CESifo WORKING PAPERS

7325
2018

November 2018

Self-Confidence and Reactions to Subjective Performance Evaluations

Charles Bellemare, Alexander Sebald

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl
www.cesifo-group.org/wp

An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: www.CESifo-group.org/wp

CESifo *Working Paper No. 7325*
Category 13: Behavioural Economics

Self-Confidence and Reactions to Subjective Performance Evaluations

Abstract

Subjective performance evaluations are commonly used to provide feedback and incentives to workers. However, such evaluations can generate significant disagreements and conflicts, the severity of which may be driven by many factors. In this paper we show that a workers' level of self-confidence plays a central role in shaping reactions to subjective evaluations - overconfident agents engage in costly punishment when they receive evaluations below their own, but provide limited rewards to principals when evaluations exceed their own. In contrast, underconfident agents do not significantly react to evaluations below their own, but reward significantly evaluations exceeding their own. Our analysis exploits data from a principal-agent experiment run with a large sample of the Danish working age population, varying the financial consequences associated with the evaluations workers receive. In contrast to existing economic models of reciprocal behavior, reactions to evaluations are weakly related to the financial consequences of the evaluations. These results point towards a behavioral model of reciprocity that intertwines the desire to protect self-perceptions with over-/underconfidence.

JEL-Codes: D010, D020, D820, D860, J410.

Keywords: subjective performance evaluations, self-confidence, reciprocity.

Charles Bellemare
Department of Economics
Université Laval / Québec / Canada
cbellemare@ecn.ulaval.ca

Alexander Sebald
Department of Economics
University of Copenhagen /Denmark
alexander.sebald@econ.ku.dk

October 17, 2018
The authors thank audiences of seminars at the universities of Cologne, Lyon, Maastricht, Pittsburgh, Michigan, Copenhagen and Mainz. Bellemare gratefully acknowledges the financial support from the Social Sciences and Humanities Research Council of Canada (Grant ID: 435-2013-1715). Sebald gratefully acknowledges the financial support from the Danish Council for Independent Research in Social Sciences (Grant ID: FSE-10-081368 and DFF-4003-00032). All errors are our own.

1 Introduction

Subjective performance evaluations are frequently used to provide feedback and incentives when objective measures of performance are costly, imperfect or lacking [see e.g. Baker *et al.* (1994), Prendergast (1999) and Gibbs *et al.* (2004)]. Because of their subjective nature, such evaluations have the potential to generate severe disagreements and conflicts in the workplace as opinions of principals and agents may differ for many reasons. Agents assessing their own performance may for example be prone to various psychological biases (e.g. self-serving [e.g. Miller *et al.* (1975) and Gervais and Odean (2001)], recency [e.g. Hogarth and Einhorn (1992)], confirmation [e.g. Rabin and Schrag (1999)], conservatism [e.g. Mobius *et al.* (2011)]).[1] These biases can lead to self-evaluations displaying over- or underconfidence about own levels of ability and performance.[2]

Biased self-perceptions and overconfidence have received a lot of attention in recent years [e.g. Hoelzl and Rustichini (2005), Burks *et al.* (2013), Malmendier and Tate (2015), Hoffman and Burks (2017)]. Many social psychologists and economists see overconfidence as a consequence of people's natural desire for self-enhancement and for their ability and performance to be positively acknowledged by others [see e.g. Baumeister and Vohs (2001), Baumeister and Finkel (2010, chapter 5), Bénabou and Tirole (2002), Köszegi (2006), Ellingsen and Johannesson (2008), Burks *et al.* (2013) and Schwardmann and Van der Weele (2016)].

Interestingly, the social psychological literature has suggested an intriguing link between overconfidence and the proneness to engage in conflicts hitherto ignored in the economic literature. Narcissistic personalities characterized by inflated self-views have been found to be prone to counterproductive behavior in the workplace in response to ego-threatening information [e.g. Campbell *et al.* (2011), Baumeister *et al.* (2000) and Campbell and Miller (2011, chapter 28)]. Many economic models predict agents would retaliate to conflicting evaluations when the later have financial consequences (e.g. Dufwenberg and Kirchsteiger (2004), Falk and Fischbacher (2006), Sebald (2010)). Absent such material consequences, models predicting agents would retaliate at a personal cost are limited, and none explicitly allows a channel through which reciprocal behavior is determined by the agents' self-confidence.

[1] For a related discussion regarding the impact of principals' or supervisors' evaluation biases see e.g. Breuer *et al.* (2013)

[2] The term overconfidence has been used to designate two other related concepts ([Moore and Healy (2008)]). First, 'better-than-average' effect leads people to believe they perform better than others. Second, miscalibration leads people to believe they have more precise knowledge.

In this paper we show that self-confidence plays a significant role in shaping how agents react to evaluations about their performance. Overconfident agents are found to significantly punish evaluations falling below their own, and refrain from rewarding principals for evaluations exceeding their own. Conversely, underconfident agents refrain from punishing principals for evaluations falling below their own, but reward principals for feedback exceeding their own. We show that these relationships are present irrespective of whether evaluations have financial consequences for agents. It follows from our results that monetary consequences, while having some significant effects on behavior, play a secondary role in shaping how agents' react to subjective evaluations.

Our analysis exploits data from an online experiment run with a large sample of the Danish working age population based on a simple principal-agent setting with subjective performance evaluations inspired by MacLeod (2003). In our main treatments agents decide how much effort to spend on a real-effort task in which their individual performance can only be evaluated subjectively by them and their matched principal. Agents subsequently receive an unverifiable performance feedback from their matched principal before deciding whether to reward or punish their principal at a personnel cost. Our two main treatments vary whether evaluations provided by principals have financial consequences for agents. While agents and principals do not observe the objective performance of the agent during the experiment, the later is observable in the data. We contrast objective performance data with the subjective evaluations agents provide about their own performance to identify over- and underconfident agents. Our results are generated by estimating a structural model using our data. Reward and sanction decisions are modelled by specifying flexible preferences for agents over their own payoff and that of their matched principal. These preferences are allowed to be correlated with over-/underconfidence measured for each agent, and vary with the level gap between evaluations of agents and principals. The estimated model is used to predict reactions (punish or reward) of agents with various levels of self-confidence across different gaps between evaluations of agents and principals, taking into account noise and measurement errors present in the data.

Our results point towards a behavioral model of reciprocity where self-confidence plays a key role in shaping responses to ego-threatening information even in the absence of financial consequences. Generating conflict in response to ego-threatening information is often cited as an interpersonal mechanism used to protect their self-perceptions [see e.g. Baumeister and Finkel (2010, chapter 5) for a discussion].[3] Our findings are consistent

[3]In contrast to this the existing economic literature has focused on prominent 'intrapersonal' strategies to maintain/protect positive self-views including staying ignorant about ones own abilities, or embarking

with such mechanisms explaining why agents (irrespective of their level of self-confidence) respond to evaluations about their performance which contain an ego-threatening component. Establishing a relation between self-confidence and reciprocity provides new insights. To our knowledge, self-confidence has not been formally incorporated as a central component of existing models of reciprocity. This has important new implications for management practices and the usefulness of subjective evaluations. Broadly speaking, we find that associating financial outcomes to these evaluations matters less than profiling workers to determine who should be provided subjective evaluations. In the limit, an implication of our findings is that subjective evaluations should play a limited role (or simply avoided) in work environments dominated by overconfident agents who respond in an unproductive way to negative assessments differing from their own.

In the next section we describe and explain the design of our experiment. In section 3 we present the data and some descriptive statistics. Sections 4 and 5 respectively explain the model and results. Section 6 concludes.

2 Experiment

The experiment consisted of 4 different parts.[4]. Each part captured a principal-agent employment relationship where an agent worked on a task generating value for a matched principal. We first discuss the common elements across all parts before discussing the part specific variations. After describing the design of the experiment we explain how the experiment was implemented online.

2.1 Design elements common to all parts

The clicking task. Each part featured a clicking task which the agent could work on and which generated payoff for the principal. The task consisted of a series of screens on which boxes were presented. The agent could click these boxes away. After a few seconds each screen disappeared with the remaining boxes and was replaced by a new screen (see Figure 1 for an example of a clicking screen). A total of 20 screens were presented to each agent in each part of the experiment, with the total number of boxes to be clicked across all screens in each part set to 400. The positioning of the boxes varied across screens, making it difficult for agents to replicate past clicking patterns. To ensure

on overly ambitious tasks on which one is sure to fail in order to protect ones self-image and confidence [see Bénabou and Tirole (2002)] or self-deception Schwardmann and Van der Weele (2016).

[4] There was a voluntary 5th part conducted a year later which is not part of the analysis here.

heterogeneity in people's evaluations that is orthogonal to ability, we randomly allocated each agent to one of three different clicking time environments, denoted main, fast, and slow. Screens in the main clicking time environment appeared between 3 and 10 seconds before being replaced.[5] Screens of agents assigned to the fast clicking time environment appeared between 2 and 7 seconds, while those of agents assigned to the slow clicking time environment appeared between 4 and 13 seconds. A video of the agent's performance on the task was recorded and shown to the principal for each part.

Importantly, screens appeared only a few seconds, making it difficult for agents to click all boxes on a given screen. Furthermore, agents and principals were unaware that there were 400 possible boxes to click in each part. What is more, both parts, i.e. principals and agents, did not have enough time to evaluate the total number of boxes which appeared on each screen. These design features served the purpose of limiting as much as possible the capacity of agents and principals to infer precisely the objective share of boxes the agent was able to click throughout a given part of the experiment.[6] For these reasons, assessments of productivity in our design focus on the percentage of boxes clicked away in a given part by a given agent.

Payoffs of principals. As mentioned above the agents generated value for their principal through their work on the clicking task. Specifically, payoffs of principals in all four parts of the experiment were based on the productivity of agents using the structure presented below:

If the agent clicked away...

0-10% of the boxes, the principal received 325 DKK
10-20% of the boxes, the principal received 350 DKK,
20-30% of the boxes, the principal received 375 DKK,
30-40% of the boxes, the principal received 400 DKK,
40-50% of the boxes, the principal received 425 DKK,
50-60% of the boxes, the principal received 450 DKK,
60-70% of the boxes, the principal received 475 DKK,
70-80% of the boxes, the principal received 500 DKK,

[5]In this environment, the exact time each screen appeared was drawn randomly from a uniform distribution on the set of integers $\{3, 4, ..., 10\}$. The same procedure was used for the other two time environments.

[6]Agents who clicked all boxes of all screens in a given part would know their objective performance. Section 3 documents that very few agents were able to accomplish this.

80-90% of the boxes, the principal received 525 DKK,
90-100% of the boxes, the principal received 550 DKK.

Performance evaluations and feedback. After respectively working and observing the agent's work on the task, the agent and the principal had to privately evaluate the performance of the agent. Both players made these assessments by selecting 1 of 10 possible productivity intervals for the agent. Specifically, both players had to indicate whether they thought the agent clicked away 0-10% of the boxes, 10-20% of the boxes, ... , or 90-100% of all boxes in a given part (see Figure 2 for a translated own evaluation screen of the agent. The principal's screen was exactly the same.). Additionally, the principal had to give a performance feedback to the agent by selecting 1 of the same 10 productivity intervals. Importantly the private evaluations of the principal and the principal's feedback to the agent were independent of each other and could differ.

2.2 Design elements varying across parts

Payoffs of agents. The agent's payoff was in all parts paid from the principal's profit generated through the agent's work on the clicking task. Agents in Parts 2 and 4 received a fixed wage of 150 DKK independent of their performance and the feedback of the principal. The agent's payoff in Part 3 was increasing in the subjective feedback off the principal using the following payoff structure:

If the principal's feedback is that the agent clicked away...

0-10% of the boxes, than the agent received 110 DKK
10-20% of the boxes, than the agent received 120 DKK,
20-30% of the boxes, than the agent received 130 DKK,
30-40% of the boxes, than the agent received 140 DKK,
40-50% of the boxes, than the agent received 150 DKK,
50-60% of the boxes, than the agent received 160 DKK,
60-70% of the boxes, than the agent received 170 DKK,
70-80% of the boxes, than the agent received 180 DKK,
80-90% of the boxes, than the agent received 190 DKK,
90-100% of the boxes, than the agent received 200 DKK.

Agents in Part 1 faced the same payoff scheme with the difference that the payment to the agent was not contingent on the principal's subjective feedback but the agent's actual performance. That is, if the agent actually clicked away 0-10% of the boxes, then the agent received 110 DKK, if the agent clicked away 10-20% of the boxes, then the agent received 120 DKK.

Reward/punishment. Parts also differed in terms of the possibility to reward or punish principals for their feedback. In contrast, agents in Parts 1 and 2 could not reward or punish principals for possible feedbacks. Agents in Parts 3 and 4 were given the possibility to reward or punish principals for different possible feedbacks they could receive. This was implemented using a strategy method approach – agents had the possibility to costly increase or decrease the payoff of the principal (between -100 and 100 DKK) for each possible feedback. Each DKK of increase or decrease cost the agent 0.25 DKK, an amount deducted from his/her payoff. Figure 3 presents a translated version of the reaction screen).

2.3 Conduct of the experiment

The experiment was conducted online by the Center of Experimental Economics at the University of Copenhagen during a period of 6 weeks during the summer of 2012. 18000 people were invited to participate with the help of Statistics Denmark - the statistical office of Denmark. People were invited by regular mail containing an invitation letter to participate in the online study. The letter invitation contained some explanations and stated a link as well as a password and login which could be used to participate (see Figure 8 for a sample invitation letter). Invitees were given 6 weeks to complete the experiment. As we invited a large and representative sample of the Danish working-age population the experiment was conducted in Danish.

Upon logging on participants were told that the experiment consisted of 4 different parts and that one part was randomly selected ex-post to be paid out. They were given explanations to Part 1 and some control questions to make sure that the strategic setting was clearly understood. Note that the explanations were kept neutral with the principal and the agent respectively being labeled Person A and Person B. Participants were also told that their role would be revealed after successfully answering the control questions. Lastly, participants were told that instructions to Parts 2, 3 and 4 would follow later on during the experiment. Participants did not receive any feedback between parts and outcomes of the experiment could be viewed 8 weeks later by logging in the same web site.[7] Role assignment was implemented using an alternating assignment scheme to limit

[7]Neither agents nor principals received information about the agent's true performance, feedback,

the number of unmatched players. The first person that answered the control questions correctly was assigned the role of an agent. The next person that logged on and answered the control questions correctly was either assigned to the role of the principal in case there was an agent that had already completed the experiment or the role of the agent. Agents and the matched principal's were matched for the entire experiment but no feedback was provided between parts. As no feedback was given between parts, our design meant that agents went through the 4 different parts before being able to observe the feedback of principals as well as their true objective productivity in each part of the experiment. The strategy method thus enabled us to elicit the complete reaction strategy of the agents while limiting spillovers across parts. Moreover, the design ensured that objective information about the productivity of agents was never provided before the whole experiment was completed.

Principals went through the 4 different parts sequentially. This meant first viewing a video of the clicking task accomplished by their matched agent in each part. They were then asked to state their private evaluations, followed by the feedback they intend to pass to agents after the end of the experiment. As outlined above, this feedback determined the agents' payoff only in Part 3.

A total of 2155 participants successfully answered the control questions, 1693 of which finished the experiment. Furthermore, due to a technical problem 27 participants were either not assigned to a partner or wrongly assigned to 2 partners. We do not consider the data from these players. The data analysis in the remainder of this paper is based on the data from 1666 participants that completed the entire experiment, i.e. 833 matched pairs consisting of one agent and one principal.

3 Data and Descriptives

We measure individual productivity using the proportion of boxes clicked by an agent. Figure 4 presents the distribution of individual productivity across Parts 1 and 2 of the experiment for each of the three clicking time environments that were randomly assigned across players.[8] The left hand side graph presents the distribution for players assigned to the fast clicking time environment. The right-hand graph presents the distribution

and payoffs in each part during the experiment. All decisions in the experiment were thus based only on the subjective evaluations of agents and principals.

[8]Within agent correlation in total number of clicks between Parts 1 and 2 is very high (0.905). Separate graphs for both parts are thus omitted.

for players assigned to the slow clicking time environment. The middle graph plots productivity for the main clicking time environment. All three graphs reveal considerable heterogeneity of clicking ability. As expected, we find that productivity tends to increase with the time given to click away boxes. The econometric analysis in Section 4 will exploit the randomisation of clicking times across agents to identify the causal effect of productivity on the agents' own evaluations.

Figure 5 presents the distributions of agents' own evaluations for Parts 3 and 4 of the experiment. We find that both distributions are very similar and concentrated near intervals [30 – 40] and [40 – 50]. The similarity of both distributions reflects the limited within agent variation in own evaluations across Parts 3 and 4. We find that 47.9% of agents used the same own evaluation category for both parts, while 42,5% of agents used adjacent categories. The left histogram in Figure 6 plots the differences between evaluations of matched principal-agent pairs for Parts 3 and 4. We find that differences are concentrated around zero in both parts of the experiment, suggesting that principals and agents tend to have similar evaluations of the productivity of agents. Our experimental design enforced that a principal's evaluation was private information that was not passed on to the agent. Each agent on the other hand received feedback from the principal that could differ from the private evaluation made by the principal. This provided each principal with the opportunity to hide their private assessment of the productivity of their agent in order to avoid conflict and possible sanctions in Parts 3 and 4. At the same time, providing generous feedback may induce reciprocal agents to reward principals at a personal cost to them.

The right histogram in Figure 6 plots the differences between private evaluations and feedback of each principal. We find that distributions in both parts are significantly concentrated near 0, an indication that principals on average passed on their private evaluations as feedback to agents. Note that this pattern holds both when feedback was costly (Part 3) or not (Part 4) to provide. Cost-minimization induces selfish principals to provide the lowest possible feedback in Part 3. This incentive is especially salient in this experiment as agents made reward/sanctions decisions contingent on all possible feedbacks they could receive rather than after observing the feedback provided by the principal. Evidence from Figure 6 thus suggests that principals wanted outcomes of the experiment to be settled using their private information.

Table 1 presents descriptive evidence on evaluations of both players as a function of the objective productivity of agents. The later – computed using the share of boxes clicked in a part – is divided across 8 intervals chosen to match the evaluation intervals used by agents and principals in the experiment. The proportion of evaluations that coincide, exceed, or fall below the objective productivity interval are presented. An evaluation is accurate if it correctly identifies the objective productivity interval associated with the clicking task that was accomplished. We find that the proportion of accurate evaluations made by agents who evaluated their own performance ranges from 28.9% to 46.8% across the different intervals. Moreover, evaluations that fall below an objective interval are more prominent than evaluations that exceed these intervals, suggesting agents are more underconfident than overconfident in the clicking task. Finally, average own evaluations within each objective interval are presented. We find a general positive relationship between objective productivity and own evaluations. This relationship also depicts the relatively stronger underconfidence present in the data – excluding boundary intervals, average evaluations are below the mid-points of each interval, and in some cases below objective lower bounds. Average evaluations were separately computed for men and women agents. We find no significant differences across the intervals considered (p-values > 0.05, individual t tests), suggesting limited differences in confidence across gender. Evaluations of principals are presented in the bottom of Table 1. Principals tend to make more accurate evaluations than agents – the proportions of evaluations that coincide with the objective productivity intervals are higher for all intervals considered. This higher accuracy is a likely reflection of the fact that principals' main task was to observe the work of agents.

An interesting observation concerning the agents' own evaluations is that own-evaluation bias is persistent across the four parts of our experiment. We find that 86 % of all agents are either over or underconfident in at least three out of four parts. More specifically, 46% of agents (i.e. 384 out of 833) are overconfident in at least three out of four parts, 40% of agents (i.e. 335 out of 833) are underconfident in at least three out of four parts. This suggests that over- and underconfidence measured by the distance between individual assessments and objective performance captures a stable characteristic rather than random deviations of assessments and objective productivity.

The main focus of this paper is on the behavior of agents in response to the feedback received. Let rew capture the reward/sanction behavior of agents, where rew ranges from -100 (maximal sanction in DKK currency) to 100 (maximal reward in DKK currency). Furthermore, let $feed$ denote the principals' feedback of the performance of the agent, and own denote the agents' own evaluation of his/her performance. Table 2 breaks down average values of rew for different ranges of $feed - own$. To simplify, values of $feed - own$ are expressed using the lower bound of each subjective performance interval players could select. Results are similar when the upper bound of each interval is used. Agents in

Part 3 sanction principals on average 25.042 points when feedback falls below the agents evaluation by at least 50 percentage points. Sanctioning behavior tends to decrease as feedback increases, given *own*. Agents reward principals with an average of 15.834 points when feedback and evaluations coincide. Interestingly, average rewards do not increase as the feedback exceeds the agents own evaluation. In fact, we observe a small decline in average reward behavior when feedbacks exceed evaluations of agents by at least 50 percentage points. Removing the monetary consequences of providing feedbacks (as was done in Part 4) results in slightly weaker average response behavior (averages are closer to zero). Similar results hold when we use a finer partition on intervals (not presented here). These aggregate results suggest that both the protection of self-esteem and reciprocal responses to the monetary consequences of the feedback provided plays a potential role in determining the response of agents to subjective feedback. The final two columns present the minimal and maximal rewards used by agents. Extreme responses of -100 (maximal sanction) and 100 (maximal reward) are always present, regardless of the value of $feed - own$. Careful examination of the data reveals that a small fraction of agents appear to be randomly selecting one of these two extreme values when making all their reward/sanction decisions. This apparent noise will require a careful analysis in the econometric model presented below.

Agents selected rewards and sanctions from a set which contained 201 possible values $\{-100, -99, ...0, ..., 99, 100\}$. In practice, observed rewards and sanctions belong to a much smaller set. The bottom panel of Table 2 illustrates this phenomena. We find that 11.45% of all rewards decisions in Part 3 and 9.31% are either -100 or 100, the lowest and highest possible values. We also find a huge bunching of responses at zero, indicating absence of rewards or sanctions (43.3% of responses in Part 3, and 53.4% in Part 4). Responses using other multiples of 10 are also very prominent (35.49% of Part 3 responses and 28.96% of Part 4 responses). Finally, decisions whose value ends in other multiples of 5 are also observed in the data, but to a lower extent. In total, we find the four response patterns presented in the panel capture 97.3% of observed decisions in Part 3, and 97.77% of observed decisions in Part 4. A natural interpretation of the data reported in the panel is that agents round their decisions to some degree (i.e. rounding to the nearest multiple of 5 or 10). Rounding introduces non-standard measurement error which may bias estimates which take the data at face value (see Kleinjans and van Soest (2014) for a recent discussion of these issues). The econometric model presented below will control for possible rounding and allow us to separate rounded responses from genuine responses.

Finally, Table 4 in the appendix replicates the aggregate statistics presented in the top part of Table 2, broken down by gender. We find little differences between the average reward/sanction behavior of men and women. We conducted formal tests of the null hypothesis that gender effects are not present in the data. In particular, we regressed rewards (sanctions coded as negative values) on a set of dummy variables for each category $feed - own$ in Table 2 (using $feed - own \leq -50$ as a reference category), a gender binary indicator, and interactions of this indicator with the category specific dummy variables. Separate regressions were run for Part 3 and Part 4, clustering standard errors at the individual level. The gender indicators and none of the interactions were significant at the 10% significance level, suggesting no gender differences in reward/sanction behavior. These descriptive results, combined with the lack of gender effects in over(under)confidence reported above, point towards limited differences in behavior between men and women in our experiment. A model-based test of gender effects will be presented in the following section.

4 Model

In this section we present an econometric model linking self-confidence to the 20 reward/sanction decisions observed in the experiment. We model reward/punishment decisions of agents by specifying preferences of agents over their own payoffs and the payoffs of the principal. These preferences are flexible and allowed to depend on the subjective evaluations of both the principal and the agent as well as on whether the agent's payoff is contingent on the principal's evaluation. This part of the model also takes into account possible rounding of reward/punishment decisions made by subjects in the experiment. We model self-confidence through a separate equation linking own subjective evaluation to objective productivity and assigned clicking time environment. Conditional on these two variables, deviations of own subjective evaluations identify under and over confidence agents. We allow these deviations to be correlated with preferences driving reward/punishment behavior.

For clarity, subscripts for agents and reward/sanction decisions are omitted in this section. These subscripts are introduced in the appendix where the likelihood function is presented. Let $feed_p$ represent the principal's feedback provided in part $p \in \{3, 4\}$ of the experiment, own_p the agent's own subjective evaluation, and rew the reward given by the agent to the principal (punishment corresponds to negative values of rew). Furthermore, define the following five dummy variables

$$\begin{aligned}
d_1 &= 1\left[feed_p - own_p \leq -50\right] \\
d_2 &= 1\left[feed_p - own_p \in \{-40, -30, -20 - 10\}\right] \\
d_3 &= 1\left[feed_p - own_p = 0\right] \\
d_4 &= 1\left[feed_p - own_p \in \{10, 20, 30, 40\}\right] \\
d_5 &= 1\left[feed_p - own_p \geq 50\right]
\end{aligned}$$

These dummy variables separate ranges of possible differences in evaluations between principal and agents $feed_p - own_p$ in five mutually exclusive categories. Moreover, let d_{part3} and d_{part4} denote binary variables separating parts 3 and 4 of the experiment.

We begin by assuming that a proportion ω of agents have convex preferences defined over own earnings m_p, and earnings of the principal y_p,

$$\begin{aligned}
U(rew) &= \alpha_1 m_p + \alpha_2 y_p + \alpha_3 m_p^2 + \alpha_4 y_p^2 + \alpha_5 m_p y_p + \varepsilon \\
\alpha_j &= \sum_{k=1}^{5} \alpha_{jk} d_k, \text{ for } j = 1, 2, 3, 4 && (1) \\
\alpha_5^p &= \sum_{k=1}^{5} \alpha_{5k}^p d_k + \mu \text{ with } p \in \{3, 4\} && (2) \\
m_p &= \left(\left[100 + 10 \times feed_p\right] d_{part3} + 150 d_{part4} - 0.25 \times |rew|\right) && (3) \\
y_p &= \left(300 + own_p \times 25 + rew\right) - \left[100 + 10 \times feed_p\right] d_{part3} - 150 d_{part4} && (4)
\end{aligned}$$

We allow the effects of the linear and quadratic terms in m_p and y_p to vary across the five levels of $feed_p - own_p$ captured by d_k with $k = 1, 2, ..., 5$. The interaction parameter α_5^p is further allowed to vary across parts 3 and 4 of the experiment. This provides a means to test whether agents behave differently when evaluations have financial consequences.[9] We also introduce unobserved heterogeneity in preferences through this parameter by adding the random unobservable variable μ. The later is assumed to be normally distributed with mean 0 and variance σ_μ^2. Finally, ε are unobservable random components determining utility. These components are added to the utility of each choice alternative and are assumed to follow independent standard Extreme Type 1 value distributions.

Observation of the data reveals two notable aspects worth addressing. First, many

[9]We attempted to estimate a more flexible specification allowing all preference parameters to shift across parts 3 and 4. Results proved numerical instable and inconclusive.

agents appear to randomly choose one of three rewards (-100,0, or 100) irrespective of their own evaluation or the principal's feedback. We interpret this behavior as noise, perhaps reflecting the agent's incomprehension of the experiment. We capture this in our model by allowing a proportion $(1-\omega)$ of agents to randomly choose rewards -100, 0, or 100 with equal probability. Second, agents could choose any of the 201 rewards in the set $\{-100,-99,...,99,100\}$. We documented in Section 3 that observed rewards/punishments are predominantly multiples of 5, 10, 25, or 50. The observed bunching of rewards at these multiples is largely inconsistent with convex preferences specified above. Instead, agents very likely round their responses in some way to simplify their decision process. Rounding of decisions represents measurement error which can potential bias estimated preferences. We follow Kleinjans and van Soest (2014) and model rounding of rewards by assuming agents use one of five rounding regimes during the experiment. Let the variable R capture the rounding pattern used, where

$$
\begin{aligned}
R &= 1: \text{ the reward is not rounded} \\
R &= 2: \text{ the reward is rounded to a multiple of 5} \\
R &= 3: \text{ the reward is rounded to a multiple of 10} \\
R &= 4: \text{ the reward is rounded to a multiple of 25} \\
R &= 5: \text{ the reward is rounded to a multiple of 50}
\end{aligned}
$$

Let $\rho_r = \Pr(R = r)$. Furthermore, another prominent feature of our data is the important occurrence of $rew = 0$, with agents not rewarding nor punishing principals. As discussed above, rounding may be one reason explaining bunching at this value. Another explanation is that agents suffer from status quo bias. Status quo biases in decision making can occur when individuals have a status quo alternative. Examples of such alternatives are doing nothing or repeating previous decisions. In our context, agents were unable to reward or punish principals in earlier parts of the experiment. Moreover, setting $rew = 0$ is payoff neutral for the principal. Both features of our design may contribute to $rew = 0$ being perceived as the status quo alternative. What is more, experimental evidence suggests that status quo biases can be particularly important when the number of alternatives in the choice set is large, as in our experiment (see Samuelson and Zeckhauser, 1988). To capture this bias, let σ denote the probability that an agent succumbs to this bias and sets $rew = 0$ irrespective of their true underlying preferences. We estimate $\{\rho_r : r = 1,2,...,5\}$ and σ along with the other model parameters. Our econometric model can thus be used for example to determine the share of observed rewards of zero reflecting a genuine zero, a rounded value, or a status quo response which is unrelated to the true underlying preferences of an agent.

We capture self-confidence my modelling own evaluations own_p using an ordered response choice model. Let own_p^* denote the latent unobserved subjective own evaluation of an agent. This latent variable is related to

$$own_p^* = \gamma_1 speed + \gamma_2 prod + \varepsilon_p + \theta\mu \tag{5}$$

$$own_p = j \text{ if } \tau_{t-1} \leq own_p^* \leq \tau_t \tag{6}$$

where $\tau_0 = -\infty, \tau_1, ...\tau_9, \tau_{10} = \infty$, are threshold parameters, *speed* denotes the clicking time randomly assigned to each agent. Relative to the main clicking time environment, screens times in the slow and fast clicking environments are multiplied by 0.75 and 1.25 respectively. The variable *speed* takes these three values (0.75,1,1.25) depending on the clicking time environment assigned to each agent. We expect a positive effect of *speed* on own evaluations. The variable *prod* represents an objective measure of the clicking ability constructed by taking the total number of clicks made by an agent in parts 1 and 2 of the experiment. Controlling for clicking ability is essential to identify under-confidence and over-confidence as deviations below and above the expected value of own for given levels of *speed* and *prod*. These deviations are captured by the error term in (5). This error term has two parts. The first part ε_p captures unobserved determinants of own evaluations in part p of the experiment and is assumed to be normally distributed with mean 0 and variance set to 1. The second part $\theta\mu$ is introduced to capture possible correlation between own evaluations and preferences of the agent (captured by μ), where θ represents an unknown parameter. A value of $\theta = 0$ implies that own evaluations are unrelated to preferences. A negative value of θ on the other hand implies that agents with a lower value of α_5^p tend to be overconfident and report higher than average own evaluations, conditional on their assigned clicking speed and objective productivity. Recall that lower value of α_5^p are associated with agents whose marginal utility of own income increases weakly or possibly decreases with income of the principal. These agents are more predisposed to use costly punishment to reduce the earnings of the principal. Hence, negative values of θ are consistent with overconfident agents sanctioning more principals for what they perceive is unfair feedback.

Intuition for identification of the model parameters follows. Data on own evaluations identify (γ_1, γ_2) as well as the vector of threshold parameters τ. Correlation between own evaluations and reward decisions identify θ. Rounding probabilities ρ_r are identified by the observed bunching of reward decisions in the data. The proportion of non rounded rewards is determined by the proportion of reward decisions which are not multiples of 5. Rounding to a multiple of 10 is determined by the difference between the number of rewards which are multiples of 10 and rewards which are multiples of 5. Other rounding probabilities are identified in a similar way. Finally, the proportion of zero rewards due to status quo bias is given by the difference between the proportion of zero rewards observed in the data and the proportion predicted allowing only for rounding and genuine zero rewards.

The model described above is estimated using reward/punishment decisions and agent's evaluations in Parts 3 and 4 of the experiment. Estimation is done by Maximum Simulated Likelihood, integrating over the distribution of μ.

5 Results

5.1 Parameter estimates

Estimated parameters are presented in Table 3. Most of the preference parameters cannot be directly interpreted, since they individually contribute only part of the marginal utility of rewarding/punishing the principal for given self-confidence and subjective feedback received. All preference parameters are precisely estimated and significant at the 5% level. We reject the null hypothesis that preferences are the same in parts 3 and 4 of the experiment (p-value = 0.000).[10] The later suggests that associating monetary consequences to subjective feedback does impact behavior. Simulations presented below will be used to quantify the effects of monetary consequences on behavior. We also find evidence of significant unobserved heterogeneity in preferences, captured by the variance parameter σ_μ^2.

Model estimates also reveal different forms of measurement error. Estimates suggest that 4.6% of agents (1-$\widehat{\omega}$) randomly choose rewards in the set $(-100, 0, 100)$ with equal probability, irrespective of their true underlying preferences. Estimates also suggest sizeable rounding of reward/sanction decisions, with only 3.3% of decisions not rounded ($R = 1$ as defined in Section 4). Dominant rounding patterns include rounding to the nearest multiple of 10 ($R = 3$ as defined in Section 4) 34.8% of the time, and rounding to the nearest multiple of 50 ($R = 5$) 46.3% of the time. Other rounding regimes are

[10]The estimated p-value is based on the Wald test statistic of the joint hypotheses that $\alpha_{3k}^3 = \alpha_{3k}^4$ $\forall k = 1, 2, .., 5$.

smaller in magnitude. Finally, we do not find evidence of residual status quo bias given the other forms of measurement errors captured by the model, with the estimated value of σ practically equal to 0.

We also find that clicking speed and objective productivity of agents both have a positive and significant impact on own evaluations. Finally, correlation between reward and sanction behavior and self-confidence is captured by θ. The estimated value of the later is negative and significantly different from zero (p-value = 0.000), suggesting a significant link between self-confidence and reward/punishment behavior. The simulated impact of this link on reward and sanction behavior is discussed below.

5.2 Model simulations

We simulated the behavioral implications of the model using parameter estimates presented in the appendix. Because these estimates are based on a model taking into account rounding and status quo biases, simulations have the advantage of capturing the predicted behavior of agents net of these considerations. Figure 7 presents simulated behavior interacting three intervals of agent own evaluation about their performance ($[10, 20], [50, 60], [80, 90]$) and three levels of their true (unobservable) objective productivity (15%, 55%, and 85%). Differences in own evaluation intervals are presented row wise in increasing values from top down. Differences in objective productivity are presented column wise in increasing value from left to right. For each combination of own evaluation and objective productivity, we report reward (positive values on the vertical axis) and punishment (negative values on the vertical axis) behavior for all possible evaluations received by the principal (horizontal axis, from 0 to 100). Simulations are presented for Part 3 (black lines) and Part 4 (red lines) of the experiment, allowing to assess the impact of monetary consequences. Dashed lines present point estimates, while solid lines present 95% confidence bands. Graphs on the diagonal of Figure 7 predict behavior when own evaluation intervals contain the true objective performance of the agent. Expectations about objective productivity are thus accurate, displaying no over- or underconfidence. Graphs off the diagonal capture behavior of over and underconfident subjects whose own evaluation interval respectively lies above and below their objective productivity.

Graphs below the diagonal capture behavior for overconfident agents. The subjective feedback threshold where agents switch from punishing to rewarding principals shifts to the right as over-confidence increases, tracking the agents' own evaluation interval about their performance. It is important to emphasize that the estimated model does not impose such thresholds nor that thresholds coincide with the agents' own evaluation interval. Severity of punishment below these thresholds is directly related to the level of over-confidence. We find that punishment is particulary severe in the most extreme case (bottom left hand corner graph), and is otherwise a salient feature of behavior for all over-confidence levels considered. Interestingly, we find little evidence that overconfident agents will reward principals for feedback exceeding their own evaluation. This is perhaps more apparent for agents who believe having clicked between 50% and 60% of all boxes, despite having objectively clicked 25% of all boxes (middle graph in the first column). In this graph, the threshold is centered on the subjective evaluation interval. We find close to no reward for subjective feedback exceeding own evaluations. Similar results emerge in the graph below. Finally, we find very little differences in behavior for overconfident subjects across parts 3 and 4. Overall these results suggest that the reciprocal response of overconfident agents is relatively independent of possible financial consequences of the subjective feedback.

Graphs above the diagonal capture behavior for different levels of underconfident agents. In sharp contrast with overconfident agents, we find that underconfident agents are reluctant to engage in punishment behavior - a very prominent feature of all three graphs above the diagonal in Figure 7. This also holds when the subjective feedback of principals falls below the agent's own evaluations. Furthermore, in contrast to overconfident agents we find that underconfident agents have a stronger propensity to reward principals for feedback that exceed their own. This is perhaps most apparent for agents who believe having clicked between 20% and 30% of all boxes, despite having objectively clicked 85% of all boxes (upper graph in the third column). Finally, the effects of monetary consequences become visible when comparing behavior in Part 3 and Part 4 for underconfident agents. We find that underconfident agents tend to reward principals more for payoff relevant subjective feedback that exceed their own (see e.g. the graph displaying agents who believe having clicked between 50% and 60% of all boxes, despite having objectively clicked 85% of all boxes (middle graph in the third column).

Lastly, agents with accurate expectations (graphs on the diagonal) punish principals for feedback below their own evaluations, especially when feedback is very negative (having clicked $[0, 10]$% of all boxes). The propensity to reward remains relatively weak however. A notable exception concerns agents who accurately evaluate that they clicked between 80% and 90% of all boxes. The later group is predicted to return close to 25 points to the principal in return for accurate feedback.

6 Conclusion

The optimal design of subjective evaluations requires understanding the forces driving how workers respond to agreeing or conflicting feedback about their productivity. This paper provided direct experimental evidence on the factors underlying how workers respond to subjective evaluations, allowing to separate financial implications from psychological forces involving self-confidence and the ego-threatening nature of these evaluations.

Overconfident workers are predicted to punish significantly more evaluations falling below their own than underconfident workers. The later group on the other hand is predicted to reward generous evaluations significantly more. These results indicate that the mix of underconfident and overconfident workers present in a firm will determine the profitability and usefulness of subjective evaluations and ultimately determine whether and how such evaluations should be provided. Firms largely populated by underconfident workers have more flexibility to evaluate workers based on their perceived merits, as evaluations falling below the workers own evaluations will unlikely be met with punishment. Firms populated with overconfident workers on the other hand may be inclined to avoid providing conflicting evaluations in order to avoid costly punishment, even if these evaluations do accurately reflect worker performance.

Our results point towards the development of new and more refined models of human behavior where self-confidence and self-image concerns play a central role in shaping how individuals interact and respond to actions of others. Existing models that already embark into this direction are Bénabou and Tirole (2006) and Ellingsen and Johannesson (2008) which allow people's self-image concerns to interact with their social preferences towards others. Such models would have predictive power in all strategic situations in which people are concerned about their self-image and are willing to invest resources to protect it independent of whether the actions of others have financial consequences for them or not.

7 Appendix : Log-likelihood function

Let $\mathbf{rew}_i$ denote the 20 x 1 vector of observed rewards/punishments of agent i across parts 3 and 4 of the experiment, where $\mathbf{rew}_i = [rew_{i1}, rew_{i2}, ..., rew_{i20}]'$, and where $rew_{ij} \in \{-100, -99, ..., 99, 100\}$.

Our distributional assumptions imply

$$\Pr(rew_{ij} = k) = \frac{\exp(U_{ij}(k))}{\sum_{m=-100}^{100} \exp(U_{ij}(m))} \tag{7}$$

denote the choice probability absent rounding or other measurement error, where conditioning on payoffs and μ are left implicit to simplify notation.

$$
\begin{aligned}
L_{ij} &= (1-\sigma)\left[\rho_1 \Pr\left(rew_{ij}=k\right)+\rho_2 \Pr\left(rew_{ij}\in\{k-2,k+2\}\right)\right] \\
&\quad \text{ if } k\in\{-95,-85,-65,-55,-45,-35,-15,-5,5,15,35,45,55,65,85,95\} \\
&= (1-\sigma)\left[\rho_1 \Pr\left(rew_{ij}=k\right)+\rho_2 \Pr\left(rew_{ij}\in\{k-2,k+2\}\right)+\rho_3 \Pr\left(rew_{ij}\in\{k-4,k+5\}\right)\right] \\
&\quad \text{if } k\in\{-90,-80,-70,-60,-40,-30,-20,-10\} \\
&= (1-\sigma)\left[\rho_1 \Pr\left(rew_{ij}=k\right)+\rho_2 \Pr\left(rew_{ij}\in\{k-2,k+2\}\right)+\rho_3 \Pr\left(rew_{ij}\in\{k-5,k+4\}\right)\right] \\
&\quad \text{if } k\in\{10,20,30,40,60,70,80,90\} \\
&= (1-\sigma)\left[\rho_1 \Pr\left(rew_{ij}=k\right)+\rho_2 \Pr\left(rew_{ij}\in\{k-2,k+2\}\right)+\rho_3 \Pr\left(rew_{ij}\in\{k-4,k+5\}\right)\right. \\
&\quad \left.+\rho_4 \Pr\left(rew_{ij}\in\{k-12,k+12\}\right)\right] \\
&\quad \text{if } k\in\{-75,-25\} \\
&= (1-\sigma)\left[\rho_1 \Pr\left(rew_{ij}=k\right)+\rho_2 \Pr\left(rew_{ij}\in\{k-2,k+2\}\right)+\rho_3 \Pr\left(rew_{ij}\in\{k-5,k+4\}\right)\right. \\
&\quad \left.+\rho_4 \Pr\left(rew_{ij}\in\{k-12,k+12\}\right)\right] \\
&\quad \text{if } k\in\{25,75\} \\
&\quad (1-\sigma)\left[\rho_1 \Pr\left(rew_{ij}=k\right)+\rho_2 \Pr\left(rew_{ij}\in\{k,k+2\}\right)+\rho_3 \Pr\left(rew_{ij}\in\{k,k+5\}\right)\right. \\
&\quad \left.+\rho_4 \Pr\left(rew_{ij}\in\{k-12,k+12\}\right)+\rho_5 \Pr\left(rew_{ij}\in\{k,k+25\}\right)\right] \\
&\quad \text{if } k\in\{-100\} \\
&\quad (1-\sigma)\left[\rho_1 \Pr\left(rew_{ij}=k\right)+\rho_2 \Pr\left(rew_{ij}\in\{k-2,k+2\}\right)+\rho_3 \Pr\left(rew_{ij}\in\{k-5,k+4\}\right)\right. \\
&\quad \left.+\rho_4 \Pr\left(rew_{ij}\in\{k-12,k+12\}\right)+\rho_5 \Pr\left(rew_{ij}\in\{k-24,k+25\}\right)\right] \\
&\quad \text{if } k\in\{-50\} \\
&\quad (1-\sigma)\left[\rho_1 \Pr\left(rew_{ij}=k\right)+\rho_2 \Pr\left(rew_{ij}\in\{k-2,k+2\}\right)+\rho_3 \Pr\left(rew_{ij}\in\{k-4,k+5\}\right)\right. \\
&\quad \left.+\rho_4 \Pr\left(rew_{ij}\in\{k-12,k+12\}\right)+\rho_5 \Pr\left(rew_{ij}\in\{k-25,k+24\}\right)\right] \\
&\quad \text{if } k\in\{50\} \\
&\quad (1-\sigma)\left[\rho_1 \Pr\left(rew_{ij}=k\right)+\rho_2 \Pr\left(rew_{ij}\in\{k-2,k\}\right)+\rho_3 \Pr\left(rew_{ij}\in\{k-4,k\}\right)\right. \\
&\quad \left.+\rho_4 \Pr\left(rew_{ij}\in\{k-12,k\}\right)+\rho_5 \Pr\left(rew_{ij}\in\{k-25,k\}\right)\right] \\
&\quad \text{if } k\in\{100\} \\
&= (1-\sigma)\left[\rho_1 \Pr\left(rew_{ij}=k\right)+\rho_2 \Pr\left(rew_{ij}\in\{k-2,k+2\}\right)+\rho_3 \Pr\left(rew_{ij}\in\{k-4,k+5\}\right)\right. \\
&\quad \left.+\rho_4 \Pr\left(rew_{ij}\in\{k-12,k+12\}\right)+\rho_5 \Pr\left(rew_{ij}\in\{k-24,k+24\}\right)\right]+\sigma \\
&\quad \text{if } k=0 \\
&= (1-\sigma)\left[\rho_1 \Pr\left(rew_{ij}=k\right)\right] \text{ otherwise}
\end{aligned}
$$

where $\Pr(rew_{ij} \in \{k-a, k+b\})$ sums $\Pr(rew_{ij} = k)$ given in (7) over the integer values in the set $\{k-a, k+b\}$.

Likelihood of own_p is given by a ordered probit model

$$\Omega_{pi} = \Lambda(\tau_{t,i} - \gamma_1 speed_i - \gamma_2 prod_i - \theta\mu_i) - \Lambda(\tau_{t-1,i} - \gamma_1 speed_i - \gamma_2 prod_i - \theta\mu_i)$$

where Λ denotes de cumulative distribution function of the standard normal distribution and $(\tau_{t,i}, \tau_{t-1,i})$ are the agent's threshold given own_p. Thresholds for both parts belong to a common set of parameters $\{\tau_t : t = 0, 1, ..., 10\}$ where $(\tau_0 = -\infty, \tau_{10} = +\infty)$ are normalizations.

The log-likelihood contribution for agent i is given by

$$LogL_i = \ln\left(\int \prod_{j=1}^{20} L_{ij}\Omega_{3i}\Omega_{4i} f(\mu_i)\, d\mu_i\right)$$

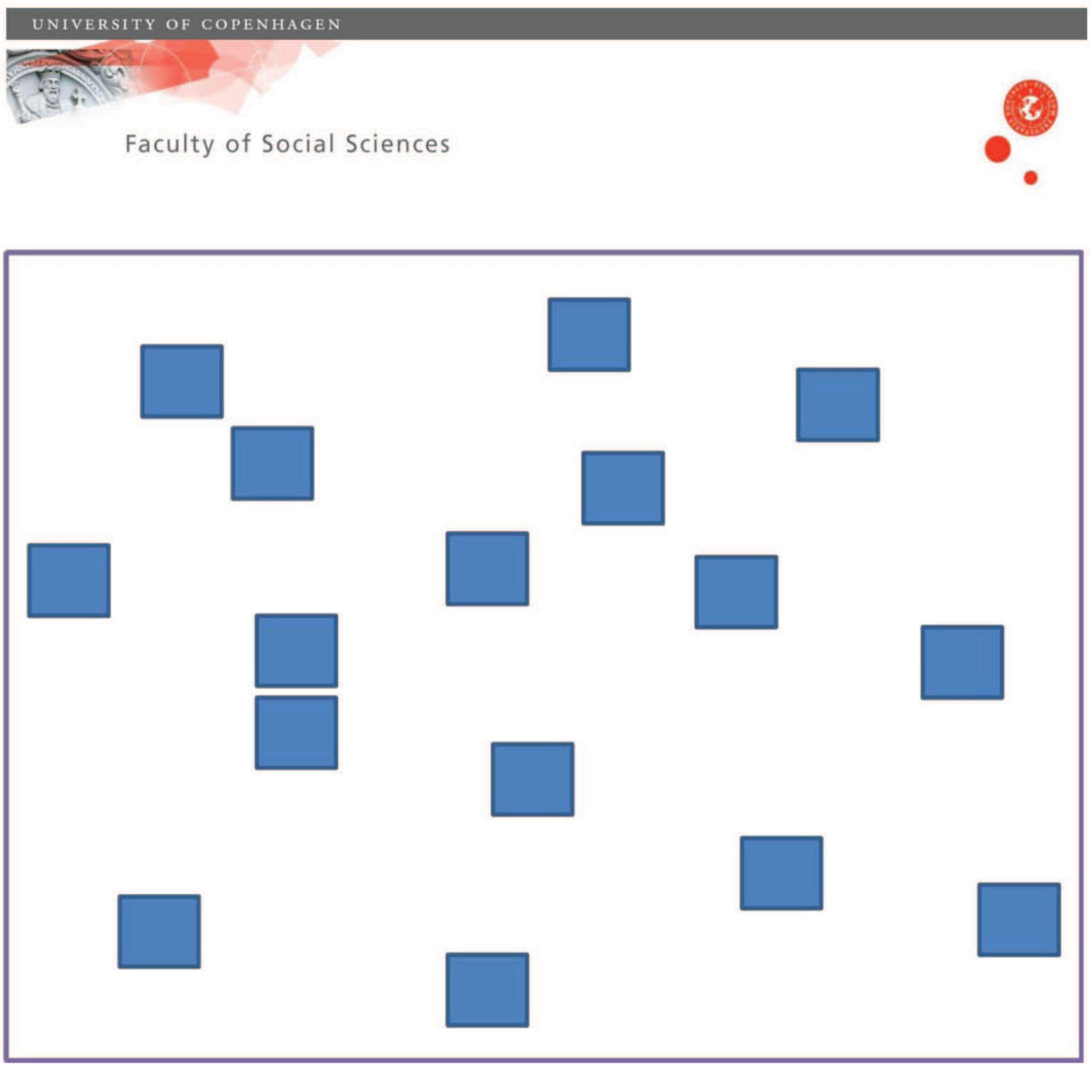

Figure 1: One screen of the clicking task.

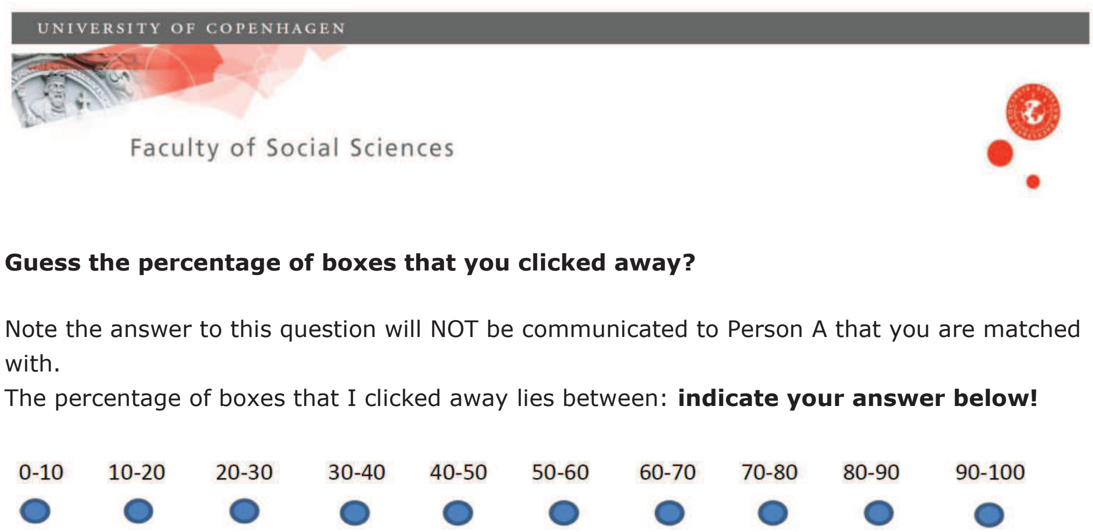

Figure 2: The evaluation screen.

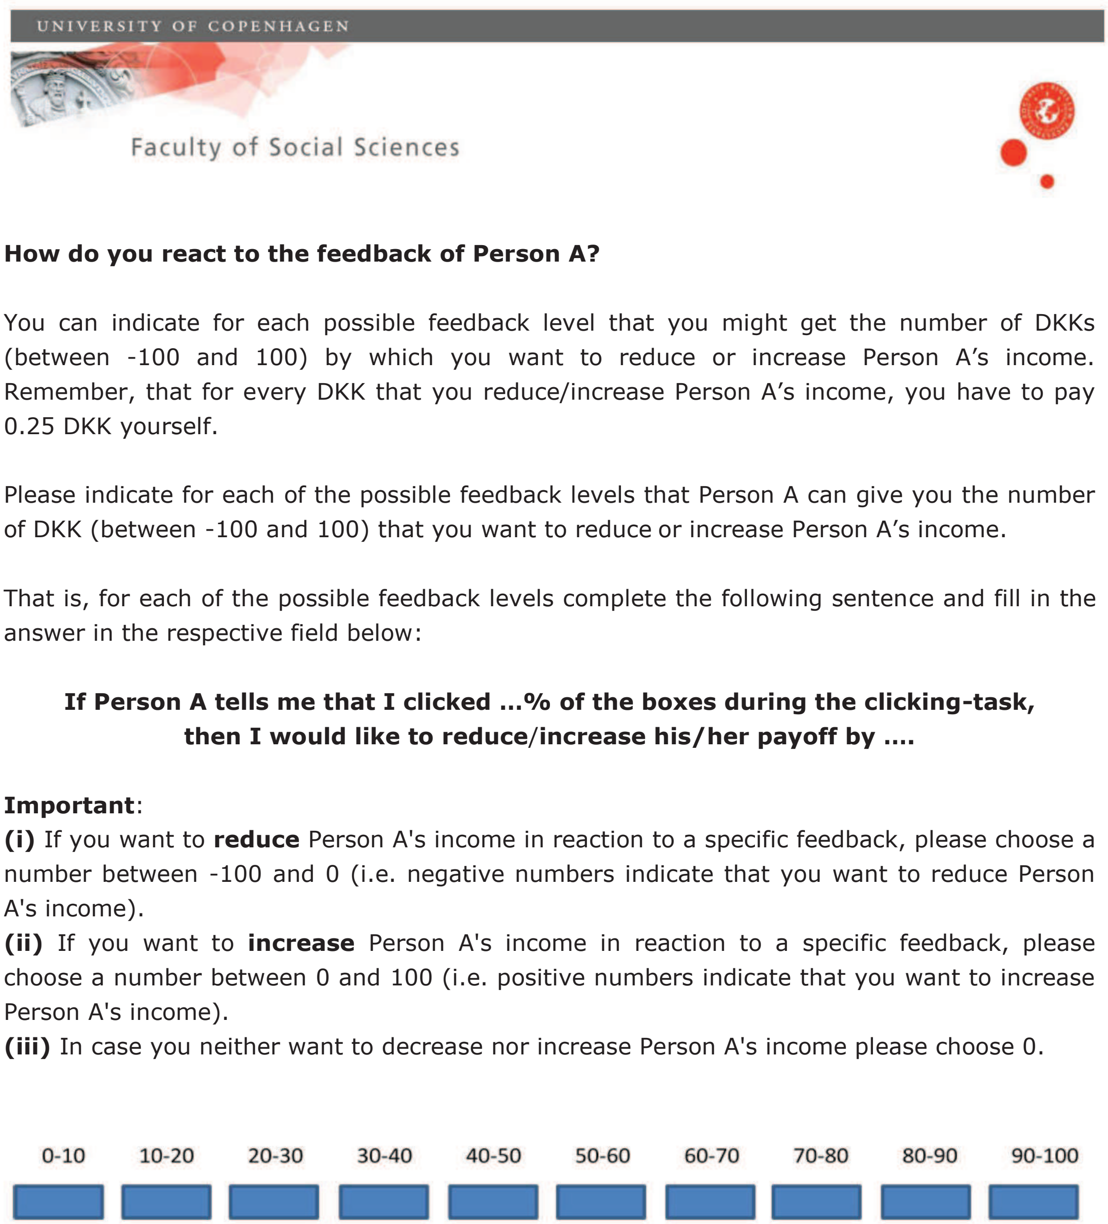

Figure 3: The reaction screen.

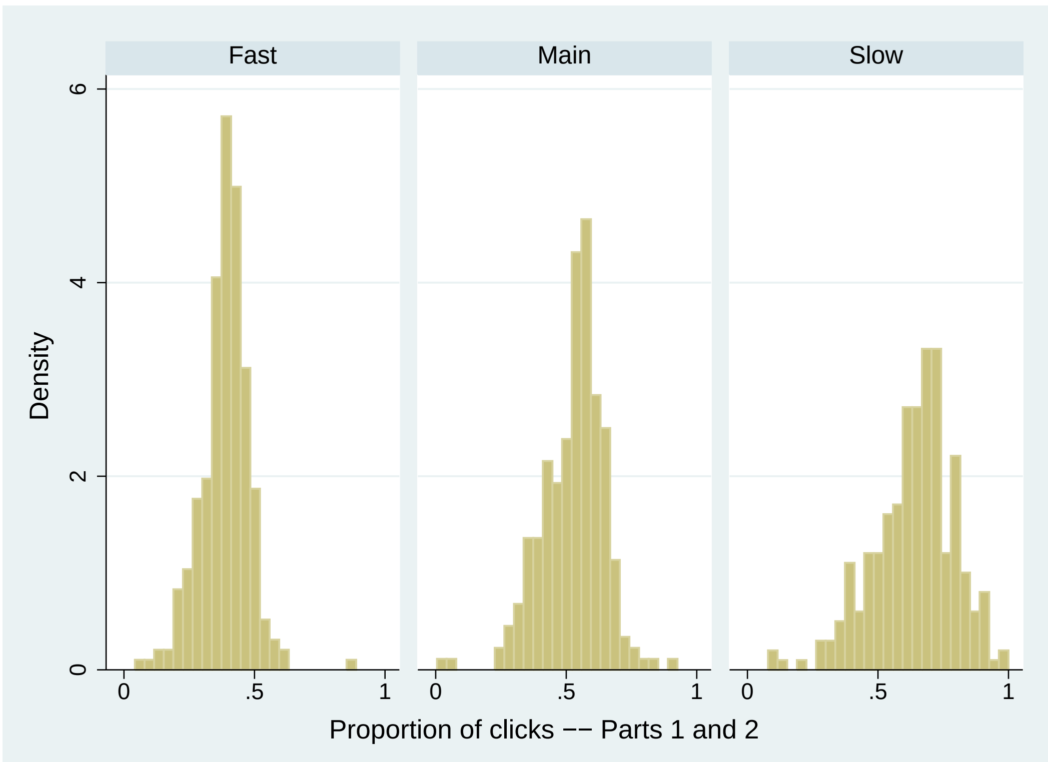

Figure 4: Proportion of clicks per player across Parts 1 and 2 of the experiment. Figures present the distributions for players assigned the lowest, intermediate, and highest time to click away boxes.

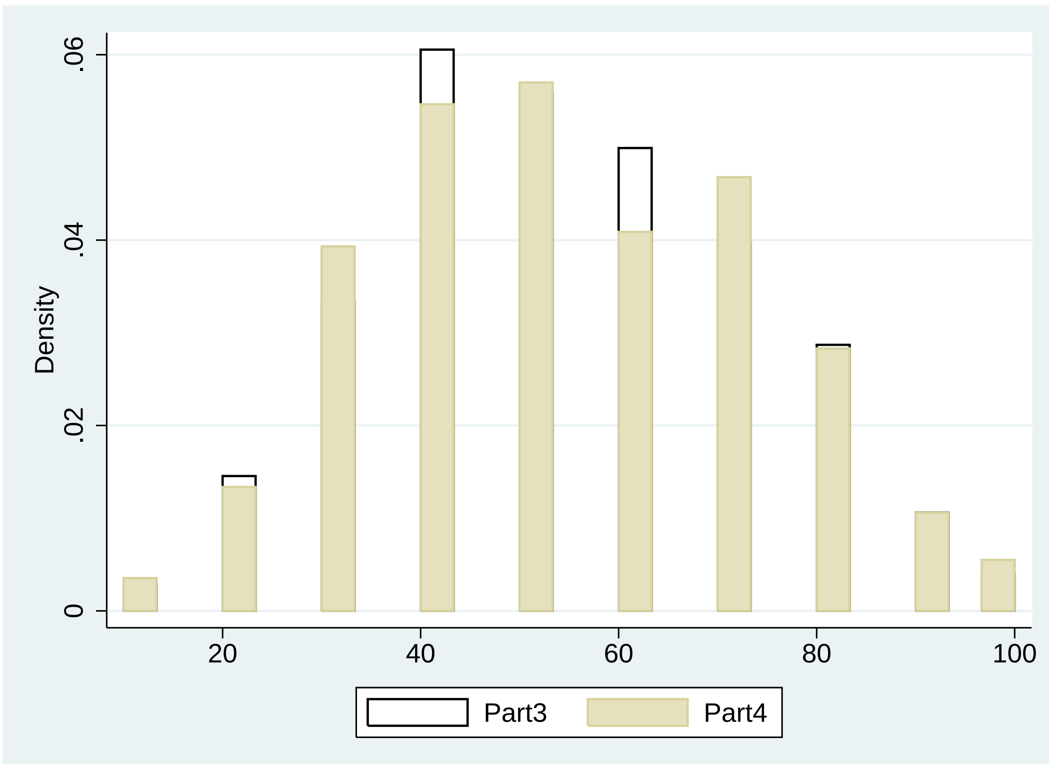

Figure 5: Distribution of agents' own evaluations for Part 3 and Part 4 of the experiment.

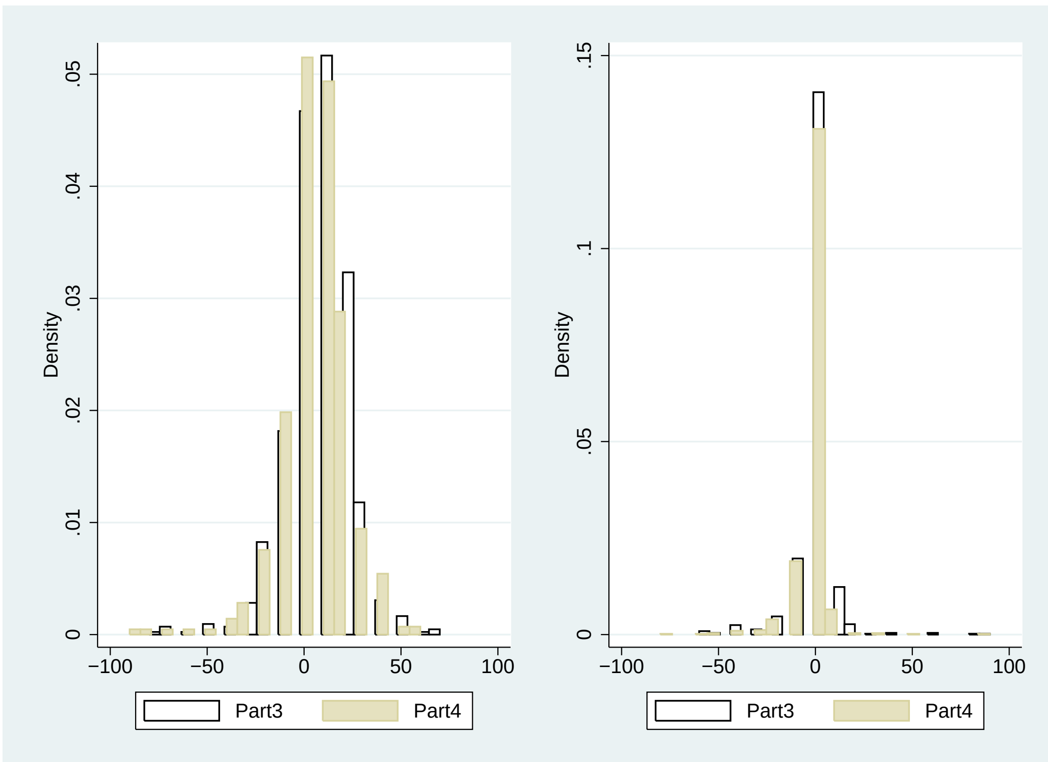

Figure 6: Left histogram plots the differences between the private evaluations of the principal-agent pairs. Right histogram plots the differences between the principal's private evaluations and the feedback provided.

	Objective productivity								
	[0,20]	(20,30]	(30,40]	(40,50]	[50,60]	[60,70]	[70,80]	[80,90]	[90,100]
Agent's evaluations									
% coincide with objective interval	0.667	0.289	0.352	0.316	0.257	0.357	0.379	0.417	0.468
% exceed objective interval	0.333	0.303	0.099	0.137	0.175	0.147	0.096	0.069	0
% below objective interval	0	0.406	0.549	0.546	0.568	0.496	0.524	0.514	0.531
Average evaluation	18.571	20.579	24.366	34.222	43.555	54.524	63.193	73.750	80.937
Principal's evaluations									
% coincide with objective interval	0.762	0.579	0.536	0.384	0.358	0.392	0.445	0.472	0.563
% exceed objective interval	0.095	0.318	0.247	0.265	0.429	0.444	0.247	0.264	0
% below objective interval	0.143	0.101	0.197	0.351	0.213	0.163	0.307	0.264	0.437
Average evaluation	4.288	23.333	32.019	39.403	52.857	62.857	66.566	76.388	77.812
Number of cases	21	69	213	386	315	252	166	72	32

Table 1: Sample averages and proportions of evaluations (agents/principals) within, above, or below the objective productivity intervals of agents, and . Proportions are computed combining data from Parts 3 and 4.

	Reward/sanction behavior							
	Part 3				Part 4			
	Avg.	Min	Max	N	Avg.	Min	Max	N
$feed - own \leq -50$	-25.042	-100	100	706	-17.573	-100	100	736
$feed - own \in \{-40, -30, -20, -10\}$	-6.720	-100	100	2582	-3.027	-100	100	2568
$feed - own = 0$	15.834	-100	100	761	14.973	-100	100	761
$feed - own \in \{10, 20, 30, 40\}$	15.649	-100	100	2676	10.943	-100	100	2644
$feed - own \geq 50$	10.014	-100	100	885	2.423	-100	100	901
	Proportion of rounding patterns							
-100 or 100	11.45				9.31			
0	43.30				53.40			
Other multiples of 10	35.49				28.96			
Multiples ending in 5	7.05				6.10			
Total	97.30				97.77			

Table 2: Top panel : Reward (positive values) and sanction (negative values) behavior in parts 3 and 4 as a function of $feed - own$. N denotes the number of cases. Bottom panel : rounding patterns present in the data for each part.

	Parameter	Std. err.		Parameter	Std. err.
α_{11}	-15.212	6.134	α_{51}	2.339	0.123
α_{12}	-25.349	2.717	α_{52}	1.542	0.076
α_{13}	-27.409	7.740	α_{53}	2.115	0.167
α_{14}	17.071	3.474	α_{54}	2.804	0.081
α_{15}	39.619	4.623	α_{55}	4.500	0.124
α_{21}	-17.805	1.154	ω	0.954	0.007
α_{22}	-13.634	0.542	ρ_1	0.033	0.001
α_{23}	-14.522	1.651	ρ_2	0.093	0.002
α_{24}	-20.055	0.861	ρ_3	0.348	0.003
α_{25}	-23.847	1.084	ρ_4	0.063	0.003
α_{31}^3	2.155	0.647	ρ_5	0.463	0.004
α_{32}^3	4.547	0.423	σ	0.005	0.005
α_{33}^3	6.165	1.012	τ_1	0.835	0.150
α_{34}^3	7.119	0.445	τ_2	1.747	0.122
α_{35}^3	4.949	0.572	τ_3	2.689	0.118
α_{31}^4	2.422	0.532	τ_4	3.616	0.119
α_{32}^4	4.506	0.388	τ_5	4.435	0.126
α_{33}^4	5.836	1.008	τ_6	5.121	0.129
α_{34}^4	6.806	0.503	τ_7	5.972	0.137
α_{35}^4	3.622	0.758	τ_8	6.950	0.149
α_{41}	15.527	2.177	τ_9	7.748	0.169
α_{42}	17.869	1.070	γ_1	1.018	0.118
α_{43}	18.921	2.459	γ_2	1.546	0.032
α_{44}	2.299	0.924	σ_μ^2	3.297	0.046
α_{45}	-3.327	1.153	θ	-0.101	0.009

Table 3: Estimated model parameters and standard errors.

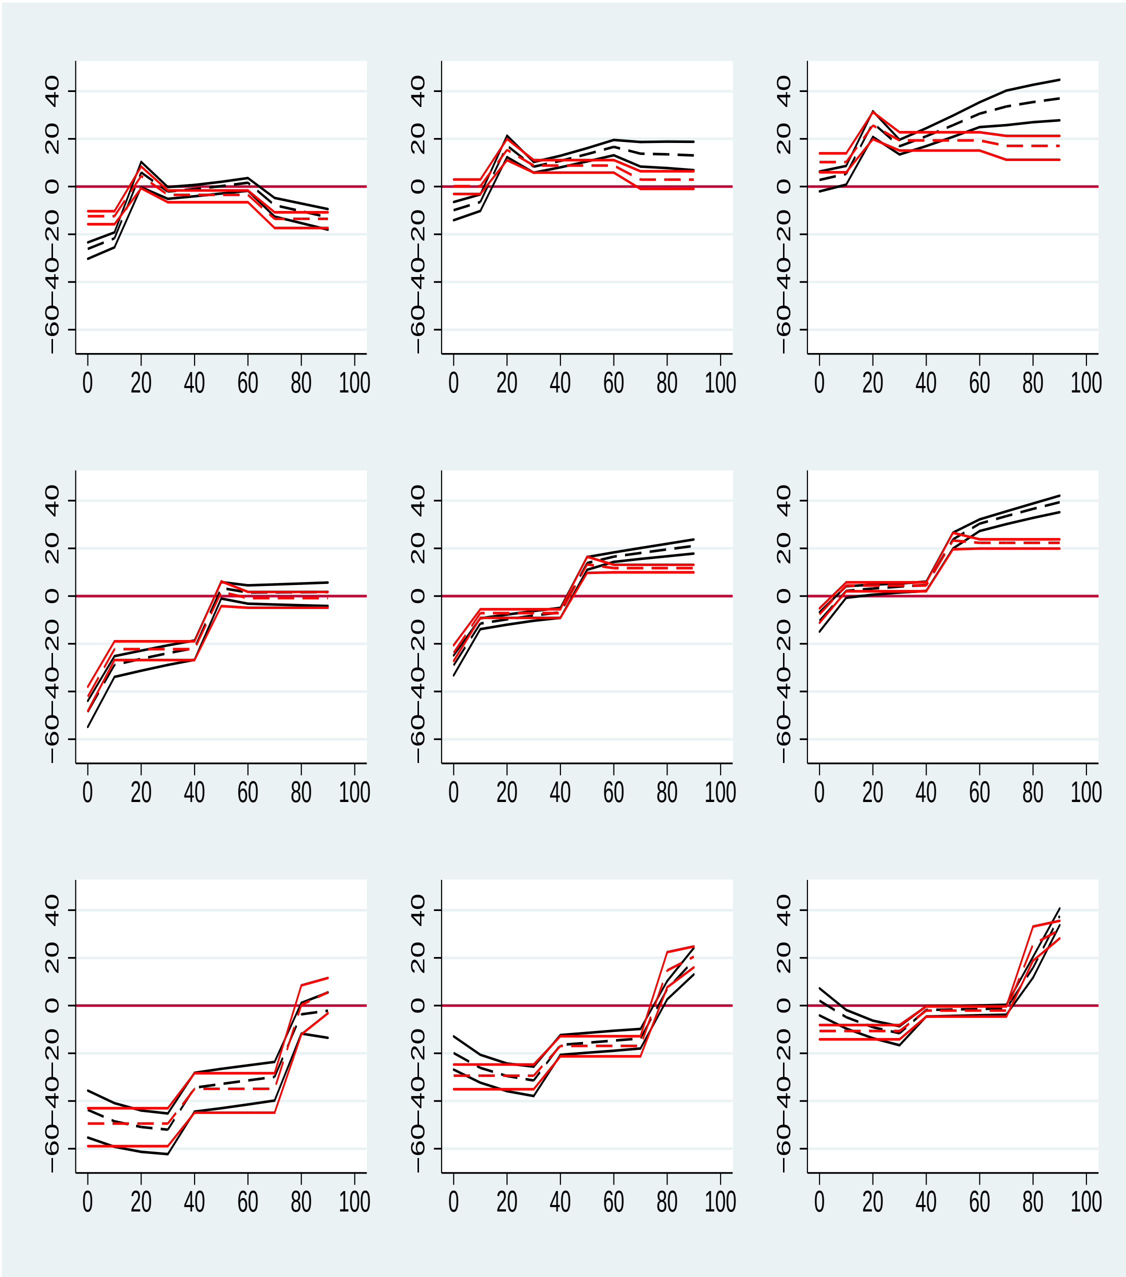

Figure 7: Simulated reward (positive values on vertical axes) and punishment (negative values on vertical axes) behavior as a function of subjective evaluations received (horizontal axes, from 0% to 100% of all boxes clicked). Graphs columns correspond to simulations when agents objectively clicked 25% (left column), 55% (middle column), and 85% (right column) of all boxes. Graphs rows correspond to simulations for agents who believed they clicked $[20, 30]$ percent (first row), $[50, 60]$ percent (second row), and $[80, 90]$ percent (last row) of all boxes. Point estimates (dashed lines) and 95% confidence intervals (solid lines) for Part 3 (black) and Part 4 (red).

XX. juli 2012

DANMARKS STATISTIK

«navn»
«adresselabel»
«POSTNR» «POSTNAVN»

Kære «Navn»
Danmarks Statistik og Internet Laboratoriet for Eksperimentel Økonomi (iLEE) ved Økonomisk Institut på Københavns Universitet inviterer dig hermed til at deltage i et eksperiment vedrørende forskellige beslutningsprocesser. Et af hovedformålene er blandt andet at undersøge hvordan vi selv reagerer når andre mennesker tager beslutninger, der har en indvirkning på os. Læs mere på bagsiden af dette brev, om det interaktive spørgeskema.

Ved at deltage i eksperimentet får du mulighed for at tjene penge. Vi kan ikke garantere dig, at du vil tjene et bestemt beløb, idet din indtjening vil afhænge af dine egne samt andre deltageres beslutninger. De nærmere regler er beskrevet på hjemmesiden for eksperimentet.

Hvordan gør jeg?
Dine beslutninger i eksperimentet bliver behandlet strengt fortroligt og anonymt. For at se detaljerne om eksperimentet, herunder opgaven, tidsforbrug mv., bedes du snarest muligt logge ind på vores hjemmeside:

https://applications.econ.ku.dk/LeeInternetSurvey

Login:	**[login]**
Password:	**[password]**

Forudsætninger for at kunne deltage i eksperimentet er at din computer skal have:
- Windows styresystem
- Microsoft Silverlight
- Explorer eller Google Chrome browser

Hvis Microsoft Silverlight ikke allerede er installeret på din computer, vil dette tilbydes gratis via et pop up vindue når der logges på hjemmesiden ovenfor.

Dette eksperiment er åbent for op til 2000 deltagere, men senest til og med mandag 6/8/2012.

Anonymitet og spørgsmål
Hvis du har problemer med at logge ind eller har yderligere spørgsmål, er du velkommen til at kontakte Økonomisk Institut ved at sende en e-mail til **ilee@econ.ku** eller ved at ringe til os på telefon 35324418**.**

Med venlig hilsen og på forhånd tak for din hjælp.

Isak Isaksen
Chef for Interviewservice
Danmarks Statistik

Alexander Sebald
Assistant Professor
Økonomisk Institut

Referencenr.: «Respnr»

Figure 8: Invitation letter

References

Baker, G., Gibbons, R. and Murphy, K. J. (1994). Subjective performance measures in optimal incentive contracts. *The Quarterly Journal of Economics*, **109** (4), 1125–1156.

Baumeister, R. F., Bushman, B. J. and Campbell, W. K. (2000). Self-esteem, narcissism, and aggression: Does violence result from low self-esteem or from threatened egotism? *Current directions in psychological science*, **9** (1), 26–29.

— and Finkel, E. J. (2010). *Advanced social psychology: The state of the science*. OUP USA.

— and Vohs, K. D. (2001). Narcissism as addiction to esteem. *Psychological Inquiry*, **12** (4), 206–210.

Bénabou, R. and Tirole, J. (2002). Self-confidence and personal motivation. *The Quarterly Journal of Economics*, **117** (3), 871–915.

— and Tirole, J. (2006). Incentives and prosocial behavior. *The American Economic Review*, **96** (5), 1652–1678.

Breuer, K., Nieken, P. and Sliwka, D. (2013). Social ties and subjective performance evaluations: an empirical investigation. *Review of Managerial Science*, **7** (2), 141–157.

Burks, S. V., Carpenter, J. P., Goette, L. and Rustichini, A. (2013). Overconfidence and social signalling. *Review of Economic Studies*, **80** (3), 949–983.

Campbell, W. K., Hoffman, B. J., Campbell, S. M. and Marchisio, G. (2011). Narcissism in organizational contexts. *Human Resource Management Review*, **21** (4), 268–284.

— and Miller, J. D. (2011). *The handbook of narcissism and narcissistic personality disorder: Theoretical approaches, empirical findings, and treatments*. John Wiley & Sons.

Dufwenberg, M. and Kirchsteiger, G. (2004). A theory of sequential reciprocity. *Games and Economic Behavior*, **47** (2), 268–298.

Ellingsen, T. and Johannesson, M. (2008). Pride and prejudice: The human side of incentive theory. *The American Economic Review*, **98** (3), 990–1008.

Falk, A. and Fischbacher, U. (2006). A theory of reciprocity. *Games and Economic Behavior*, **54** (2), 293–315.

Gervais, S. and Odean, T. (2001). Learning to be overconfident. *The Review of Financial Studies*, **14** (1), 1–27.

Gibbs, M., Merchant, K., Van der Stede, W. and Vargus, M. (2004). Determinants and effects of subjectivity in incentives. *Accounting Review*, pp. 409–436.

Hoelzl, E. and Rustichini, A. (2005). Overconfident: Do you put your money on it? *The Economic Journal*, **115** (503), 305–318.

Hoffman, M. and Burks, S. V. (2017). *Worker Overconfidence: Field Evidence and Implications for Employee Turnover and Returns from Training*. Tech. rep., National Bureau of Economic Research.

Hogarth, R. M. and Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, **24** (1), 1–55.

Kleinjans, K. J. and Soest, A. V. (2014). Rounding, focal point answers and nonresponse to subjective probability questions. *Journal of Applied Econometrics*, **29** (4), 567–585.

Köszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association*, **4** (4), 673–707.

MacLeod, W. (2003). Optimal contracting with subjective evaluation. *The American Economic Review*, **93** (1), 216–240.

Malmendier, U. and Tate, G. (2015). Behavioral ceos: The role of managerial overconfidence. *Journal of Economic Perspectives*, **29** (4), 37–60.

Miller, D. T., Ross, M. *et al.* (1975). Self-serving biases in the attribution of causality: Fact or fiction. *Psychological Bulletin*, **82** (2), 213–225.

Mobius, M. M., Niederle, M., Niehaus, P. and Rosenblat, T. S. (2011). Managing self-confidence: Theory and experimental evidence. *National Bureau of Economic Research*.

Moore, D. A. and Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, **115** (2), 502.

Prendergast, C. (1999). The provision of incentives in firms. *Journal of economic literature*, **37** (1), 7–63.

Rabin, M. and Schrag, J. L. (1999). First impressions matter: A model of confirmatory bias. *The Quarterly Journal of Economics*, **114** (1), 37–82.

Samuelson, W. and Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty*, **1** (1), 7–59.

Schwardmann, P. and Van der Weele, J. J. (2016). Deception and self-deception. *Mimeo*.

Sebald, A. (2010). Attribution and reciprocity. *Games and Economic Behavior*, **68** (1), 339–352.

	Part 3				Part 4			
	Reward/sanction behavior							
Men	Avg.	Min	Max	N	Avg.	Min	Max	N
$feed - own \leq -50$	-27.057	-100	100	433	-18.761	-100	100	431
$feed - own \in \{-40, -30, -20, -10\}$	-8.061	-100	100	1415	-4.518	-100	100	1415
$feed - own = 0$	18.333	-100	100	412	16.412	-100	100	412
$feed - own \in \{10, 20, 30, 40\}$	18.097	-100	100	1408	11.350	-100	100	1402
$feed - own \geq 50$	11.743	-100	100	452	2.154	-100	100	460
Women	Avg.	Min	Max	N	Avg.	Min	Max	N
$feed - own \leq -50$	-21.846	-100	100	273	-15.895	-100	100	305
$feed - own \in \{-40, -30, -20, -10\}$	-5.094	-100	100	1167	-1.197	-100	100	1153
$feed - own = 0$	12.882	-100	100	349	13.275	-100	100	349
$feed - own \in \{10, 20, 30, 40\}$	12.931	-100	100	1268	10.487	-100	100	1242
$feed - own \geq 50$	8.210	-100	100	433	2.705	-100	100	441

Table 4: Top panel : Reward (positive values) and sanction (negative values) behavior in parts 3 and 4 as a function of $feed - own$. N denotes the number of cases. Bottom panel : rounding patterns present in the data for each part.