Accepted Manuscript

Title: Career concerns incentives: An experimental test

Authors: Alexander K. Koch, Albrecht Morgenstern, Philippe Raab

PII: S0167-2681(09)00123-1
DOI: doi:10.1016/j.jebo.2009.04.014
Reference: JEBO 2382

To appear in: *Journal of Economic Behavior & Organization*

Received date: 18-4-2008
Revised date: 22-4-2009
Accepted date: 23-4-2009

Please cite this article as: Koch, A.K., Morgenstern, A., Raab, P., Career concerns incentives: An experimental test, *Journal of Economic Behavior and Organization* (2008), doi:10.1016/j.jebo.2009.04.014

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Career concerns incentives: An experimental test*

Alexander K. Koch^a, Albrecht Morgenstern^b, and Philippe Raab^{c†}

^aAarhus University, Royal Holloway University of London and IZA

^bGerman Federal Chancellery and IZA

^cAllianz Versicherungs-AG

April 22, 2009

Abstract

Holmström's (1982/99) career concerns model has become a workhorse for analyzing agency issues in many fields. The underlying signal jamming argument requires players to use information in a Bayesian way, which is difficult to directly test with field data: typically little is known about the information that individuals base their decisions on. Our laboratory experiment provides prima facie evidence: i) the signal jamming mechanism successfully creates incentives on the labor supply side; ii) decision errors take time to decrease; iii) while subjects' average beliefs are remarkably consistent with play, a mild winner's curse arises on the labor demand side.

JEL-Classification: C91, D83, L14

Keywords: Incentives; Reputation; Career concerns; Signal jamming; Experiments

*We are grateful to Armin Falk, Simon Gächter, Lars Koch, David Myatt, Hans Normann, Jörg Oechssler, Klaus Schmidt as well as an Associate Editor and two anonymous referees for valuable comments. Moreover, we thank Bernd Irlenbusch and Dirk Sliwka for providing us with information on their experiments. Financial support from the Bonn Graduate School of Economics is gratefully acknowledged. The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the German Federal Chancellery or the Allianz Insurance Company.

^{†a} corresponding author: School of Economics and Management, Aarhus University, Building 1322, 8000 Aarhus C, Denmark. Tel.: +45 8942 2137, E-mail: akoch@econ.au.dk; ^b Bundeskanzleramt (German Federal Chancellery), 11012 Berlin, E-mail: albrecht.morgenstern@gmx.net; ^c Allianz Versicherungs-AG, An den Treptowers 3, 12435 Berlin, Germany. E-mail: philipperraab@yahoo.de.

1 Introduction

Building on Holmström's (1982/99) seminal work, the career concerns model has become an important workhorse for the analysis of agency issues. Both the model itself and the signal jamming logic that it develops have been applied by economic theorists in a wide array of fields, including labor relationships and organizational economics (for a survey see Borland 1992), industrial organization (e.g., Fudenberg and Tirole 1986, Samuelson et al. 1994), financial economics (e.g., Stein 1989), and political economy (e.g., Lohmann 1998, 1999, LeBorgne and Lockwood 2006). However, the abundance of theoretical work has been met by relatively few empirical studies on the effectiveness of career concerns (Gibbons and Murphy 1992, Fee and Hadlock 2003).

A major difficulty in using field data to test the career concerns model is that it is virtually impossible to control for the information that decision makers have at hand. This paper reports results from a laboratory experiment designed to fill this gap: it provides *prima facie* evidence on career concerns incentives by *directly* testing the empirical viability of the underlying signal jamming mechanism.

In its most basic version, the career concerns argument goes as follows. An agent (e.g., a manager) can take an action (exert costly effort) to improve the signal (output) that principals (potential employers) receive about her productive abilities. Even though all parties are symmetrically informed, the agent has an incentive to interfere with the principals' updating by exerting effort to make the signal look more favorable, i.e., engage in *signal jamming*. In equilibrium, principals take this into account and the agent's signal is discounted accordingly. As a consequence, the agent is trapped into providing the effort to meet principals' tougher standards and not to look worse than she actually is.

Our experimental design implements a simplified version of Holmström's (1982/99) model. Managers are endowed with a random ability level which is unknown to them and the other subjects. In the first period they choose the level of costly effort. Together with the manager's random ability level this effort determines the performance signal which is observed by the firm subjects. Firms are interested only in a manager's ability and compete in the second period by making wage offers. In the *Hidden Ability* treatment, which implements the simplified Holmström (1982/99) model, only the sum of each manager's (unknown) ability and effort is revealed to the firm subjects before they make wage offers. In the *Revealed Ability* treatment, which serves as a control, both ability and effort are revealed separately before second period wage offers are made. The theoretical prediction is that effort incentives should exist only in the Hidden Ability treatment: ability is (symmetrically) unknown to the labor market participants and managers can thus push the firms' beliefs about their ability upwards by exerting effort.

For the sake of comparability, we chose a design and parameters that are similar to the experiment by Irlenbusch and Sliwka (2006), who investigate the impact of transparency on incentives in a more complex informational environment that allows for both signaling and signal jamming (see below).

The first contribution of our experiment is that it delivers the first clean test of whether the qualitative predictions on signal jamming from the career concerns model hold up in the laboratory. This provides evidence on how solid the foundations are of the theoretical models that build on career concerns being effective in creating incentives. The advantage of our controlled laboratory environment over field data is that we can even move beyond this. Finding that subjects tend to play according to the predictions derived using the signal jamming logic does not actually tell us about the beliefs that led subjects to such behavior. Theoretical career concerns models rely on the Bayesian Equilibrium concept, which puts high demands on the rationality of players and requires sophisticated reasoning about other players' strategies (e.g., Fudenberg and Tirole 1991). The second contribution of our experiment is that it elicits subjects' first- and second-order beliefs, and investigates whether players correctly conjecture each others' actions and beliefs and engage in behavior consistent with this. This tells us whether the concept of Bayesian Equilibrium provides a good approximation of actual play.

To preview, our main finding is that career concerns are effective in providing effort incentives. In line with the theoretical predictions, our experimental managers invest more costly effort when their ability is not revealed to firms. We obtain both direct and indirect evidence that subjects adjust their beliefs in the direction predicted by the theory. Firm subjects make roughly zero profits on average. This requires accurate estimates of managers' ability levels, despite managers' attempts to jam the signal that inference is based upon. Stated first- and second-order beliefs are remarkably consistent with each other and actual play in the game. This is a precondition for the Bayesian Equilibrium concept to produce a reasonable approximation of the outcome of individual behavior under incomplete information. There are however decision errors for some subjects on both sides of the market. For managers these take time to decrease. Firms correct their behavior more quickly because decision errors cause relatively big losses. Decision errors by managers complicate the inference task for firm subjects, which is reflected in a mild winner's curse.

Our results are in contrast to Irlenbusch and Sliwka (2006), who report for their experiment that average effort is higher whenever effort and ability can be observed separately. A comparison of our setups suggests that these effects are due to the signaling opportunities arising from repeated gift

exchange in their experimental labor market, and do not contradict the relevance of career concerns incentives as such. We discuss this in more detail in Section 6.

The paper is organized as follows. Section 2 introduces the simple variant of Holmström's (1982/99) career concerns model that we implement in the laboratory. Section 3 derives hypotheses to be tested and Section 4 presents the experimental design and reports procedures. The analysis of results is contained in Section 5 and Section 6 compares our findings with those in Irlenbusch and Sliwka (2006). Section 7 concludes.

2 Theoretical background

The central feature of Holmström's (1982/99) model is that the labor market determines a manager's wage based on an imperfect signal about her ability. Therefore, we will refer to the simplified version of this setup as the *Hidden Ability* model. Although career concerns occur in a number of situations and our experiment has a neutral frame, we illustrate the argument using Holmström's original interpretation of a manager-labor market relationship.

2.1 The Hidden Ability (HA) model

A manager of uncertain ability first chooses an effort level which influences the signal observed by the labor market. More specifically, the manager is endowed with an ability level¹ $a \in \{0, 1, 2, \dots, 19\}$, which is unobservable both to the manager and the principals. What is common knowledge is that a is uniformly distributed on the set $\{0, \dots, 19\}$. The manager chooses some effort level $e \in \{0, 1, 2, \dots, 19\}$ at a private cost $c(e)$, which is increasing and convex in e (for details see Appendix A). Together with the manager's random ability level this effort determines the performance signal which is observed by the labor market:

$$y = a + e, \quad a, e \in \{0, 1, 2, \dots, 19\}. \quad (1)$$

Upon receiving this information, $N \geq 2$ firms $i \in \{1, \dots, N\}$ active in that market offer the manager a wage $w_i(y)$. If the manager accepts an offer $w_i(y)$ she has a payoff of

$$u(e, w_i(y)) = w_i(y) - c(e). \quad (2)$$

¹We depart from Holmström's version of the model by imposing a finite support on the ability parameter a to be able to implement the model in a laboratory setting. For the sake of comparability of experimental results, our model and parameter choices follow Irlenbusch and Sliwka (2006).

If she declines all offers she receives her outside option (normalized to zero) and still bears the cost $c(e)$. The employing firm i 's profit is the manager's ability less the wage offered:²

$$\pi_i(a, w_i(y)) = a - w_i(y). \quad (3)$$

Firms that do not succeed in hiring a manager make zero profits.

Equilibria

To construct a Perfect Bayesian Equilibrium with effort level e^* , we proceed by solving backward. In the continuation game following firms' wage offers, the manager maximizes her payoff by accepting the highest offer $w_{max}(y) = \max\{w_1(y), \dots, w_N(y)\}$. Since there are multiple firms, which all attach the same value to the manager's services, there is Bertrand-like competition if all firms hold the same beliefs. Thus, in such an equilibrium at least two firms offer a wage equal to the expected ability of the manager conditional on the signal y given market belief e^* :

$$w_{max}(y) = E[a | y, e^*]. \quad (4)$$

Since $y = a + e$ and firms hold the belief that the manager exerted effort $e = e^*$, we have $w_{max}(y) = E[a | y, e^*] = y - e^* = a$ for realizations that are on the equilibrium path. Firms' beliefs about a when they observe performance signals $y \in \{0, \dots, e^* - 1, e^* + 20, \dots, 38\}$ that are off the equilibrium path are not pinned down by Bayes' rule, however. To be consistent with the model structure these beliefs only have to fulfill the following conditions:

$$E[a | y, e^*] \leq \min\{y, 19\}, \quad (C 1)$$

$$E[a | y, e^*] \geq \max\{y - 19, 0\}. \quad (C 2)$$

For example, a manager with $y = 2$ can at most have ability 2 (implying that zero effort was provided); and a manager with $y = 38$ must have ability 19 (implying that the maximum possible effort level 19 was chosen).

Given firms' beliefs, the manager chooses effort $e \in \{0, \dots, 19\}$ to maximize her future wage minus the effort costs, $E_y[w_{max}(y)|e] - c(e)$. The manager's effort e greater (less) than the market's belief e^* increases (decreases) the perceived ability of the manager for all y which can be observed under the equilibrium effort e^* , i.e., for $y \in \{e^*, \dots, e^* + 19\}$. Note that for signals y which are not in the

²Note that the manager's first period effort does not affect firm profits and is purely wasteful. This eliminates gift exchange motives (see Section 6).

equilibrium set this need not hold, because beliefs are only constrained by the conditions (C 1) and (C 2). In equilibrium, e^* and the corresponding beliefs must be such that the manager has no incentive to deviate from e^* :

$$e^* = \operatorname{argmax}_{e \in \{0, \dots, 19\}} \{E_y[w_{\max}(y)|e] - c(e)\}. \quad (5)$$

Although this means that, in equilibrium, the manager cannot fool firms in their assessment of her ability a , she still has no incentive to deviate to effort levels below e^* . The labor market expects exactly e^* and discounts the signal y accordingly, so that any lower effort level would make a manager look worse than she actually is.

Without further assumptions on the beliefs off the equilibrium path, it turns out that this leaves room for multiple equilibria. In particular, it can be shown that there exist beliefs which support pure strategy equilibrium effort levels $e^* = 0, \dots, 12$ (see Appendix B).³ A unique equilibrium effort level $e^* = 12$ can be achieved by imposing the following monotonicity restriction on beliefs:

$$E[a | y, e^*] = \begin{cases} 0 & \text{for } 0 \leq y < e^*, \\ y - e^* & \text{for } e^* \leq y < e^* + 20, \\ 19 & \text{for } e^* + 20 \leq y \leq 38. \end{cases} \quad (\text{MR})$$

This yields the following result:

Proposition 1 (Hidden Ability model)

If firms observe signal $y = a + e$, there exist pure strategy Perfect Bayesian Equilibria under which the manager chooses an effort level $e^ \in \{0, \dots, 12\}$ and at least two firms offer $w_{\max}(y) = y - e^*$. Under belief restriction (MR), all Perfect Bayesian Equilibria involve effort level $e^* = 12$.*

Hence, career concerns suffice to motivate the manager to exert costly effort if all parties only observe an imperfect measure of ability, that the manager can manipulate.

2.2 The Revealed Ability (RA) model

In what we refer to as the *Revealed Ability* model, firms not only observe the performance signal y but also the manager's ability a prior to their wage offers. As this is a game of complete information, the comparable solution concept to the one used in the Hidden Ability model is Subgame Perfect Nash Equilibrium. Solving backward, $w_{\max}(y, a) = a$ because of Bertrand competition among firms. The manager's wage hence is independent of effort, and because effort is costly, all equilibria have $e^* = 0$.

This is summarized in the following proposition.

³This is in contrast to Holmström (1982/99), who derives a unique equilibrium effort level in a model where a has full support. However, multiple equilibria occur frequently also in other models of career concerns (e.g., Dewatripont et al. 1999).

Proposition 2 (Revealed Ability model)

If firms observe separately a and y , all Subgame Perfect Nash Equilibria have the following properties: the manager chooses an effort level $e^ = 0$ and at least two firms offer $w_{\max}(y, a) = a$.*

In contrast to Proposition 1, without uncertainty about ability there is no reason for the manager to manipulate her performance signal by exerting costly effort, because firms are only interested in ability.

3 Research questions

The Revealed and the Hidden Ability models serve as basis for finding evidence on the effectiveness of career concerns in our experimental setting. Accordingly, the treatments which implement them are labeled *Hidden Ability* (HA) and *Revealed Ability* (RA), respectively. Section 4 describes in detail the experimental procedures.

3.1 Effort

Our first goal is to test whether career concerns create effort incentives. That is, we want to establish whether uncertainty about the manager's ability in the HA treatment induces her to exert costly effort to generate a more favorable performance signal. The RA treatment serves as control. Our theoretical predictions in Propositions 1 and 2 lead to the following hypothesis:

Hypothesis 1

In the Hidden Ability (HA) treatment, effort is higher than in the Revealed Ability (RA) treatment: $e^{HA} > e^{RA}$.

In our setup we are able to establish the validity of the career concerns signal jamming mechanism whenever HA subjects choose a strictly larger effort than RA subjects. There is one caveat though, arising from the multiple equilibria for the HA setting. If one found in the experiment that $e^{HA} = e^{RA} = 0$, this could mean either i) that there are no incentives from career concerns, or ii) that there are incentives from career concerns, but HA subjects happened to coordinate on the equilibrium in which $e^* = 0$.

3.2 Beliefs

The theoretical concept of Perfect Bayesian Equilibrium imposes certain requirements on players' beliefs in the Hidden Ability model. First, players' beliefs have to be consistent with the model

structure and satisfy (C 1) and (C 2). Second, beliefs should be mutually consistent and in line with equilibrium actions. To analyze whether this is a good approximation of behavior in our experiment, we elicit firms' beliefs and managers' second-order beliefs in the HA treatment (see Section 4).

Hypothesis 2

In the Hidden Ability (HA) treatment, firms' beliefs $b^F(y)$ and managers' second-order beliefs $b^M(y)$ for performance signals $y \in \{0, 1, 2, \dots, 38\}$

- (a) *satisfy conditions (C 1) and (C 2);*
- (b) *are mutually consistent: $b^M(y) = b^F(y)$;*
- (c) *are consistent with firms' two maximum wage offers: $w_{max}(y) = b^F(y)$;*
- (d) *are consistent with equilibrium expected ability levels: $E[a|y, e^*] = b^F(y)$;*
- (e) *are consistent with managers' effort choices:*

$$e = \operatorname{argmax}_{e \in \{0, \dots, 19\}} \{E_y[b^M(y)|e] - c(e)\}.$$

3.3 Wage offers

Career concerns incentives are indirect: a higher performance signal in the HA treatment is not valuable per se, but only has a value for the manager because it leads to higher wage offers. This mechanism must break down whenever ability is revealed. Hence, in the HA treatment wage offers should depend on the performance signal (i.e., equally on ability and effort), while in the RA treatment only ability should matter.

Hypothesis 3

In the HA treatment, wage offers depend on the performance signal, i.e., equally on effort and ability, while in the RA treatment, wage offers depend only on ability.

3.4 Profits

HA managers' average profits should be (weakly) lower than in the RA treatment: average wages are the same but managers exert costly effort to manipulate the performance signal in the former and none in the latter. Bertrand competition implies that average firm profits should be zero, regardless of the information structure. Firm profits in the HA treatment are likely to exhibit considerable variance due to uncertainty about ability. In the RA treatment, however, firm profits should be less volatile because ability is revealed prior to wage offers.

Hypothesis 4

In the RA treatment, managers earn (weakly) more than managers in the HA treatment. Firms in both treatments earn zero profits on average, but the variance of firm profits in the HA treatment is larger than in the RA treatment.

4 Procedures

The experiment was run computer based with the experimental software z-tree (Fischbacher 2007) in the Experimental Laboratories of the University of Bonn and Royal Holloway, University of London. We conducted six sessions with a total of 84 participants, which provided us with six independent observations for each treatment.⁴ The participants were students from all subjects, of which less than one third were Economics majors. 44 percent were female and for 27 percent this was the first experiment they participated in.

In the beginning of each session, subjects were randomly assigned to separate cubicles, instructions were distributed and read aloud.⁵ Any questions concerning the experiment were only allowed on a bilateral basis. The experiment on the computer was only started after all subjects had successfully answered a questionnaire requiring them to calculate simple examples on how actions determine the round profit for a manager and for a firm.

A session consisted of two separate groups with seven participants each. Within each group, three subjects were assigned the manager role (*A-players*) and four subjects were assigned the firm role (*B-players*). Roles were kept fixed during the whole experiment. In each session, subjects completed 12 rounds of play. Each round consisted of three stages, as illustrated in Figure 1.

Figure 1 about here.

In the first stage, each manager chose her effort level e (*number*) from the set $\{0, 1, \dots, 19\}$ (effort costs are given in Appendix A). Each manager was endowed with an ability parameter a (*base value*) drawn from the set $\{0, 1, \dots, 19\}$, where each value had equal probability.⁶ This ability parameter was not revealed to anyone. That is, neither the firms nor the managers themselves knew the ability parameter.

In the second stage, firms made wage (*transfer*) offers from the set $\{0.00, 0.01, \dots, 38.00\}$ to each of the three managers in their group. Treatments differed with regard to the information available to firms.

⁴An additional session for each treatment exists where subjects choosing effort of 19 did not get the correct cost deducted due to a computer bug. The London sessions replace these. The qualitative findings in the paper obtain, however, even when including the two observations per treatment affected by the bug.

⁵To avoid framing effects we used a neutral representation of the model. Instructions are given in Appendix C.

⁶The three sequences of ability parameters for managers were drawn before the experiments and kept fixed across all observations to facilitate comparisons within and between treatments.

HA firms were shown the performance signal y of each manager (the sum of the current round's base value a and number e , labeled *result*). RA firms observed both the performance signal y and the ability parameter a for each manager. We randomized the order of appearance of managers' performance signals/ability levels on the screen in every round. Instructions mentioned explicitly that firms could therefore not infer the identity of managers from the position of their signal/ability level on the screen. In the third stage, every manager was presented with her four wage offers and could decide whether to accept one of them or reject all offers. Again, we randomized the ordering of wage offers in each round, and communicated this to the subjects.

A round concluded with feedback and display of profits. Managers saw their signal, effort, cost of effort, accepted wage offer, and the resulting round profit. This was computed as the accepted wage offer (or zero if all offers were rejected) minus the cost of her chosen effort level. Firms saw for each accepted wage offer the respective manager's signal, ability, wage paid as well as the resulting profit from that manager and the round profit. The later was the sum of the ability parameters of all the managers who accepted this firm's respective wage offers, minus the wages paid out to these managers. If all three wage offers were not accepted the total round profit displayed was 0.

BELIEF ELICITATION. After the eighth round in the HA treatment, we elicited subjects' beliefs by questionnaire. Firms were asked to indicate for every possible realization of the performance signal $y \in \{0, 1, 2, \dots, 38\}$ what ability parameter they would expect such a manager to have on average (first-order beliefs). In a similar fashion, managers were asked to report what belief they would expect from a firm on average for a given signal realization (second-order beliefs).

A comment is in order. Belief elicitation was not incentivized to rule out that subjects state distorted beliefs as a hedge against adverse outcomes from other actions in the experiment. To allow us to ascertain that results are not driven by other effects of belief elicitation, the belief questionnaire was administered after round 8 only. For a recent discussion of these issues and evidence on belief elicitation effects see Blanco et al. (2008).

INITIAL ENDOWMENT AND EXCHANGE RATE. Each firm could hire as many of the three managers as it wanted to, and the number of firms exceeded that of managers. We thus expected aggressive competition in wage offers, resulting in roughly zero round profits for firms. To finance their wage offers, each firm received an initial endowment of 180 experimental currency units in the first round. Managers were expected to earn significant round profits. Because we wanted firms and managers to earn roughly the same amount in the experiment, the managers' initial endowment was adjusted to 80 experimental currency units.⁷

⁷The instructions mentioned initial endowments, but levels were only disclosed to the subjects assigned the respective

Table 1: Summary statistics for effort choices of managers

	Hidden Ability (HA)	Revealed Ability (RA)
Average effort	6.96	4.17
Standard deviation	4.89	4.34
Zero effort played (percent)	17.13	40.28

At the end of the experiment, the total profit was converted converted into Euros (Bonn)/Pounds Sterling (London) at an exchange rate of EUR/£ 0.05 per experimental currency unit. In the HA treatment, we additionally paid EUR/£ 3 to compensate for the longer duration of sessions due to the belief elicitation after the eighth round. In total, HA sessions lasted for roughly 1 hour and 20 minutes with average earnings of EUR 11.45 (EUR 13.27 managers/EUR 10.09 firms) in Bonn and £ 11.67 (13.86/10.02) in London. RA sessions lasted for about 1 hour and average earnings were EUR 9.05 (10.49/7.97) and £ 8.77 (8.84/8.73).

5 Results

5.1 Effort (Hypothesis 1)

Average effort for HA managers is 6.96 – almost three effort units more than their RA counterparts’ average of 4.17 (see Table 1). The difference is highly significant (p-value 0.01, Wilcoxon-Mann-Whitney (WMW) test)⁸ and supports Hypothesis 1. This pattern is stable across periods, as average HA effort exceeds average RA effort in any given round (see Figure 2).

Figure 2 about here.

The experimental setup is relatively complicated. One manifestation of this seems to be that average effort in the RA treatment differs from the theoretical prediction of $e^{RA} = 0$. But this prediction is a corner solution, so decision errors cannot wash out (as they would with an interior solution). Mistakes therefore push the average effort above the predicted level (Andreoni 1995). Figure 3 gives a more detailed picture of the distribution of effort choices. Indeed, RA effort levels are concentrated on zero. In the HA treatment effort is more dispersed, and the distribution first-order stochastically dominates that under the RA treatment (p-value < 0.001, Kolmogorov-Smirnov test).

Figure 3 about here.

role via the current balance on the computer screen, to preclude any impact of perceived inequalities in endowments.

⁸In the paper, we report results based on all periods. All our qualitative findings also obtain when using only the periods 1-8 before beliefs are elicited by questionnaire.

What drives the behavior of RA subjects who make mistakes? Apparently, learning during the game is mainly a reaction to negative round profits, reflecting loss aversion. When offering too high wages, RA firms immediately realize a loss. The average round profit following such ‘decision errors’ is -16.39 (see Section 5.4). In contrast, managers’ mistakes very rarely result in losses. Their average round profit following ‘decision errors’ is 8.31 . Out of 129 rounds where RA managers chose $e > 0$, in only 11 did they experience a loss; but even then the average loss of -4.46 is small relative to that experienced by firms who offer too high wages. In this sense, psychological feedback from round payoffs is rather weak for manager subjects.

Figure 4 about here.

Indeed, as Figure 4 shows, learning eliminates ‘errors’ more quickly for firms than for managers. Nevertheless, the incidence of positive effort drops to around 50 percent, and average effort for all managers decreases to around 3. This decline is driven by a group of 12 managers whose effort exceeded 1 at least once (high effort group). Their average effort decreases from 8.92 to 4.92. For the remaining 6 managers (low effort group) average effort moves from 0.50 to 0.17. Overall, in 40 percent of all cases managers provided zero effort, in accordance with the theoretical prediction. This is significantly more than in the HA treatment, where zero effort is provided in only 17 percent of all cases (p-value 0.08, WMW test).⁹

5.2 Beliefs

5.2.1 Hypothesis 2 (a) and (b)

Figure 5 about here.

Do beliefs satisfy the consistency conditions from Hypothesis 2 (a)? In Figure 5 the average of firms’ beliefs, $b^F(y)$, and managers’ second-order beliefs, $b^M(y)$, almost always lie in the hatched area where beliefs satisfy the consistency conditions (C 1) and (C 2). Even at an individual level beliefs are surprisingly consistent: only 46 out of the 1,638 reported beliefs – less than 3 percent – violate these conditions.

In line with Hypothesis 2 (b), $b^F(y)$ and $b^M(y)$ lie very close to each other in Figure 5. The hypothesis that $b^M(y) = b^F(y)$ cannot be rejected at the 5%-level based on the 39 paired overall averages for each observation (p-value 0.09, Wilcoxon-Signed-Ranks (WSR) test). An alternative hypothesis is

⁹This test compares the rankings of the number of times managers play zero in each group for each treatment. Note that a chi-squared test comparing the number of times zero effort is chosen relative to the number of times positive effort is chosen in the two treatments would require the assumption that effort decisions *within* each group are independent. If one is willing to make this assumption, the p-value is < 0.001 (Chi-squared test).

that stated beliefs are just a naïve average of all ability levels that a signal realization y admits (e.g. for $y = 2$ this would be $(0 + 1 + 2)/3 = 1$).¹⁰ Figure 5 however shows that these ‘naïve averaging’ beliefs tends to be lower than firms’ and managers’ beliefs, and we reject the hypothesis that stated beliefs coincide with the ‘naïve averaging’ beliefs (both p-values < 0.001 , WSR test). Beliefs also do not satisfy the restriction (MR), which would lead to a unique equilibrium prediction (both p-values < 0.001 , WSR test).

According to the theoretical prediction, beliefs should have a slope equal to unity on the equilibrium path, i.e., over a range of 19 different values of y . But reported individual beliefs and their averages typically do not display this feature. There are two explanations for this, besides simple mistakes. First, a firm may expect managers to experiment or make errors when choosing effort. In forming beliefs a firm may therefore use different conjectures about the expected effort level exerted depending on the realization of the signal y (e.g., while $y = 23$ rules out effort of less than 4, $y = 6$ does not). The ‘naïve averaging’ beliefs discussed above are an example for this. Second, stated beliefs are averages over the beliefs that a firm holds for three different managers within the same experimental group, and there is strategic uncertainty regarding the equilibrium effort. Consistent with these two explanations, firm subjects appear to realize that managers play different effort levels¹¹ and thus discount higher realizations of the signal more than lower ones.¹²

5.2.2 Hypothesis 2 (c) and (d)

Figure 6 about here.

Do firms’ beliefs accurately reflect average actual play, and is firms’ behavior consistent with stated beliefs? From Figure 6 it appears that average wage offers (grey circles) lie below average beliefs (grey line) for higher signal values. Indeed, the hypothesis $w(y) \geq b^F(y)$ is rejected (p-value 0.006, WSR test). According to Hypothesis 2 (c), the *maximum* two wage offers made to a manager should satisfy $w^{max}(y) = b^F(y)$. The highest wage offers, however, are larger than average firm beliefs (p-value < 0.001) – pointing to the winner’s curse phenomenon discussed below.

In the HA post-experiment questionnaire we ask for firm subjects’ estimates of the average manager effort level: the average response of 9.38 (standard deviation 4.86) overestimates the actual average

¹⁰We are grateful to a referee for pointing out this possibility.

¹¹In the post-experiment questionnaire 12 out of 24 firms report that they thought managers chose “very different effort levels” rather than “very similar effort levels”.

¹²To see this point, assume that a firm thinks that each manager plays according to a different ‘equilibrium’ strategy, say $e_1 = 0$, $e_2 = 6$, and $e_3 = 12$. Moreover, suppose that for out-of-equilibrium signals y the firm’s beliefs are such that the overall profile is monotone. Then the belief about the average manager has a slope of $1/3$ in the range $[0, 6]$, of $2/3$ in the range $[6, 12]$, of 1 in the range $[12, 19]$, of $2/3$ in the range $[19, 25]$, of $1/3$ in the range $[25, 31]$, and is flat in the remaining part.

Table 2: Best replies versus actual effort choices (HA treatment)

	Best reply against		Actual effort
	own stated beliefs	firms' average beliefs	
Average	6.08	6.33	6.96
Standard deviation	2.12	1.37	4.89
Best reply against average wage offers			5
Best reply against average maximum wage offers			5

effort level of 6.96 (standard deviation 4.89). This point estimate raises the question whether stated firm beliefs are a good estimate of the managers' average ability parameters, plotted as black line with spades in Figure 6. Firms' average beliefs (grey line) match the true underlying ability fairly well, and are not significantly different (p-value 0.76, WSR test). This supports Hypothesis 2 (d) and suggests that firms' beliefs do imply an accurate assessment of managers' worth on average.

Figure 7 about here.

However, there is also evidence of a winner's curse:¹³ The two highest wage offers exceed average ability (p-value 0.003, WSR test). The same holds for the average accepted wage, plotted as empty circles in Figure 6, and the average of the maximum reported firm beliefs in each observation, plotted as grey dashed line (both p-values < 0.001, WSR test). Thus, the most optimistic firms often make losses, which is also reflected in the analysis of profit data below. Figure 7 provides a sense of the individual heterogeneity in firms' beliefs. The solid line traces the average manager second-order belief in the respective HA-experiment, and the black triangles mark the actual ability levels.

Figure 8 about here.

5.2.3 Hypothesis 2 (e)

Figure 8 shows that managers' second-order beliefs (black line) correspond roughly to average wage offers (grey circles). However, they slightly overestimate wages for higher signal values and $b^M(y) \leq w(y)$ is rejected (p-value 0.09, WSR test). Is managers' play consistent with their stated beliefs? To test Hypothesis 2 (e) we compare for individual subjects their actual average effort levels e with their best response to own second-order beliefs e^* : $e = e^*$ cannot be rejected (p-value 0.44, WSR test). The average best reply against own stated second-order beliefs is 6.08, slightly lower than the average best

¹³A winner's curse results if individuals underappreciate the informational content of their actions in relation to other people's behavior (see Eyster and Rabin, 2005, for an equilibrium concept embedding this) and has been widely documented in other settings (e.g., Thaler, 1988, and for a survey of laboratory experiments Kagel and Levin, 2002).

Table 3: Wage regressions

	Hidden Ability (HA)	Revealed Ability (RA)
constant	1.98* (0.83)	1.72 (1.12)
ability	0.38*** (0.08)	0.58*** (0.09)
effort	0.43*** (0.04)	0.02 (0.05)
R ²	0.15	0.21
N	864	864

Numbers in parentheses are Huber-White robust standard errors, treating each experimental observation as a separate cluster. *, **, *** denote significance at the 10-, 5-, 1-percent levels, respectively.

reply against firms' stated beliefs within the respective managers' experiment of 6.33 (see Table 2). Both are close to the actual effort chosen by subjects, 6.96. However, stated manager beliefs reflect only average expected wage offers, and the relevant maximum wage offers (i.e., the accepted wages) are somewhat higher (see empty circles in Figure 8). Nevertheless, judging by firms' actual wage setting behavior, a manager's expected profit is maximized at effort level 5 against either the average wage offers or the average of firms' maximum wage offers (see Table 2). While actual average effort exceeds 5, this does not affect managers' profits much because profit schedules are relatively flat around the optimum effort level. But managers do respond to differences across firms: regressing effort on what would be the best response to the stated beliefs by the most optimistic firm in that experiment, the coefficient estimate is 1.01 (p-value 0.004, standard errors clustered at the experiment level; $R^2 = 0.07$).

5.3 Wage offers (Hypothesis 3)

Table 3 reports regressions of firms' wage offers on the respective manager's ability and effort levels. In the HA treatment, both effort and ability are significant. Moreover, the hypothesis that their coefficients are equal cannot be rejected ($F(1,5)=0.32$, p-value 0.60), which is in line with the theoretical prediction. In contrast, in the RA treatment effort is not significant and wages react only to ability. Thus, the experimental data support Hypothesis 3.

A more fine-tuned theoretical prediction is that the coefficients on ability/effort (HA) and ability (RA) should be equal to one.¹⁴ This has to be rejected for the HA treatment (ability: $F(1,5)=61.62$,

¹⁴Strictly speaking, Bertrand competition only predicts that this ought to hold for the highest two wage offers. But even then these coefficients are all less than one: the ability coefficients are 0.40*** for HA ($F(1,5)=112.01$, p-value < 0.001) and 0.80*** for RA ($F(1,5)= 8.01$, p-value 0.037); the effort coefficients are 0.47*** for HA ($F(1,5)=53.85$, p-value < 0.001) and -0.04 for RA.

Table 4: Summary statistics for manager and firm profits

	Hidden Ability (HA)	Revealed Ability (RA)
Average profit of managers	10.77	9.89
Standard deviation	8.83	8.58
Average profit of firms	-3.26/-0.79 ^a	-1.32/0.44 ^a
Standard deviation	13.61/6.33 ^a	10.71/1.72 ^a

^a Excluding bankrupt firms and instances where all three wage offers > 19 (see Footnote 15).

p-value < 0.001; effort: $F(1,5)=226.77$, p-value < 0.001). We discussed explanations for this above: managers' decision errors and strategic uncertainty about 'equilibrium' effort levels complicate the inference problem for firms (see Section 5.2). The fact that the RA ability coefficient is significantly lower than unity ($F(1,5)=21.08$, p-value 0.01) partly reflects outlier behavior (see Footnote 15).

Nevertheless, Bertrand competition appears to occur in both treatments. Managers accept the highest wage offers in 96.76 percent (HA) and 96.30 percent (RA) of all cases, respectively, and overall firm profits are close to zero (see below). Moreover, the difference between first- and second-highest wage offers in the RA treatment is just 0.87 on average, excluding rounds where firms post wage offers exceeding the *visible* ability parameter of the manager (and 2.93 for the overall sample, driven mainly by the outliers discussed in Footnote 15).

5.4 Profits (Hypothesis 4)

Table 4 summarizes profits for manager and firm subjects. Contrary to Hypothesis 4, HA managers make a bit more than RA managers, partly reflecting the winner's curse of HA firms. The difference is not significant, however (p-value 0.24, WMW test). Correspondingly, firm profits are lower in the HA treatment (p-value 0.09, WMW test). The standard deviation of firm profits is slightly lower in the RA treatment.

These results, however, are driven by outliers. As already noted in Section 5.1, psychological feedback from firm profits is strong because firms that offer too high wages realize losses. In particular, this is the case in the RA treatment where 'errors' in wage offers, $w(y, a) > a$, quickly disappear (see Figure 4). Firms that bid wages exceeding the *visible* ability parameter a of managers realize an average loss of -16.39. While such losses are sustained by firms only in 27 out of 288 rounds, they nevertheless drive down the overall average firm profit to -1.32.

Figure 9 about here.

Figure 9 plots average profits for firms across rounds excluding outliers.¹⁵ Profits stabilize in both

¹⁵Several bankruptcies occurred: 2 (HA)/ 1 (RA) firm subject used up the initial endowment and received additional

Table 5: Probability of a loss for a firm in round t (HA treatment)

	Logit ^a	Fixed-effect logit
constant	17.13** (0.21)	—
loss in round $t - 1$	0.53*** (0.19)	-0.45** (0.22)
round	-0.13*** (0.03)	-0.17*** (0.03)

^a If round profit in round $t < 0$: loss in round $t=1$, otherwise loss in round $t=0$.

Numbers in parentheses are standard errors. **, *** denote significance at the 5- and 1-percent levels, respectively.

treatments after a few rounds and the variance of profits is lower in the RA treatment. Table 4 shows that, excluding bankruptcy candidates and outlier rounds, average profits are -0.79 (HA) and 0.44 (RA), respectively. The difference is significant (p-value 0.004, WMW test) and reflects the winner's curse for HA firms. Moreover, the standard deviation of RA profits (1.72) is much lower than in the HA treatment (6.33), in line with Hypothesis 4. As Figure 9 shows, in the HA treatment initially firms incur substantial losses on average – manifested also in the winner's curse for firms from average accepted wages in Figure 6. But firms quickly correct mistakes and learn to account for managers' signal jamming, making roughly zero profits on average then. Out of the wage offers that firms made in rounds 1-6, those that were accepted led to losses in 76.64% of the cases (of those offers that were not accepted 33.54% would have been loss making). But in rounds 7-12, firms made a loss on only 36.11% of the accepted wages (and of those offers that were not accepted only 13.89% would have led to losses).

The logit regressions in Table 5 provide further evidence on firm learning. They show that the probability of making a loss decreases with experience. The results nicely capture that some firms are more prone to making losses than others. The logit in the first column shows that a loss in a given period implies a higher chance of a loss in the next period also. The second column controls for individual heterogeneity with subject-level fixed effects in a panel regression. Here, at the individual level, a loss in a round actually makes it less likely that the firm suffers a loss in the next round as

credit to be able to continue to play. Unless explicitly stated, we include these subjects in the analysis to avoid sample selection bias. These subjects clearly misunderstood the experimental setup. The two HA firm subjects offered wages *larger* than a manager's signal y or the maximum possible ability level $a = 19$ in 53 and 47 percent of all their offers, respectively. They thus outbid all other firms and lowered average firm profits to -5.39, compared to an average of -2.20 in the remaining HA observations. The RA firm subject bid more than the *visible* ability parameter a in 44 percent of his offers. This lowered the average firm profits to -4.26, compared to an average of -0.73 in the remaining RA observations. Among the observations for firm subjects who did not suffer bankruptcy there are several outliers: in 4 (HA)/ 5 (HA) rounds firms offered three wages exceeding the maximum ability level $a = 19$, resulting in massive losses in that round.

Table 6: Comparison of results with Irlenbusch and Sliwka (2005a)

	Hidden Ability (HA)		Revealed Ability (RA)	
	<i>our design</i>	<i>IS^a</i>	<i>our design</i>	<i>IS^a</i>
	<i>Average effort</i>			
	6.96	3.56	4.17	5.50
	<i>Coefficients in wage regressions</i>			
ability	0.38***	{ 0.40**	0.58***	0.84***
effort	0.43***		0.02	0.37***

^a Irlenbusch and Sliwka (2006): first period ability/effort and second period wages.

*, **, *** denote significance at the 10-, 5-, 1-percent levels, respectively.

well. So firms in the HA treatment seem to be learning from losses, in line with our discussion for the RA treatment in Section 5.1.

Overall, the results from wage offers and firm profits support the theoretical prediction of Bertrand competition in both treatments, and are in line with other experimental evidence on Bertrand competition with more than two players (e.g., Dufwenberg and Gneezy 2000).

6 Career concerns and gift exchange

Our setup allows directly testing the signal jamming mechanism at the heart of career concerns models. We chose a design that is close to that of Irlenbusch and Sliwka (2006) to be able to relate our findings to theirs. Their setup focuses on the impact of transparency on incentives in the specific context of an experimental labor market as pioneered by Fehr et al. (1993), which leads to a more complex informational environment combining opportunities for both signaling and signal jamming. Specifically, in their design firms set wages both in the beginning of the first and the second periods (as opposed to only in the second period in our design).¹⁶ In both periods managers who accept a wage offer make an effort decision (our managers only choose effort in the first period). The employing firm receives as payoff the sum of the manager's ability parameter and effort in the current period (whereas to our firms only the manager's ability matters). For selfish preferences, the theoretical predictions for effort decisions in both setups are obviously identical.

Table 6 summarizes experimental results. In Irlenbusch and Sliwka first period average effort levels in the HA treatment are lower than those in the RA treatment. This is exactly the opposite of our finding. Differences in the dynamic structure of the two experiments in conjunction with other

¹⁶More subtle differences to our design are that their setup allows transfers proposed by firms to managers to be either positive or negative and restricts them to the integer set $\{-38, \dots, 38\}$.

experimental evidence suggest that these results are due to additional influences that interfere with the signal jamming mechanism in their setup. In contrast to our design, firms in Irlenbusch and Sliwka (2006) can ‘invite’ reciprocal behavior by making generous wage offers; and managers can reward such behavior through higher effort, which translates into more profit for the firm. It is well established from static versions of such experimental labor markets à la Fehr et al. (1993) that they encourage gift exchange in the form of more effort and higher wages than predicted by models based on selfish preferences. This finding is quite robust to the market institutions and their degree of competitiveness.¹⁷

The dynamic structure introduces the possibility of additional influences. In Irlenbusch and Sliwka (2006) subjects play a repeated gift exchange game (with the added feature that there is also learning about a manager’s unknown ability). If subjects perceive the experimental setup as a game of incomplete information about other players’ willingness to reciprocate, actions in the first period serve as signals about a player’s ‘type’. In such games, cooperative behavior can often be sustained for some time because players will engage in behavior in the first rounds to build a reputation for being a ‘cooperative’ type (e.g., Kreps et al. 1982).¹⁸

Indeed, there is experimental evidence for such reputation building behavior in gift exchange games (Gächter and Falk 2002, Irlenbusch and Sliwka 2005). The latter experiment implements a repeated gift exchange model similar to the one in Irlenbusch and Sliwka (2006), replacing the permanent random ability component by a transitory productivity shock to output. In one treatment the firms can observe both profit components separately, while in the other treatment only the sum of effort and the productivity shock is revealed. The experiment shows that players’ reciprocal behavior is stronger when the firm receives a less noisy signal about the effort that a manager provided in response to the accepted wage. This ‘transparency effect’ on signaling about reciprocity is manifested in a significant increase in the wage-effort correlation when transparency on effort is increased.¹⁹

One might therefore expect that in the Irlenbusch and Sliwka (2006) experiment the difference in effort levels across treatments is driven by two countervailing forces: the impact of transparency on career concerns (as witnessed in our experiment) and the transparency effect on signaling about reciprocity. Indeed, the higher average effort level in their RA treatment compared to our design supports this view (see Table 6). Moreover, first period effort has no impact on second-period wages in our design, whereas the coefficient on effort is significant in the wage regression of Irlenbusch and Sliwka (2006).

¹⁷E.g., Fehr, Kirchsteiger, and Riedl (1993, 1998), Fehr, Kirchler, Weichbold, and Gächter (1998), Fehr and Falk (1999), and Brandts and Charness (2004).

¹⁸For a survey of experiments on reputation models and signaling in experiments more generally see Camerer (2003). Kübler et al. (2008) provide an experimental comparison of signaling and screening models.

¹⁹However, average first period effort levels do not change significantly when transparency is increased, only the dispersion of effort levels increases.

This suggests that effort provides a means to signal the willingness to reciprocate in the second period – which is absent in our design.

A notable difference between the experiments, however, is that in our HA treatment average effort is higher. This suggests that the effect from signaling about the degree to which a manager is going to reciprocate in the second round cannot be neatly disentangled from the impact of signal jamming on effort induced by Holmström-type career concerns. As noted above, intransparency has a dampening effect on the use of effort as a signal of willingness to reciprocate in a repeated gift exchange setting. It is possible that this provides players with a means to coordinate on a low effort signal jamming equilibrium in the HA treatment.

In sum, the comparison of the results from the two experiments shows that the findings of Irlenbusch and Sliwka (2006) are overturned once one strips away the gift exchange components and allows only for signal jamming. This is an important insight. Clearly, the interaction of career concerns and gift exchange is relevant in many settings. Nevertheless, even in labor relations future career prospects may be very strongly influenced by the market's perception of an individual's ability – for example, think of the market for CEOs. In such circumstances, there is little reason to expect that gift exchange is an important factor in determining effort. The same is true whenever signal jamming effort does not generate any additional surplus.²⁰ Moreover, gift exchange is not relevant at all in many settings to which the career concerns model has been applied to.²¹

7 Conclusion

Models of signal jamming have become an important part of the toolkit of economic theorists. Despite the importance of this mechanism for theoretical work, empirical evidence has been scant, due to limitations of field data. Our experiment provides a direct test of the signal jamming mechanism that forms the core of Holmström's (1982/99) model. The results support the theoretical premise that economic actors play according to the signal jamming logic. Moreover, subjects' beliefs that were elicited during the experiment are remarkably consistent with each other and with actual behavior, which is a necessary condition for the concept of Bayesian Equilibrium to produce a reasonable approximation of the outcome of individual behavior under incomplete information.

²⁰For example, consider a graduate student who prepares a job application. She may not know her ability (with respect to the tasks ahead) but her resumé provides a signal to potential employers. Polishing the resumé is a costly activity that is likely to enhance the impression that potential employers get while it does not improve the applicant's ability or create any extra value to the future employer.

²¹See the references listed in the introduction.

References

- Andreoni, J., 1995. Cooperation in public-goods experiments: Kindness or confusion? *American Economic Review* 85, 891–904.
- Blanco, M., Engelmann, D., Koch, A.K., Normann, H.T., 2008. Belief elicitation in experiments: Is there a hedging problem? IZA Discussion Paper No. 3517.
- Borland, J., 1992. Career concerns: Incentives and endogenous learning in labour markets. *Journal of Economic Surveys* 6, 251–270.
- Brandts, J., Charness, G., 2004. Do market conditions affect gift exchange? Evidence from experimental markets. *Economic Journal* 114, 684–708.
- Camerer, C.F., 2003. *Behavioral Game Theory*. Princeton: Princeton University Press.
- Dewatripont, M., Jewitt, I., Tirole, J., 1999. The economics of career concerns, part I: Comparing information structures. *Review of Economic Studies* 66, 183–198.
- Dufwenberg, M., Gneezy, U., 2000. Price competition and market concentration: An Experimental study. *International Journal of Industrial Organization* 18, 7–22.
- Eyster, E., Rabin, M., 2005. Cursed equilibrium. *Econometrica* 73, 1623–1672.
- Fee, C.E., Hadlock, C.J., 2003. Raids, rewards and reputations in the market for CEO talent. *Review of Financial Studies* 16, 1315–1357.
- Fehr, E., Falk, A., 1999. Wage Rigidities in a competitive incomplete contract market. An experimental investigation. *Journal of Political Economy* 107, 106–134.
- Fehr, E., Kirchler, E., Weichbold, A., Gächter, S., 1998. When social norms overpower competition: Gift exchange in experimental labor markets. *Journal of Labor Economics* 16, 324–351.
- Fehr, E., Kirchsteiger, G., Riedl, A., 1993. Does fairness prevent market clearing? An experimental investigation. *Quarterly Journal of Economics* 108, 437–460.
- Fehr, E., Kirchsteiger, G., Riedl, A., 1998. Gift exchange and reciprocity in competitive experimental markets. *European Economic Review* 42, 1–34.
- Fischbacher, U., 2007. Z-Tree - Zurich Toolbox for Ready-Made Economic Experiments. *Experimental Economics* 10, 171–178.

- Fudenberg, D., Tirole, J., 1986. A “signal-jamming” theory of predation. *RAND Journal of Economics* 17, 366–376.
- Fudenberg, D., Tirole, J., 1991. *Game Theory*. Cambridge: MIT Press.
- Gächter, S., Falk, A., 2002. Reputation and reciprocity: Consequences for the labour relation. *Scandinavian Journal of Economics* 104, 1–27.
- Gibbons, R., Murphy, K.J., 1992. Optimal incentive contracts in the presence of career concerns: Theory and evidence. *Journal of Political Economy* 100, 468–505.
- Holmström, B., 1982/99. Managerial incentive problems: A dynamic perspective. *Review of Economic Studies* 66, 169–182. Originally published in: *Essays in Economics and Management in Honour of Lars Wahlbeck*, Helsinki.
- Irlenbusch, B., Sliwka, D., 2005. Transparency and reciprocal behavior in employment relations. *Journal of Economic Behavior and Organizations* 56, 383–403.
- Irlenbusch, B., Sliwka, D., 2006. Career concerns in a simple experimental labour market. *European Economic Review* 50, 147–170.
- Kagel, J.H., Levin, D., 2002. *Common Value Auctions and the Winner’s Curse*. Princeton: Princeton University Press.
- Kreps, D., Milgrom, P., Roberts, J., Wilson, R., 1982. Rational cooperation in the finitely repeated prisoner’s dilemma. *Journal of Economic Theory* 27, 245–252.
- Kübler, D., Normann, H.T., Müller, W., 2008. Job market signalling and screening in laboratory experiments. *Games and Economic Behavior* 64, 219–236.
- LeBorgne, E., Lockwood, B., 2006. Do elections always motivate incumbents? Experimentation vs. career concerns. *Public Choice* 129, 41–60.
- Lohmann, S., 1998. Rationalizing the political business cycle: A workhorse model. *Economics and Politics* 10, 1–17.
- Lohmann, S., 1999. What price accountability? The Lucas island model and the politics of monetary policy. *American Journal of Political Science* 43, 396–430.
- Samuelson, L., Mirman, L.J., Schlee, E.E., 1994. Strategic information transmission in duopolies. *Journal of Economic Theory* 62, 363–384.

Stein, J., 1989. Efficient capital markets, inefficient firms: A model of myopic corporate behavior. *Quarterly Journal of Economics* 104, 655–669.

Thaler, R.H., 1988. Anomalies: The winner's curse. *Journal of Economic Perspectives* 2(1), 191–202.

Accepted Manuscript

Appendix

A Cost function

Table 7: Cost function for effort

	Effort (e)									
	0	1	2	3	4	5	6	7	8	9
c(e)	0	0.26	0.57	0.93	1.34	1.8	2.33	2.92	3.57	4.29
	Effort (e)									
	10	11	12	13	14	15	16	17	18	19
c(e)	5.08	5.93	6.84	7.94	9.23	10.7	12.35	14.19	16.22	18.43

B Derivation of equilibria

In this section we show how to construct Perfect Bayesian Equilibria with a pure strategy equilibrium effort level $e^* \in \{0, 1, \dots, 12\}$. Let $b(y|e^*)$ denote firms' beliefs about the manager's ability a , given that signal y has been observed and that the firms expect the manager to choose effort e^* . Table 8 provides examples of such beliefs supporting pure strategy equilibria. To illustrate the construction of an equilibrium, let us focus on the candidate equilibrium $e^* = 2$ (column four in Table 8). Suppose that a signal $y = 10$ realizes. This is on the (proposed) equilibrium path and thus $b(10|2) = 10 - 2 = 8$. In contrast, $y = 1$ would be off the equilibrium path and beliefs only need to satisfy restrictions (C 1) and (C 2), as $b(1|2) = 0$ does for instance.

Given that firms hold beliefs $b(y|e^*)$, a manager chooses effort e to maximize the expected wage net of effort costs, $E_y[b(y|e^*)|e] - c(e)$, using the fact that a is uniformly distributed on $\{0, 1, \dots, 19\}$. In our example with candidate equilibrium $e^* = 2$, these net expected payoffs are given in column four in Table 9. The net expected payoff increases when moving from $e = 0$ to $e = 2$ and then always decreases, confirming that a manager will actually choose the conjectured equilibrium effort level $e^* = 2$. The equilibrium with $e^* = 12$ corresponds to the unique equilibrium effort level if beliefs are restricted by assuming that the monotonicity condition (MR) holds. It is easy to verify that it is not possible to construct equilibria with effort levels $e^* > 12$.

Table 8: Examples of firm beliefs supporting pure strategy equilibria

y = a + e	Equilibrium effort level (e*)												
	0	1	2	3	4	5	6	7	8	9	10	11	12 ^a
	Firm beliefs b(y e*) ^b												
0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	1	0	0	0	0	0	0	0	0	0	0	0	0
2	2	1	0	0	0	0	0	0	0	0	0	0	0
3	3	2	1	0	0	0	0	0	0	0	0	0	0
4	4	3	2	1	0	0	0	0	0	0	0	0	0
5	5	4	3	2	1	0	0	0	0	0	0	0	0
6	6	5	4	3	2	1	0	0	0	0	0	0	0
7	7	6	5	4	3	2	1	0	0	0	0	0	0
8	8	7	6	5	4	3	2	1	0	0	0	0	0
9	9	8	7	6	5	4	3	2	1	0	0	0	0
10	10	9	8	7	6	5	4	3	2	1	0	0	0
11	11	10	9	8	7	6	5	4	3	2	1	0	0
12	12	11	10	9	8	7	6	5	4	3	2	1	0
13	13	12	11	10	9	8	7	6	5	4	3	2	1
14	14	13	12	11	10	9	8	7	6	5	4	3	2
15	15	14	13	12	11	10	9	8	7	6	5	4	3
16	16	15	14	13	12	11	10	9	8	7	6	5	4
17	17	16	15	14	13	12	11	10	9	8	7	6	5
18	18	17	16	15	14	13	12	11	10	9	8	7	6
19	19	18	17	16	15	14	13	12	11	10	9	8	7
20	5	19	18	17	16	15	14	13	12	11	10	9	8
21	6	6	19	18	17	16	15	14	13	12	11	10	9
22	7	7	7	19	18	17	16	15	14	13	12	11	10
23	8	8	8	8	19	18	17	16	15	14	13	12	11
24	9	9	9	9	9	19	18	17	16	15	14	13	12
25	10	10	10	10	10	10	19	18	17	16	15	14	13
26	11	11	11	11	11	11	11	19	18	17	16	15	14
27	12	12	12	12	12	12	12	12	19	18	17	16	15
28	13	13	13	13	13	13	13	13	14	19	18	17	16
29	14	14	14	14	14	14	14	14	15	15	19	18	17
30	15	15	15	15	15	15	15	15	16	16	16	19	18
31	16	16	16	16	16	16	16	16	17	17	17	18	19
32	17	17	17	17	17	17	17	17	18	18	18	19	19
33	18	18	18	18	18	18	18	18	19	19	19	19	19
34	19	19	19	19	19	19	19	19	19	19	19	19	19
35	19	19	19	19	19	19	19	19	19	19	19	19	19
36	19	19	19	19	19	19	19	19	19	19	19	19	19
37	19	19	19	19	19	19	19	19	19	19	19	19	19
38	19	19	19	19	19	19	19	19	19	19	19	19	19

^a Satisfy the monotonicity condition (MR). ^b Beliefs off the equilibrium path are written in italics.

Table 9: Manager's effort decision facing equilibrium beliefs $b(y|e^*)$ from Table 8

e	Equilibrium effort level e^*												
	0	1	2	3	4	5	6	7	8	9	10	11	12
	Expected wage net of effort costs ^a $E_y[b(y e^*) e] - c(e)$												
0	9.50	8.55	7.65	6.80	6.00	5.25	4.55	3.90	3.30	2.75	2.25	1.80	1.40
1	9.49	9.24	8.29	7.39	6.54	5.74	4.99	4.29	3.64	3.04	2.49	1.99	1.54
2	9.43	9.23	8.93	7.98	7.08	6.23	5.43	4.68	3.98	3.33	2.73	2.18	1.68
3	9.32	9.17	8.92	8.57	7.62	6.72	5.87	5.07	4.32	3.62	2.97	2.37	1.82
4	9.16	9.06	8.86	8.56	8.16	7.21	6.31	5.46	4.66	3.91	3.21	2.56	1.96
5	8.95	8.90	8.75	8.50	8.15	7.70	6.75	5.85	5.00	4.20	3.45	2.75	2.10
6	8.67	8.67	8.57	8.37	8.07	7.67	7.17	6.22	5.32	4.47	3.67	2.92	2.22
7	8.33	8.38	8.33	8.18	7.93	7.58	7.13	6.58	5.63	4.73	3.88	3.08	2.33
8	7.93	8.03	8.03	7.93	7.73	7.43	7.03	6.53	5.93	4.98	4.08	3.23	2.43
9	7.46	7.61	7.66	7.61	7.46	7.21	6.86	6.41	5.91	5.21	4.26	3.36	2.51
10	6.92	7.12	7.22	7.22	7.12	6.92	6.62	6.22	5.82	5.17	4.42	3.47	2.57
11	6.32	6.57	6.72	6.77	6.72	6.57	6.32	5.97	5.67	5.07	4.37	3.57	2.62
12	5.66	5.96	6.16	6.26	6.26	6.16	5.96	5.66	5.46	4.91	4.26	3.56	2.66
13	4.81	5.16	5.41	5.56	5.61	5.56	5.41	5.16	5.06	4.56	3.96	3.36	2.51
14	3.77	4.17	4.47	4.67	4.77	4.77	4.67	4.47	4.47	4.02	3.47	2.92	2.12
15	2.55	3.00	3.35	3.60	3.75	3.80	3.75	3.60	3.65	3.25	2.75	2.25	1.50
16	1.10	1.60	2.00	2.30	2.50	2.60	2.60	2.50	2.60	2.25	1.80	1.35	0.65
17	-0.59	-0.04	0.41	0.76	1.01	1.16	1.21	1.16	1.31	1.01	0.61	0.21	-0.44
18	-2.52	-1.92	-1.42	-1.02	-0.72	-0.52	-0.42	-0.42	-0.22	-0.47	-0.82	-1.17	-1.77
19	-4.68	-4.03	-3.48	-3.03	-2.68	-2.43	-2.28	-2.23	-1.98	-2.18	-2.48	-2.78	-3.33

^a The maximum value of the expected wage net of effort costs is written in bold.

C Instructions

Welcome to the experiment! In this experiment you can earn money. The amount of your earnings depends on your decisions and the decisions of the other participants. All decisions will be taken anonymously, i.e., nobody gets to know what decisions you have taken at any point. You will be paid at the end of the experiment. Payments are done one by one so that nobody gets to know the earnings of other participants. These instructions are the same for all participants. From now on, please refrain from talking to other participants and ask questions only to the experimenters directly.

ROUNDS AND TYPES OF PARTICIPANTS.

- The experiment lasts for 12 rounds.
- In the beginning of the experiment, you will be randomly assigned to one of two groups which consist of seven members each. You will belong to that group during the whole experiment. Neither during the experiment nor after, will any participant get to know who else belonged to her group.
- In each group, there are two types of participants – A-participants and B-participants. There are 3 A-participants and 4 B-participants in each of the groups. In the beginning of the experiment, your role will be determined randomly and communicated to you. During the experiment, the roles of the participants remain unchanged.

THE BASE VALUE OF A-PARTICIPANTS. In the beginning of each round, each A-participant will be assigned a base value. This base value is drawn randomly from the numbers 0, 1, 2, 3, ... to 19, where each number is equally likely to occur. This base value remains the same throughout a round and will not be communicated to any participant (not even the A-participant herself). [RA-treatment: This base value remains the same throughout a round and will not be communicated to A-participants.] Note that each A-participant receives a new base value at the beginning of each new round.

SEQUENCE OF EVENTS DURING A ROUND.

- A-participants
 - A new base value from 0 to 19 is randomly assigned to each A-participant. This base value cannot be observed by any participant.
 - Each A-participant chooses a number 0, 1, 2, 3, ... to 19. Each of these numbers is associated with a cost to the A-participant (see cost sheet).
 - The base value and the chosen number by the A-participant determine the so-called result for the A-participant. This result equals: base value + chosen number.

- B-participants

- The results [RA-treatment: the results and the base values] of the 3 A-participants in a group are displayed to the B-participants in this group.

Note: The order of appearance of results [RA-treatment: and base values] on the screen is changed randomly in every round. It is therefore impossible to use the ranking of the results to infer the identity of an A-participant.

- Each B-participant makes a transfer offer for each A-participant. This transfer offer has to lie between 0.00 and 38.00.

- A-participants

- Each A-participant observes her chosen number, her result, and her transfer offers.

Note: The order of appearance of transfer offers on the screen is changed randomly in every round. It is therefore impossible to use the ranking of the transfer offers to infer the identity of a B-participant.

- Each A-participant accepts one (or no) transfer offer.

- The payoffs per round are calculated and displayed.

END OF THE ROUND. At the end of the round the payoffs for the round are displayed. They are calculated as follows.

- Payoff per round for A-participants.

- If one transfer offer was accepted: payoff for this round = transfer offer – cost for chosen number.
- If no transfer offer was accepted: payoff for this round = 0 – cost for chosen number.

- Payoff per round for B-participants.

- If one transfer offer of the B-participant was accepted: payoff for this round = base value of the accepting A-participant – transfer offer to this A-participant
- If several transfer offers of the B-participant were accepted: payoff for this round = sum of the base values of the accepting A-participants – sum of the transfer offers to these A-participants.
- If no transfer offer of the B-participant was accepted, the payoff for this round is 0.

At the end of the experiment, the sum of the 12 round payoffs and an initial endowment will be exchanged into Euros (Bonn)/Pounds Sterling (London) (at a rate of 5 Euro-Cent/pence per experimental currency unit) and paid out to you.

Please note: during the experiment, you will be asked to answer a few questions on the screen. Your answers to these questions will have no influence on your payments.

Feedback about round:

Managers: Signal $[y_m]$, effort $[e_m]$, cost of effort $[c(e_m)]$; accepted wage offer (w_m), or if all rejected $[w_m=0]$; resulting round profit $[w_m - c(e_m)]$.
Firms: For each accepted wage offer: signal $[y_m]$ and ability $[a_m]$ of the manager, accepted wage offer $[w_m]$, resulting round profit from manager $[a_m - w_m]$; total round profit. If all three wage offers were not accepted the total round profit displayed is 0.

Figure 1: Stages of play in each of the 12 rounds in a session

Figure 2: Average effort of managers across all rounds

Figure 3: Effort choices of managers over all rounds

Figure 4: 'Errors' in manager and firm decisions (RA treatment)

Figure 5: Averages of firms' beliefs and managers' second-order beliefs (HA treatment)

Figure 6: Firms' beliefs and actual play (HA treatment)

Figure 7: Firms' beliefs (individual subjects in each observation)

Figure 8: Managers' second-order beliefs and actual play (HA treatment)

(a) Profits of firms (HA treatment) – excluding outliers^a

(b) Profits of firms (RA treatment) – excluding outliers^a

^aExcluding bankrupt firms and instances where all three wage offers > 19 (see Footnote 15).

Figure 9: Profits of managers and firms across all rounds (excluding outliers)