments. That is, regarding average employment rates per treatment, the theory works well even in a quantitative sense. (The exception is *Noise.67*.) A Jonckheere-Terpstra test on the average employments rates confirms that the ranking we observe rejects the null hypothesis (p=0.004). This significant test result is in favor of Hypothesis 4(c).

Employment decisions obviously depend on workers' effort levels. Figure 5 shows the likelihood of getting employed as a function of the effort signals. The probabilities are obtained from simple probit regressions where the response decision $r \in \{0, 1\}$ is a function of s, the signal received. For each treatment there is a separate regression. The probits are based on data from periods 21 to 40 and are clustered at the group level (Woolridge, 2003). As expected, the probability of an $r = 1$ choice increases with the received signal s. Indeed, in all treatments, most of the increase occurs between $s = 40$ and $s = 55$. In all four cases, both the constant and the marginal effects of the effort signal are significant at $p < 0.001$.

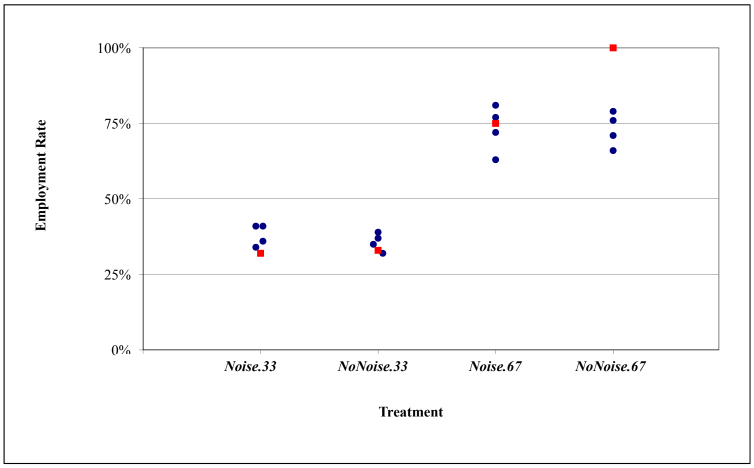

Figure 4: *Employment Rates* (equilibrium: squares; group data: bullets).

All regressions are highly significant with the pseudo R^2 varying between 0.190 (*Noise.67*) and 0.413 (*NoNoise.33*).

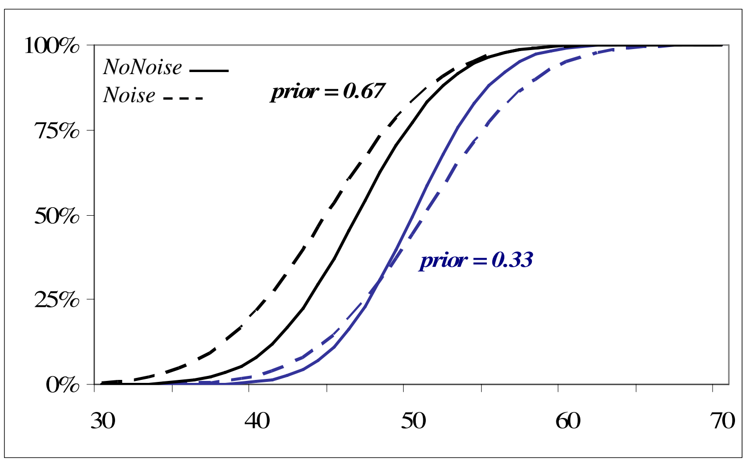

Figure 5: *Probability of Employment as a Function of the Signal*

There are two further intuitive observations from Figure 5. First, the employment likelihood is higher with $\rho_0 = 2/3$. Both with and without noise, $r = 1$ responses are more likely with the high prior. As can be seen in Figure 5, the $\rho_0 = 2/3$ treatments first-order stochastically dominate those with $\rho_0 = 1/3$. Second, the figure shows that the curves in the noise treatments are flatter than their *NoNoise* counterparts for the middle range of effort choices.

For the $\rho_0 = 1/3$ prior, $\Pr(r = 1)$ is larger in the noise treatment for $s \leq 49$ and smaller otherwise. Similarly, for the $\rho_0 = 2/3$ prior, $\Pr(r = 1)$ is larger with noise as long as $s \leq 55$ and smaller otherwise. This is intuitive. A signal of, say, 45 is rather unlikely to have been sent by $\bar{t}$ without noise. With noise, there is some chance $\bar{t}$'s choice was distorted negatively to the level of 45. The reverse is true for choices larger than 50 and 55, respectively. Such high choices are almost surely sent by $\bar{t}$ when there is no noise. With noise, there is still the chance that noise caused the high signal and thus managers are less likely to choose $r = 1$ compared to *NoNoise*, given the same effort choice.

The probit regressions in Table 6 reveal that these two effects are significant. The regressors used are as follows. *Effort Signal* is the signal the manager receives. *Noise* is a dummy equal to one if and only if there is noise in a treatment. Similarly, *Prior67* is equal to one if and only if $\rho_0 = 2/3$. Regression (1) shows that *Prior67* has a significant impact on the likelihood of getting employed whereas *Noise* per se does not. Once we include the interaction terms (see regression (2)), the *Noise* dummy is positive and significant and the *Noise*×*Effort* interaction is negative and significant. The *Prior67*×*Effort* interaction is insignificant but *Prior67* remains significant even when we include the interactions. That is, the likelihood of getting employed is higher with $\rho_0 = 2/3$, but the slope does not change.

We also use the probit regressions to test Hypothesis 5 which is on the switching strategies. Specifically, we run the probits above separately for each group and calculate the median accepted effort choice for each group, that is, the effort choice under which there is a 50% probability of being employed. We compared these median threshold effort levels (one for each group, four for each treatment) to the ranking of switching points given in Hypothesis 5. (If 100% of our subjects behaved consistently with the theory, they would reject every signal below the switching point and employ for every signal above that point, and thus the median signal that results in employment would be qual to the predicted switching point.) A Jonckheere-Terpstra rejects the null hypothesis at p=0.001, supporting the ordinal ranking in Hypothesis 5. Figure 6 shows the empirical median switching points per group in conjunction with the equilibrium switching points.

	(1)	(2)
Effort Signal	0.192	0.228
	(0.013)**	(0.019)**
Noise	0.071	2.557
	(0.168)	(1.122)*
Prior67	0.949	0.772
	(0.175)**	(0.237)**
Noise×Effort		-0.061
		(0.018)**
Prior67×Effort		0.000
		(0.000)
Constant	-9.73	-11.43
	(0.659)**	(0.916)**
LR χ^2	218.6	511.41
p value	0.000	0.000
Pseudo R^2	0.418	0.424

Table 6: *Probit Regressions of Accept Decisions* (clustered at the group level, standard errors in parenthesis, *(**) indicates significance at the 5% (1%) level.)

5.3 How do players respond to others' empirical behavior?

Above, we saw that behavior does not seem to converge fully to equilibrium benchmarks. On the other hand, the support for some implications of the theory indicates that play is far from erratic. This raises the question of how subjects respond to the actual empirical behavior of the other subjects.

We first analyze how workers' actions correspond to the employment decisions. To do this, we determine the optimal effort level given managers' actual employment decisions in periods 21 to 40, separately for each treatment and type. Specifically, we use the acceptance

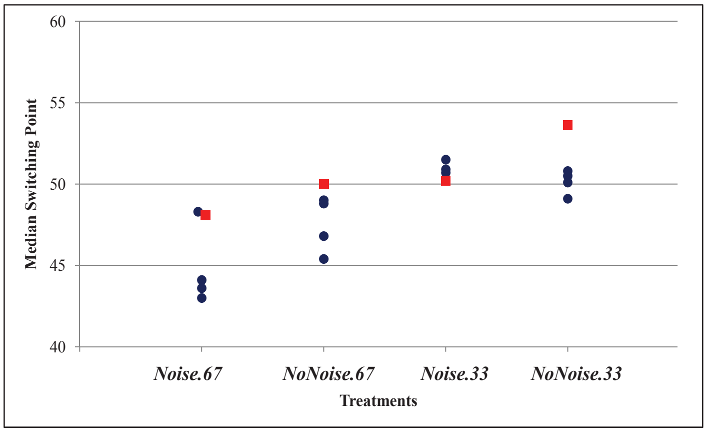

Figure 6: *Switching Points* (equilibrium: squares; group medians: bullets).

probabilities from the probit regressions underlying Figure 5 and calculate workers' expected payoffs from this. For some effort level a,

$$100 - {}^1\!/_2\,(a-t)^2 + \Pr(r=1|a) \times 100$$

is t's expected payoff.

The main insight from this exercise is that average effort choices often do not differ much from the value that maximizes expected payoffs and, whenever they do differ, this can be explained by "flat" expected payoff maxima. Table 7 shows the results.

Prior	*NoNoise*		*Noise*	
	$\underline{t}$	$\bar{t}$	$\underline{t}$	$\bar{t}$
$^1/_3$	40.8	54.9	42.0	55.2
	+3.5	−3.3	+3.2	−2.4
$^2/_3$	48.0	53.2	46.6	52.6
	−1.1	−2.2	−0.9	−1.1

Table 7: *Optimal Effort Choices Given Employment Decisions* (and difference between actual average effort and optimal choices.)

There are eight cases, one for each treatment and each type. For example, the first entry (top left) indicates that, given the empirical receiver behavior in the second half of the experiment, $\underline{t}$'s optimal choice in *NoNoise.33* was 40.8, yet the actual average choice was 3.5 higher (44.3) than the optimal choice of that type in that treatment (where the actual averages are obtained from Table 4). The (absolute) minor differences suggest a certain coincidence of optimal and actual average effort choices. Moreover, note that larger differences (say, three effort units) are subject to the disclaimer that differences in expected payoffs are truly minor. In *Noise.33*, the payoff loss from not playing optimally is merely 1.35% and 1.75% in *NoNoise.33*. The biggest loss in expected payoff (5.6%) occurs for $\bar{t}$ in *Noise.33*. The fact that deviating from optimal behavior causes only minor losses of expected payoffs suggests that the discrepancy of optimal and actual average effort choices, if they occur at all, should be interpreted with caution (Harrison, 1989). The minor losses in expected payoffs also explain why play does not converge to equilibrium more quickly.

Table 7 reveals another result. The optimal effort choices given empirical behavior are surprisingly close to the equilibrium benchmarks (see Table 4). In *NoNoise.33* and *Noise.33*, they almost exactly coincide for both types. In *Noise.67*, the gap between optimal effort choices and the equilibrium benchmark is 1.4 for both types which does not seem to be too far off the mark. Only the pooling prediction in *NoNoise.67* differs from the optimal effort choices.

Next, we analyze how managers' decisions correspond to actual worker behavior. To this end, we calculate at which signal the probability that it was sent by $\bar{t}$ is $1/2$—this is the empirically optimal switching point for risk neutral receivers. This is done with probit regressions ($t = \bar{t}$ as a function of the observed s), separately for each treatment and based on the signals in periods 21 to 40. The actual switching point for each treatment is the lowest signal s that, based on the probits, leads to employment with probability of at least $1/2$. These actual switching points can be taken from Figure 5. Table 8 shows the results of this analysis. It reveals that, although the empirically optimal switching points are larger than the actual switching points throughout, they do correspond closely to one another. In *NoNoise.67* and *Noise.33*, they differ by only 0.5 and 0.6 units of effort, respectively, and in *NoNoise.33* they differ by 1.2 only. In *NoNoise.67* they differ by three units of effort. Again, the small differences between optimal and actual behavior are remarkable.

Prior	*NoNoise*	*Noise*
1/3	50.7	47.7
	−0.6	−1.2
2/3	53.8	45.0
	−3.0	−0.5

Table 8: *Optimal Switching Points Given Effort Decisions* (and difference between actual median and optimal switching points.)

5.4 Discussion

The above results by-and-large confirm the theory in a qualitative sense. Sometimes, the theory also has predictive power in a quantitative sense. A final issue that we discuss here is that, quantitatively, the theory appears to be more closely in line with behavior in the treatment with noise compared to the deterministic setting.

To make this statement precise, consider the absolute difference between equilibrium values given in Table 3 and the treatment averages. As for the effort levels chosen (Table 4), the *Noise* treatment averages are closer to the equilibrium in three of four cases. In Table 5 (employment rates), the comparison reveals that the averages with noise are closer to the equilibrium in all four cases. Similarly, when we look at optimal decisions given empirical behavior of the other players, the noisy variants have a smaller gap between optimal choice and average choices in all four cases of Table 7 and in one of two cases in Table 8.

What could be driving this result? One possibility is that because there are multiple equilibrium configurations in the deterministic case but a unique equilibrium in the stochastic version of our game, coordination on equilibrium might be easier in the noisy case. However, there is only limited evidence that subjects play any of the non-refined equilibria, as seen above. We believe that was may be driving the results is the fact that the stochastic variant captures aspects of decision making that the deterministic variant fails to address. Consider, for example, employment rates. In the deterministic game, in equilibrium, there are no errors in hiring, that is, in a separating equilibrium 100% of high types are employed and 0% of low types; whereas in a pooling equilibrium 100% of workers (i.e., both types) are employed. However, in the data both Type-I errors and Type-II errors occur. That is, high types are sometimes erroneously not hired and low types are erroneously employed. In contrast, with noise, Type-I and Type-II errors are an equilibrium phenomenon, because of the noise. Consequently employment rates are never extreme. Empirically, Type-I and Type-II errors are rather frequent—an aspect of the data that is well accounted for by the stochastic version of the model.[9] As a result, the theory performs better when noise is explicitly modeled.

6 Conclusion

We consider a sender-receiver signalling game in an environment in which the signal-generating mechanism is subject to homoskedastic noise. This noisy setup differs markedly from the standard deterministic case. With deterministic signals, a unique perfect Bayesian equilibrium can only be obtained after the application of equilibrium refinements and, depending on the prior belief, there is either a separating equilibrium or a pooling equilibrium. With noise, the unique equilibrium is separating and the equilibrium actions vary according to the level of the noise and the prior belief.

We further contrast the differences between deterministic and noisy environments by reporting on subject behavior in experiments that we ran. Specifically, we study a frame where workers can choose effort levels as a signal and personnel managers decide whether or not to employ the workers. We compare games without noise to those with noise, and games with a "high" and "low" prior. Many predictions are confirmed in the data in qualitative terms and some are even relatively close in quantitative terms. In particular, the theory has predictive power regarding the main variables of interest, *viz.* effort levels, employment rates, and employment cut-offs (i.e., switching points). As predicted, given a high prior there is more pooling behavior without noise compared to noisy environments. Also consistent with theory is that subjects choose their myopic best action less frequently in the noisy treatments.

[9] Of course, the noisy variant does not include errors in decision making (as, for example, a quantal response equilibrium does), but Type-I and Type-II errors occur both because of noise and because of decision errors. Our point is, hence, that the stochastic model correctly predicts Type-I and Type-II errors even if, partly, they occur for the wrong reason.

Absent ample learning opportunities, signalling experiments usually do not converge fully and often myopic choices and naive mimicking rather than sophisticated play is observed (Cooper and Kagel, 2005). However, even though we also do not see complete convergence in our data, we do find remarkable support of the theory, and where we find deviations these regularly result in near-negligible differences in payoffs compared to optimal play. In particular, subject behavior is distinct across treatments and in line with equilibrium differences of the two model specifications.

Furthermore, while the stochastic model may analytically be more challenging than the deterministic model, subject behavior seems more in line with the equilibrium in the stochastic treatment and model in contrast to a comparison of the empirical data of the deterministic treatment and the deterministic equilibrium. We conjecture that this is due to the fact that stochastic (noisy) settings may be similar to stochastic (non-uniform) play by subjects, leading to greater congruence between the equilibrium of the stochastic game and the empirical data.

References

[1] Ashenfelter, O., Currie, J., Farber, H.S., Siegel, M. (1992). An experimental comparision of dispute rates in alternative arbitrage systems. *Econometrica* 60: 1407–1433.

[2] Backus, D., and Driffill, J. (1985). Inflation and reputation. *American Economic Review* 75(3): 530–38.

[3] Bagwell, K. and Riordan, M.H. (1991). High and declining prices signal product quality. *American Economic Review* 81(1), 224–239.

[4] Brandts, J., and Holt, C.H. (1992). An experimental test of equilibrium dominance in signaling games. *American Economic Review* 82: 1350–1365.

[5] Brandts, J., and Holt, C.H. (1993). Adjustment patterns and equilibrium selection in experimental signaling games. *International Journal of Game Theory* 22: 279–302.

[6] Carlsson, H. and Dasgupta, S. (1997). Noise-proof equilibria in two-action signaling games. *Journal of Economic Theory* 77: 432–460.

[7] Cho, I.-K. and Kreps, D.M. (1987). Signaling games and stable equilibria. *Quarterly Journal of Economics* 102: 179–221.

[8] Cooper, D.J., Garvin, S. and Kagel, J.H. (1997a). Adaptive learning vs. equilibrium refinements in an entry limit pricing game. *The Economic Journal* 107: 553–575.

[9] Cooper, D.J., Garvin, S. and Kagel, J.H. (1997b). Signaling and adaptive learning in an entry limit pricing game. *RAND Journal of Economics* 28: 662–683.

[10] Cooper, D.J. and Kagel, J.H. (2005). Are two heads better than one? Team versus individual play in signaling games. *American Economic Review* 95: 477–509.

[11] de Haan, T., Offerman, T. and Sloof, R. (2008). Noisy Signaling: Theory and Experiment. *Mimeo*, University of Amsterdam.

[12] Fischbacher, U. (2007). Zurich toolbox for readymade economic experiments. *Experimental Economics* 10: 171–178.

[13] Grossman, S., and Perry, M. (1986). Perfect sequential equilibrium. *Journal of Economic Theory* 39(1): 120–154.

[14] Harrison, G. (1989). Theory and misbehavior of first-price auctions. *American Economic Review* 79: 749–762.

[15] Hertzendorf, M.N. (1993). Im not a high-quality firm, but I play one on TV. *RAND Journal of Economics* 24: 236-247.

[16] Hollander, M. and Wolfe, D.A. (1999). *Nonparametric Statistical Methods*, John Wiley and Sons.

[17] Kübler, D., Müller, W. and Normann, H.T. (2008). Job market signaling and screening: an experimental comparison. *Games and Economic Behavior* 64: 219–236.

[18] Mailath, G.J., Okuno-Fujiwara, M. and Postlewaite, A. (1993). Belief-based refinements in signaling games. *Journal of Economic Theory* 60: 241–276.

[19] Matthews, S. and Mirman, L.J. (1983). Equilibrium limit pricing: The effects of private information and stochastic demand. *Econometrica* 51: 981–996.

[20] Milgrom, P. and Roberts, J. (1982). Limit pricing and entry under incomplete information. *Econometrica* 50: 443–459.

[21] Milgrom, P., Roberts, J. (1986). Price and advertising signals of product quality. *Journal of Political Economy* 94: 796-821.

[22] Miller, R. and Plott, C. (1985). Product quality signaling in experimental markets. *Econometrica* 53: 837–872.

[23] Potters, J. and van Winden, F. (1996). Comparative statics of a signaling game: An experimental study. *International Journal of Game Theory* 25: 329–353.

[24] Reinganum, J. and Wilde, L. (1986). Settlement, Litigation, and the Allocation of Legal Costs. *RAND Journal of Economics* 17: 557–566.

[25] Sosis, R.H. and Ruffle, B.J. (2003). Religious ritual and cooperation: testing for a relationship on Israeli religious and secular kibbutzim. *Current Anthropology* 44(5): 713–722.

[26] Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics* 88: 355–74.

[27] Spence, M. (1974). *Market signaling–informational transfer in hiring and related processes*. Harvard University Press, Cambridge.

[28] Wooldridge, J.M. (2003). Cluster-sample methods in applied econometrics. *American Economic Review* 93(2): 133–138.

[29] Zahavi, A. (1975). Mate selection — a selection for a handicap. *Journal of Theoretical Biology* 53: 205–214.

Appendix: Instructions for *Noise.33*

Welcome to our experiment. Please read the following instructions carefully. Please do not talk to your neighbors during the experiment. Should you have any questions, please ask us. We will then come to your booth and answer the question privately.

In this experiment you can earn some real money. How much you will earn depends on your decisions and the decisions of the other participants. Your earnings will be denoted in "points". You payment at the end is equal to the sum of your earnings in each period plus a one-time payment of 4 Euros. For every 500 points, you will receive 1 Euro.

In total, there will be 40 periods.

In today's experiment, a worker and a personnel manager meet in every period. The worker can either be suitable for some task or not. Whether the worker is suitable is randomly decided by the computer in each period, and chances are 1/3 (or 33.33%) that the worker is suitable.

You will be worker or personnel manager, respectively, for five periods each time. Then roles are switched and you have the other role for five periods.

Workers and personnel managers will be randomly matched by the computer in every period. That is, you cannot tell whom you are matched with in each period.

Only the worker is informed whether or not he or she is suitable. The personnel manager has to decide whether or not to employ the worker without knowing if the worker is actually suitable. We will explain the details of your payoffs in detail below but, in principle, the worker gets 100 points if he is employed. The personnel manager gets 100 points if he employs a suitable worker; he also gets 100 points if he does not employ a non-suitable worker.

Before the personnel manager decides, the workers have to do a "test". For the purpose of this experiment, workers have to decide how much "effort" they want to invest when doing the test. We will simply ask for the amount of effort workers want to invest, but we will not do a real test with you. The personnel manager will be informed about the effort the worker invests in the test.

The effort in the test will affect the worker's payment. Have a look at the table below. The left column indicates the effort chosen; the middle column indicates the payoff from the test if the worker is suitable; and the middle column indicates the payoff from the test if the worker is not suitable.

Worker's Payoff in the Test

Effort	Payoff in Points Suitable Worker	 Non-Suitable Worker
25	-212.50	-12.50
26	-188.00	2.00
27	-164.50	15.50
28	-142.00	28.00
29	-120.50	39.50
30	-100.00	50.00
31	-80.50	59.50
32	-62.00	68.00
33	-44.50	75.50
34	-28.00	82.00
35	-12.50	87.50
36	2.00	92.00
37	15.50	95.50
38	28.00	98.00
39	39.50	99.50
40	50.00	100.00
41	59.50	99.50
42	68.00	98.00
43	75.50	95.50
44	82.00	92.00
45	87.50	87.50
46	92.00	82.00
47	95.50	75.50
48	98.00	68.00
49	99.50	59.50
50	100.00	50.00
51	99.50	39.50
52	98.00	28.00
53	95.50	15.50
54	92.00	2.00
55	87.50	-12.50
56	82.00	-28.00
57	75.50	-44.50
58	68.00	-62.00
59	59.50	-80.50
60	50.00	-100.00
61	39.50	-120.50
62	28.00	-142.00
63	15.50	-164.50
64	2.00	-188.00
65	-12.50	-212.50

For example, an effort of 42 gives 68 points to the suitable worker and 98 points to the non-suitable worker; an effort of 56 yields 82 points to the suitable worker and a loss of 28 points to the non-suitable worker. The worker's payoffs in the table will be realized regardless of the personnel manager's decision.

The personnel manager will be informed worker's effort in the test, but not the payoff the effort level chosen yields for the worker. Note that the personnel manager will not be told the effort level chosen with perfect accuracy (more on this below). The personnel manager's payoff does not depend on the effort chosen by the worker.

After the personnel manager gets the information about the worker's effort in the test, he has to decide whether or not to employ the worker. The personnel manager gets paid for this as follows. He gets

- 100 points if he (a) employs a suitable worker or if he (b) does not employ a non-suitable worker,
- 0 points if he (c) does not employ a suitable worker, or if he (d) employs a non-suitable worker.

The worker gets

- the payoff (positive or negative) from the test in the table in every period,
- plus 100 points if he gets employed by the personnel manager and
- plus 0 points if not,

no matter if he is suitable or not.

As mentioned, the worker's effort in the test will not be told the personnel manager with perfect accuracy. How the effort choice of the worker will be communicated to the personnel manager is subject to some random disturbances.

Consider an example. Suppose the worker's effort actual choice is 20. On top of this choice, the computer now adds, or subtracts, a random number. For example, the computer may subtract 4, in which case the personnel manager is told that the effort choice is 16 (rather than being told the actual choice, 20). Or the computer might add 2 on top of the chosen 20. In that case, an effort level 22 is communicated to the personnel manager.

Note that the actual effort level chosen in the test determines for the worker's payoff, not the disturbed effort level which the personnel manager learns.

As a general rule, smaller distortions are more likely than larger ones, and disturbances are possible in either direction (adding to or subtracting from the effort level chosen).

For your information, the figure below shows the deviations from the actual choices in 100 cases in the past. You should expect similar distortions in today's experiment.

Note that the magnitude of the deviation does not depend on the effort level chosen. Indeed the level chosen has no impact whatsoever on the distortions the computer adds or subtracts.

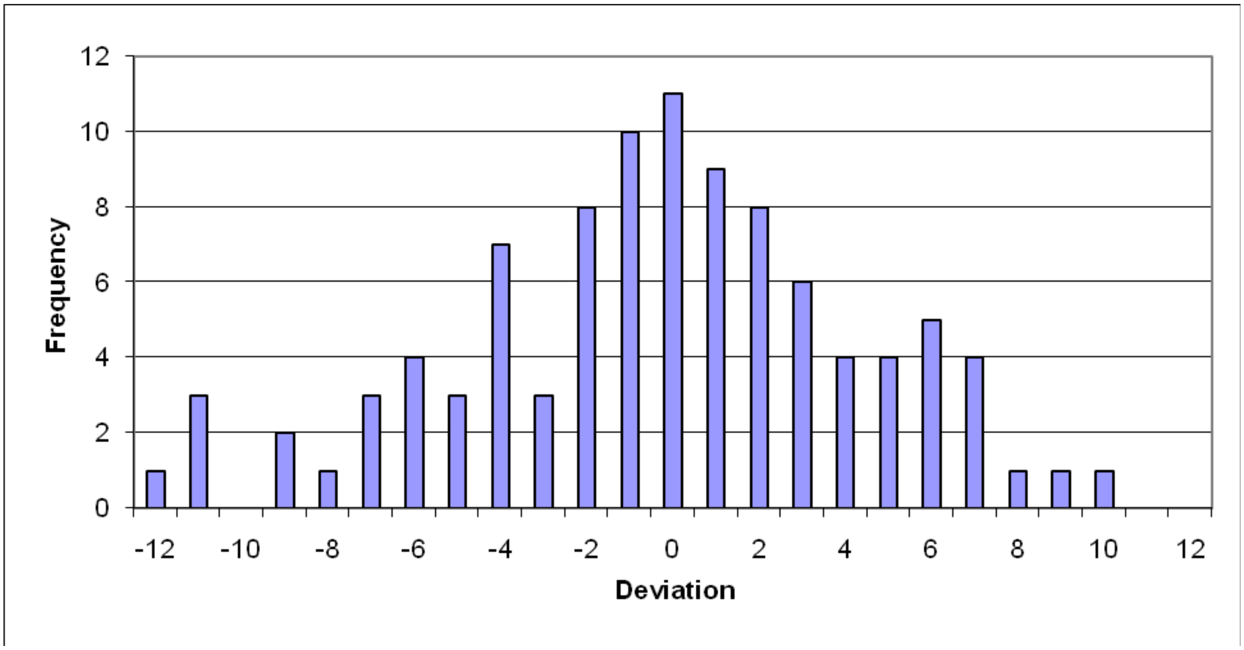

The figure shows the deviation from the true effort choice on the horizontal line and the frequency of these deviations (out of the 100 cases) with the help of the vertical bars.

In the figure you can see that, for example,

- in 4 of these 100 past cases, the computer added 4 to the chosen effort level
- in 1 of these 100 past cases, the computer added 9 to the chosen effort level
- in 3 of these 100 past cases, the computer subtracted 7 to the chosen effort level
- in 11 of these 100 past cases, the computer did not change the effort at all (the "0" column) so that the personnel manager learned the actual effort choice.
- deviations larger than +11 or smaller than -12 did not occur in 11 of these 100 past cases; such deviations are possible but unlikely.

The personnel manager will be informed about the randomly disturbed effort value before making the employment decision. The personnel manager will find out about the actual effort level chosen only at the end of a period.

Workers get their payments for the test according to their actual effort choice, not the disturbed message the personnel manager gets.

Let us summarize the experiment

1. The computer decides randomly whether or not the worker is suitable for the job. The chances for suitability are 1/3 (33.33%) for workers in each period. The worker (but not the personnel manager) is informed about the suitability of the worker.
2. The worker has to choose an effort level in the test (see the table), and the computer informs the personnel manager about the test effort. The effort level is subject to random disturbances.
3. The personnel manager decides whether or not to employ the worker
4. Payoffs are:
   - If you are a worker: the payment from the effort choice in the test plus 100 points if you are employed
   - If you are an personnel manager: 100 points if you employ a suitable worker, or if you do not employ a non-suitable worker