Informational Differences and Performance: Experimental Evidence*

Iris Bohnet† and Farzad Saidi‡

August 23, 2012

Abstract

This paper provides experimental evidence on how informational differences may translate into performance differences in competitive environments. In a laboratory tournament setting, we manipulate beliefs about the effort-reward relationship by varying how much information people receive on the potential impact of luck on outcomes. We find that an informational disadvantage worsens the understanding of the effort-reward relationship, and significantly lowers performance. Our study is inspired by informational differences in the labor market where some individuals have less data on the determinants of economic success than others – due to social networks or the availability of similar others to learn from. (*JEL* C91, D81, M50)

*We wish to thank David Cesarini, Armin Falk, Guillaume Fréchette, Xavier Gabaix, Andy Schotter, and Richard Zeckhauser, along with seminar audiences at Universitat Autònoma de Barcelona (MOVE Workshop), Harvard Kennedy School, and the 2011 Stanford Institute for Theoretical Economics Summer Workshop, for helpful comments. Financial support from the Women and Public Policy Program and the Women's Leadership Board at Harvard Kennedy School is gratefully acknowledged.

†Harvard University, Kennedy School of Government, 79 John F. Kennedy Street, Cambridge, MA 02138. iris_bohnet@harvard.edu.

‡New York University, Department of Economics, 19 W. 4th Street, New York, NY 10012. saidi@nyu.edu.

1 Introduction

People differ in the degree to which they attribute economic outcomes to effort and luck. For example, Europeans are more likely than Americans to believe that luck, rather than effort or education, determines income (Alesina et al. 2001), and similarly women are more likely than men to attribute success to luck (Fisman and O'Neill 2009). In addition, differences in information may affect how people perceive the effort-reward relationship. Specifically, some individuals might have more precise information on the determinants of economic success at their disposal than others, depending on how many people they have available to learn from.

The number of similar others seems crucial for such information transmission. For example, in organizations, the demographic mix determines the set of comparable others. Role models and mentors tend to have the same demographic characteristics as their mentees (e.g., Holmes and O'Connell 2007; Ibarra 1992, 1993; Ragins 1999 for a review). Country-of-origin or same-language social networks facilitate job seeking (Edin et al. 2003, Munshi 2003), business relationships (Jackson and Schneider 2010), as well as the participation in welfare and social programs (Bertrand et al. 2000, Figlio et al. 2011). Information derived from a smaller sample will likely be less precise, i.e., have a higher variance, than when derived from a larger sample, potentially creating an informational disadvantage for members of smaller groups.

We conduct a laboratory experiment inspired by such labor market reality where some individuals have less information on the determinants of economic success than others. We assign information conditions randomly, and form two groups that differ in the amount of information they receive on the effort-reward relationship. The effect of such informational differences on performance is measured in a competitive setting. Similar to Orrison et al. (2004), we view (promotion) tournaments as an essential incentive device in modern hierarchical organizations. In tournaments, the impact of informational differences on performance also depends on whom people compete with and what they know about their competitors. To this end, we theoretically and experimentally examine competition between equally informed as well as between differentially informed agents.

In the experiment, performance was measured in a real-effort task, namely by the number of words found in a word-find task, where subjects were assigned to pairs and competed against their anonymous counterpart for a tournament prize. Tournament outcomes were determined by both individual effort and a random bonus component. We created two information conditions regarding the bonus component and, thus, the effort-reward relationship: one group received information from a large sample of data, and another group received information from a small sample. Depending on the information condition, a person was either well or poorly informed on the potential impact of the random component on tournament outcomes, and competed either with an equally or with a differentially informed counterpart. We conjecture that individuals that are better informed on the role of the luck component for economic success will exert more effort.

Experimental participants receiving less information on the effort-reward relationship indeed perceived the variance of their bonus component to be larger, and subsequently performed worse than better informed participants. Furthermore, the performance of an informationally disadvantaged agent was particularly depressed when competing with a better rather than an equally informed counterpart. This suggests that information potentially affects performance through two channels: by affecting how well agents understand the effort-reward relationship and by creating "informational injustice" that additionally discourages the disadvantaged from exerting effort.

Our work contributes to earlier experimental studies examining the role of uncertainty and information on individual performance in competitive environments, which is not conclusive. Bull et al. (1987) and Freeman and Gelber (2010) varied the amount of information on the past performance of competitors that their experimental subjects had available. For a hypothetical effort task, Bull et al. (1987) reported that subjects who were informed of their counterparts' decisions after each round exerted less effort than those who did not receive any information. In contrast, Freeman and Gelber (2010) who used a real-effort task (mazes) found that providing more information on the historical performance of competitors led to higher effort on average. In both cases, the uncertainty about the effort-reward relationship was influenced by the amount of information subjects had available on their competitors' past performance, making best-responding a difficult problem, as past performance may not necessarily map directly into future performance and responses are interdependent. We vary information exogenously – unlike, for instance, Celen and Hyndman forth., who allowed their subjects to costly acquire information about the decisions of predecessors – and independently of subject performance, enabling us to determine the impact of the perceived importance of luck on effort. In our setup, informational differences result from differences in sample sizes, which – in addition to mimicking labor market reality – is an intuitive way to communicate differential degrees of uncertainty to subjects.

The remainder of the paper is organized as follows. In Section 2, we discuss the theoretical framework, and derive hypotheses for our experiment. The experimental design is presented in Section 3. Section 4 reports the experimental results, and Section 5 concludes.

2 Theoretical Framework

We model effort choices among competing agents using a tournament-theoretical framework (Lazear and Rosen 1981). We propose a setting in which two agents compete against each other. They belong either to the better informed group receiving a *large* data sample or the worse informed group with a *small* data sample. Each agent can control the mean of the output distribution μ by means of effort exertion, which is costly ($C(\mu) > 0$). Furthermore, a stochastic luck component ε is realized. This leads to the following observable output:

$$q_i = \mu_i + \varepsilon_i, i = L, S \tag{1}$$

where L and S stand for two different agents with a large and a small data sample, respectively.

In this rank-order tournament, agent L differs from agent S in the perceived distribution of the luck component, so that L is better informed about the effort-reward relationship. As in the experiment, we distinguish between two cases: we compare the performance of L- and S-players when each type of player either assumes to compete with an equally informed counterpart, or assumes to compete with a differentially informed counterpart.

2.1 Heterogenous Beliefs under the Assumption of Identical Information

Our general setup is akin to that in Lazear and Rosen (1981). The specificity of our model lies in the beliefs of L and S:

- Agent L believes all luck components to be independent s.t. $\varepsilon_{L,L} \sim N\left(0, \sigma_L^2\right)$ and $\varepsilon_{S,L} \sim N\left(0, \sigma_L^2\right)$ where the variance reflects the precision of the perceived relationship between effort and pay, and $\varepsilon_{x,y}$ is agent y's belief about agent x's luck component.

- Agent S believes the respective luck components to be independent s.t. $\varepsilon_{S,S} \sim N\left(0, \sigma_S^2\right)$ and $\varepsilon_{L,S} \sim N\left(0, \sigma_S^2\right)$ where $\sigma_L^2 < \sigma_S^2$.

- However, the principal, i.e., the firm, believes that both agents perceive their luck components to be independent and normally distributed with zero mean and variance σ_L^2. This is a simplified way of saying that agent L's beliefs are better aligned with the actual effort-reward relationship, whereas agent S underestimates the extent to which effort translates into rewards.[1]

[1] This approach seems similar to the model of Lundberg and Startz (1983) – however, we do not assume statistical discrimination on the part of the employer, but heterogenous beliefs among employees depending on their group affiliation. Moreover, our model is different in that both agents believe the distribution of the luck component to be identical for everyone.

Furthermore, the following general assumptions apply:

- We consider a single tournament round.
- The cost function $C(\cdot)$ is quadratic and not a source of agent heterogeneity.
- The firm derives a marginal social return V from each unit of effort that an agent exerts, and acts in a perfectly competitive market.
- We denote the tournament prize spread by $\Delta W \equiv W_1 - W_2$ where W_1 and W_2 are the winner and loser prizes, respectively.

Before we move to the analysis of the game, we introduce the notion of a performance gap in this tournament.

Definition *A performance gap exists iff* $\mu_L^* \neq \mu_S^*$.

If equilibrium effort choices differ between agents, and $\Delta W > 0$, this results in a pay gap, as different levels of effort exertion imply different probabilities of winning the tournament. The probability of winning the tournament is equal to:

$$prob_i\left(q_i > q_j\right) = prob_i\left(\mu_i - \mu_j > \varepsilon_j - \varepsilon_i\right) \tag{2}$$

where $i \neq j$.

Given the above-mentioned distributional assumption, $E\left[\varepsilon_j - \varepsilon_i\right] = 0$, with the variance depending on the beliefs of the respective player (L, S). If $\sigma_L^2 < \sigma_S^2$, agent S underestimates the impact of effort on actual pay. As we shall see, the equilibrium investment in effort is a function of $prob_S\left(q_S > q_L\right) = g\left(\mu_S - \mu_L\right)$ for S and $prob_L\left(q_L > q_S\right) = h\left(\mu_L - \mu_S\right)$ for L. Here, $g\left(\mu_S - \mu_L\right)$ and $h\left(\mu_L - \mu_S\right)$ are the probability density functions of a normal distribution with zero mean and variance $2\sigma_S^2$ and $2\sigma_L^2$, respectively, so this is the channel through which the perceived variance of the luck component impacts effort choice.

From Lazear and Rosen (1981) we know that a decrease in the precision with which the agents understand the effort-reward relationship leads to reduced effort provision by risk averse agents. In the case of risk neutrality, however, this effect would be offset by an increase in the prize spread, assuming homogenous agents. Hence, in Lazear and Rosen (1981), for risk neutral agents with homogenous beliefs about the error term, the optimum investment in effort does not vary with the variance of the luck component.

Given that in our model we have two types of agents with heterogenous beliefs, this result does not hold. Prize spreads cannot be optimally adjusted for both groups at the same time,[2] and thus, even under risk neutrality, we expect a worse understanding of the effort-reward relationship to result in less effort. Accordingly, we assume risk neutrality, and continue with the analysis of the game.

Unlike agent S, agent L perfectly observes the variance (this assumption can be relaxed, as we simply require L's belief to be closer to the firm's reality than S's belief), but both agents assume identical information conditions, i.e., $\varepsilon_{x,y}$ and $\varepsilon_{y,y}$ for $x, y \in \{L, S\}$, $x \neq y$ have the same distribution. In this setup, a performance gap follows from Lazear and Rosen (1981). The proof of the following proposition is in Appendix A.

Proposition 1 *If $\sigma_L^2 < \sigma_S^2$ and both agents assume that each of them faces identical information conditions, a performance gap exists s.t. $\mu_S^* < \mu_L^*$.*

Agent S does not invest efficiently and $\mu_S^* < \mu_L^*$, i.e., the equilibrium investment in effort of L is greater than that of S. This is due to the fact that S underestimates the responsiveness of pay to effort, whereas L knows the correct distribution of the luck component. Hence, there is a *performance gap* in equilibrium, and S is less likely to win the tournament than L.

Note that for the above result to hold, we do not require the firm to adjust the prize spread ΔW optimally for any one of the agents (cf. proof of Proposition 1): as can be seen in (A.11), a performance gap persists irrespective of the exact prize spread as long as the latter is the same for both agents. For our experiment, this implies that the theory can be tested using a fixed prize scheme for the tournament.

2.2 Impact of Awareness of Informational Differences on Beliefs

So far, we have assumed that the agents have different beliefs due to different data samples, but are unaware of these differences, which is equivalent to each of them assuming identical information conditions. In the first treatment of our experiment, we induce the assumption of identical information conditions by letting L- and S-players actually compete with equally informed counterparts, and compare the performance difference between the two types of players across pairs. In the second

[2] We take this to be a realistic assumption given that legal constraints typically keep employers from discriminating according to the traditional model. In addition, the constancy of the prize spread for all agents also reflects our experimental setup.

treatment, each player competes with a differentially informed counterpart, and is aware of that. To address the latter situation, we now turn to the analysis of the performance gap when the agents are increasingly aware of the informational differences. In particular, we consider two possibilities as to how agents react to the fact that they have more or less information than their counterparts. First, they might rationally incorporate such insights, which leads to converging beliefs, and economize on their effort levels accordingly. Second, the mere fact that information conditions are not the same for both agents might evoke feelings of injustice, the implications of which we analyze by adopting a specification of inequity aversion (Fehr and Schmidt 1999).

2.2.1 Stepwise Convergence of Beliefs

Consider, first, the case where only agent S assumes the distribution of the luck component to be homogenous for both agents, i.e., agent S still overestimates the variance of the luck component, but this time agent L is aware of it. Thus, agent L knows both perceived distributions of the respective luck components, which are independent such that $\varepsilon_{L,L} \sim N\left(0, \sigma_L^2\right)$ and $\varepsilon_{S,L} \sim N\left(0, \sigma_S^2\right)$ where $\sigma_L^2 < \sigma_S^2$. In the next proposition, we show that, given an upper bound on W_2 (the loser prize), a performance gap persists, i.e., $\mu_S^* < \mu_L^*$,[3] but the gap is smaller than the one in the baseline case of Proposition 1.

Next, assume that S-players recognize that they have a noisier notion of the effort-reward relationship, so they understand that their estimate of the variance of the luck component is actually an upper bound, and update their beliefs about the variance of their luck component to some lower variance $\widetilde{\sigma}_S^2$ such that $\sigma_L^2 < \widetilde{\sigma}_S^2 < \sigma_S^2$. If L-players are unaware of this learning process on the part of the S-players, the performance gap becomes even smaller.

In the course of learning dynamics (e.g., L-players could correct their estimate of μ_S and adjust their effort accordingly), the tendency will be for the performance gap to shrink the more L- and S-players learn about each other's notions of the effort-reward relationship – until the agents' beliefs eventually converge and the performance gap vanishes.

The above discussion can be summarized in the following proposition, the proof of which is in Appendix A.

Proposition 2a *If $\sigma_L^2 < \sigma_S^2$ but the players partially update their beliefs such that (1) they no longer assume identical information conditions and (2) $\sigma_L^2 < \widetilde{\sigma}_S^2 < \sigma_S^2$ for S, a performance gap exists, so*

[3] Note that the equilibrium investments in effort μ_L^* and μ_S^* are different from those in Proposition 1.

that $\mu_S^* < \mu_L^*$ *as long as* $W_2 < \sigma_L \left(4\sigma_L \ln \frac{\widetilde{\sigma}_S}{\sigma_L} - \sqrt{\pi}V\right)$, *and both agents invest inefficiently. With full updating and* $\widetilde{\sigma}_S^2 = \sigma_L^2$ *for* S, *the performance gap vanishes.*

2.2.2 Informational Injustice and Fairness Considerations

Given our proposed information conditions (with one group being worse informed than the other), it is conceivable that particularly the informationally disadvantaged agents might deviate from the behavior laid out above, and – instead of acknowledging that their perceived variance of the luck component is an upper bound – feel discouraged, which would in turn lead to a worsening of the perceived effort-reward relationship. That is, the informationally disadvantaged group would implicitly derive disutility in the form of a *higher* perceived variance of the luck component, discouraging effort and therefore smothering economic prospects. For simplicity, we assume that L does not exhibit positive inequity aversion (which is an extreme case of the typical assumption that agents care more about negative than about positive inequity). Then, denote by $\overline{\sigma}_S$ agent S's perceived variance incorporating the option to rationally update her beliefs, but adjusted by a penalty (leading to a higher variance) due to negative inequity aversion:

$$\begin{aligned} \overline{\sigma}_S^2 &\equiv \sigma_S^2 + \alpha \max\left\{\sigma_S^2 - \widetilde{\sigma}_S^2, 0\right\} - \beta \max\left\{\sigma_S^2 - \widetilde{\sigma}_S^2, 0\right\} \\ &= \sigma_S^2 + (\alpha - \beta) \max\left\{\sigma_S^2 - \widetilde{\sigma}_S^2, 0\right\}, \end{aligned} \tag{3}$$

where α, $\beta \in \mathbb{R}^+$ are weights for negative inequity aversion and rational variance correction, respectively, and $\sigma_L^2 \leq \widetilde{\sigma}_S^2 < \sigma_S^2$ with $\widetilde{\sigma}_S^2$ as S's updated estimate of the actual variance σ_L^2.

We have already covered the cases where $\alpha \leq \beta$ in Proposition 2a, so we are left with $\alpha > \beta$, which implies that $\overline{\sigma}_S^2 > \sigma_S^2$. In the next proposition, we demonstrate that even with full updating the performance gap widens if S exhibits negative inequity aversion and L does not take it into account. This setup can be interpreted as follows: S has a small data sample suggesting some $\sigma_S^2 > \sigma_L^2$, but – given that she realizes that her perceived variance of the luck component is an upper bound – fully updates her beliefs to $\widetilde{\sigma}_S^2 = \sigma_L^2$. S exhibits negative inequity aversion because she knows that the provided information is less precise. This leads to the following proposition, with the corresponding proof in the Appendix.

Proposition 2b *If* $\sigma_L^2 < \sigma_S^2$ *but (1) the players fully update their beliefs such that* $\widetilde{\sigma}_S^2 = \sigma_L^2$ *for* S *and (2)* S *exhibits negative inequity aversion such that* $\overline{\sigma}_S^2 > \sigma_S^2$, *which* L *is not aware of, then a performance gap exists, so that* $\mu_S^* < \mu_L^*$, *and exceeds the one in Proposition 1, where both agents assume that each of them faces identical information conditions.*

As in Section 2.2.1, the long-run tendency of the performance gap will be to shrink when L becomes aware of S's inequity aversion. However, this process hinges on L's estimate of $\overline{\mu}_S^*$ based on L's belief about $\overline{\sigma}_S^2$ – an intricate updating process that involves much uncertainty. While the theoretical performance gap would be subject to further specification of both players' updated beliefs, we have shown that if negative inequity aversion is at play, this may result in a larger performance gap than in the baseline case of Proposition 1.

We conclude that, except in the case of full updating described in Proposition 2a, overestimating the variance of the luck component leads to underinvestment in effort and – in expectation – to reduced pay. Hence, the actual incentive effect of a rank-order tournament is deemed to be weak for agents who have less precise information on the relationship between effort and pay. This has the following experimentally testable implication:

Implication 1 *If the amount of available information on the luck component impacts the agents' perceived variance of the luck component, then informationally disadvantaged agents (S-players) perform worse than informationally advantaged agents (L-players).*

As discussed in this subsection, when agents do not assume identical information conditions, the effect on the gap in effort exertion depends on whether informationally disadvantaged agents rationally update their beliefs about the effort-reward relationship, or feel treated unfairly. This leads to:

Implication 2a *In the absence of (negative) inequity aversion, the performance gap between informationally disadvantaged agents (S-players) and informationally advantaged agents (L-players) is largest if the agents are not aware of the informational differences, i.e., if they assume that each of them faces identical information conditions, and zero only if the agents are fully aware of these differences and, as a consequence, have the same beliefs.*

Alternatively, S-players might feel discouraged by the fact that they have less information than their counterparts, and subsequently exert less effort than in the baseline case where they assume identical information conditions.

Implication 2b *If the informationally disadvantaged agents (S-players) exhibit negative inequity aversion, the performance gap between S-players and informationally advantaged agents (L-players)* *can be larger when the agents are aware of the informational differences than when they assume identical information conditions.*

3 Experimental Design and Procedures

The goal of our experiment is to test Implications 1, 2a, and 2b. In particular, we examine whether individuals who make inferences about the effort-reward relationship based on smaller samples perform worse in a tournament than their counterparts who make inferences based on larger samples.

Our experiment involved a tournament where subjects competed in pairs of two in a word-find task. Subjects were confronted with a matrix containing letters: most letters appeared in random order but some formed words. Provided with a list of 20 words, subjects could find these words in the matrix by marking sequences of letters horizontally, vertically, or diagonally. An example of such a letter matrix is given in Appendix B. The subjects' task was to mark as many words as possible in three minutes. For every correct word marked, subjects received 10 points. In addition to their task score, subjects received a random bonus, i.e., their luck component.

Bonuses were drawn from a hat containing 18 different bonus values which, in turn, were drawn from a uniform distribution. In every round, subjects were informed of the highest and the lowest possible bonus value in the distribution, plus they were given a randomly drawn sample of actual bonus values in the hat (without replacement). The sample was either large (for L-players), namely 12 out of 18 possible bonus values, or small (for S-players), namely 3 out of 18 bonus values. Being worse informed about the nature of the random bonus, S-players should have a less precise understanding of the effort-reward relationship than L-players.

The final score was equal to the sum of the task score and the value of the bonus. The subject with the higher final score in a pair won the tournament and received a large prize (10 tokens), whereas the subject with the lower score received a small prize (2 tokens). One token was worth $1.

To examine the role of knowledge of one's competitor's information condition, we created two treatments: the baseline where subjects assumed identical information conditions, and a second treatment where they did not (and were aware of the informational differences). To assure that subjects indeed believed that information conditions were identical in the baseline treatment, we in fact created identical information conditions: both subjects in a pair received information on either the small or the large sample of bonus values in the hat, and we informed all subjects of that. We refer to the baseline treatment as *identical-information tournament*. In the second treatment, one person per pair received information on the large sample and the other person in the pair received information on the small sample, and this was common knowledge. We refer to this as *different-information tournament*.

After an initial practice round, the task was repeated four times (with a different letter matrix and word list in every round). Subjects remained in the same pair for the duration of the experiment. In rounds 1 and 2, subjects were confronted with a wide range of potential bonus values from 0 to 100. In rounds 3 and 4, we decreased the range of bonus values by limiting them to be between 30 and 70. Performance is likely responsive to both experience with the task and the range of potential bonus values (as predicted by the theory). At the end of each round, subjects were informed of their task score, their final score, their counterpart's final score, and the tokens they won. They did not receive information on their counterpart's task score, and were thus unable to determine with certainty whether they won/lost because of their counterpart's performance or the randomly drawn bonus. As an example, Appendix B provides the instructions for S-players in the different-information tournament.

Besides the subjects' performance on the word-find task, we collected three additional pieces of information. After subjects had been informed of the highest and the lowest possible bonus values, and had seen their random draw of sample values, we asked them to guess the mean bonus value and report a 90% confidence interval for their guess. Reporting a larger confidence interval is equivalent to perceiving a larger variance. Furthermore, after the completion of the main experiment, we had subjects participate in a risky-choice task to measure their risk preferences, which could also explain differential effort choices.[4] We employed the task introduced by Holt and Laury (2002) with identical incentives. The resulting measures for risk aversion that we will use in the empirical analysis (in Section 4) are binary variables for the categories from 1 (very risk loving) to 10 (very risk averse). Finally, the study concluded with a short questionnaire collecting demographic information: in the empirical analysis, we will use binary variables for the categories capturing the subjects' economic background, namely from 1 (very poor) to 6 (very rich), matching the six options given in the questionnaire. We also create dummy variables for gender and for whether a subject was a student at the time of the experiment.

Subjects were paid for their performance in all four rounds and, in addition, received their earnings from the risky-choice task. Average earnings, including a $10 show-up fee, were about $36

[4]Given the uncertain nature of the tournament, more risk averse players are hypothesized to exert less effort and thus score lower.

for a study that lasted one hour.

We ran the experiments in the Harvard Decision Science Laboratory in the spring of 2010. 206 subjects participated in nine sessions with 22 or 24 subjects in each of them, and we have valid score data for 812 individual outcomes.[5]

4 Results

We first report descriptive statistics. Then, we examine our central prediction, Implication 1: agents who are provided with a smaller sample of information on possible bonus values (S-players) perform worse than their counterparts with more precise information (L-players), and this effect is due to the perceived variance of the bonus component. Finally, we test Implications 2a and 2b, i.e., whether the performance gap varies depending on whether subjects assume identical information conditions (in the identical-information tournament) or are aware of informational differences (in the different-information tournament).

[Insert Figure 1 about here]

On average, subjects found 10.13 words (with a standard deviation of 3.88) out of a total of 20 words available in a given letter matrix. Women and men differed slightly in their performance, with women marking 10.35 words correctly and men finding 9.78 words on average ($p < 0.05$). This difference was entirely driven by performance in the first round, and women and men did not differ at all in their performance in the remaining three rounds. Figure 1 presents the distribution of the number of words people found in the pooled sample. Typical outcomes ranged from 5 to 16 words per matrix. Four participants, i.e., roughly 2% of our subjects, found the maximum of 20 words in at least one round.

Examining Implication 1, we first review differences in the mean number of words found by L- and S-players. Table 1 reports the data pooled across both treatments (cf. first panel) and separately for each treatment condition (cf. second and third panels). Within each panel, in the first row we present performance levels aggregated over all four rounds, in the second for the wide-range rounds (rounds 1 and 2), in the third for the narrow-range rounds (rounds 3 and 4), and in the last row for the rounds where people had already gained one round's experience within a given range condition (rounds 2 and 4). L-players found about one word more than S-players on average ($p < 0.01$).

[5] We dropped all scores of a subject *after* incidents involving IT or other problems during the experiment, which explains the loss of 12 out of 824 outcomes.

[Insert Table 1 about here]

Considering identical-information and different-information tournaments separately in the second and third panels, the performance gap between L- and S-players is exacerbated in the different-information tournament. In the third panel, L-players found 1.3 words more than S-players on average, which corresponds to an increase of one-third of a standard deviation. Experience increased the performance gap to 1.8 words in rounds 2 and 4. On average (cf. first row of the second and third panels), the performance gap was mainly driven by S-players who performed significantly worse when competing against L-players rather than against identically informed counterparts ($p < 0.05$). In contrast, the better informed group was not differentially affected by the two treatment conditions.

[Insert Table 2 about here]

People's scores improved over time, but no clear learning pattern is observable (see Table 2, which presents mean scores by round). In particular, scores decreased between rounds 2 and 3 for S-players in the identical-information and for L-players in the different-information tournament, refuting simple learning but suggesting the existence of adjustment costs to the new bonus range in round 3. We do not assign particular importance to this, other than noting that learning alone cannot explain the dynamics we observe. On average, performance levels were significantly higher in the narrow-bonus-range rounds 3 and 4 as compared to the wide-range rounds 1 and 2, which is in line with the theory. Also, performance increased from round 1 to 2 and from round 3 to 4, reflecting learning within each range condition.

To more precisely measure these effects and test whether the performance difference between L- and S-players was driven by their perceptions of the variance of the bonus component (as assumed in our model and in Implication 1), we run regressions (cf. Tables 3a and 3b). We approximate the perceived variance by subjects' reported confidence intervals of their estimates of the mean bonus.[6] As the range of potential bonus values differed between the first two and the last two rounds, we use a standardized measure of the reported confidence interval between 0 and 1 – the reported confidence interval over the actual range (100 in the first two and 40 in the last two rounds) – to make it comparable across all rounds. This yields *Perceived range* with a mean of 0.46 and a standard deviation of 0.25. On average, L-players perceived the range of bonus values to be 0.41, whereas

[6]We dropped 44 observations by subjects who gave responses that were invalid by reasonable standards, e.g., their reported confidence interval exceeded the range of potential bonus values. This explains the discrepancies in sample sizes between Tables 3a (3b) and 4. All results in Table 4 are robust to limiting its samples to the corresponding samples in Table 3a.

S-players reported the range to be 0.50 (with perceived, normalized mean bonus values of 0.51 and 0.45, respectively).[7] To more easily interpret the effect on performance, we include 1−*Perceived range* in our regressions, as subjects should exert more effort the smaller they perceive the variance to be.

A simple test of the theory would be to estimate the impact of 1−*Perceived range* on task performance, and we expect that impact to be positive. However, the reported confidence intervals used to construct *Perceived range* are the result of some intellectual effort, and are thus potentially subject to endogeneity with respect to task performance. Indeed, the p-values from Durbin-Wu-Hausman tests (reported at the bottom of Table 3b) indicate that simple OLS estimates would not be consistent. Therefore, we test Implication 1 using two-stage least-squares regressions. A valid first stage in line with Implication 1 requires that receiving a large sample of possible bonus values, which is exogenous by design, have a strong, positive impact on 1−*Perceived range*. Then, for the second stage, we hypothesize that the projected values of 1−*Perceived range* are associated with higher task scores. Tables 3a and 3b report the results of the first and second stage, respectively.

[Insert Table 3a about here]

[Insert Table 3b about here]

Our results support Implication 1. As can be seen in Table 3a, L-players indeed had a more precise understanding of the effort-reward relationship than S-players: receiving a larger sample of information implies attaching a lower variance to the value of the bonus, which – as suggested by Table 3b – encourages effort and leads to higher task performance. The first stage is strong, and the effect of receiving a larger sample on 1−*Perceived range* is, on average, equal to two-fifths of a standard deviation.[8] The results hold controlling for risk aversion.[9]

[7] Note that, theoretically, different perceptions of the mean bonus value do not impact effort choice, as all players know that they receive draws from the same urn and the values of the random components cancel each other out in the expected utilities (cf. (A.1) and (A.2)). Furthermore, all of the regressions presented in this paper are robust to including the perceived mean bonus values, which have no explanatory power for performance.

[8] While the positive effect of the narrow range of potential bonus values on performance in rounds 3 and 4 in the second stage is in line with the theory, the negative coefficient of the indicator for these rounds on 1−*Perceived range* in the first stage (cf. Table 3a) might be confusing. Subjects reported significantly smaller perceived *absolute* ranges in rounds 3 and 4 as opposed to rounds 1 and 2 (namely, 20.91 compared to 39.81). However, the perceived *relative* ranges reverse, as we divide perceived ranges by the size of the actual given ranges to make the measures comparable. We cannot use the absolute range in all our variables, although the respective coefficient in unreported regressions suggests a significantly smaller perceived absolute range in rounds 3 and 4, as that correlation is partly mechanical. Still, the correlation between the perceived relative and absolute ranges is 0.69. We therefore use the relative range in all regressions to avoid mechanical correlations. Also, our results are virtually unchanged if one drops the indicator for rounds 3 and 4 from the right-hand side of the regression specifications.

[9] The number of observations drops once we include the categories for risk aversion, as some responses were either incomplete or invalid by reasonable standards. However, all results in this paper are robust to dropping these

Having shown that informational differences impact performance through the perceived variance of the bonus component, we now discuss Implications 2a and 2b, namely whether the performance gap shrinks or widens when informational differences are public information. Based on the theoretical discussion in Section 2.2, we hypothesize that the difference in perceived variance in the different-information tournament is not the same as the one in the identical-information tournament.

Given the mean scores in Table 1, informational injustice and fairness considerations (as in Implication 2b) might affect performance directly rather than through a rational understanding of the effort-reward relationship.[10] Across all four rounds, the informationally disadvantaged S-players exerted less effort when competing against a player with a large sample of information, and found 9.21 words as compared to 10.14 words when competing against a counterpart of the same type (the difference is significant at the 2% level). In contrast, the behavior of the L-players was virtually unaffected by the treatment conditions.

This pattern of behavior corresponds with other theories and field observations: perceived informational injustice leading to a pessimistic outlook (e.g., Bénabou and Tirole 2006), lower-ability agents decreasing their effort in tournaments when they are informed of the ability distribution (Freeman and Gelber 2010), and less confident agents – namely women – working less hard when competing against men who are more confident of their ability (Gneezy et al. 2003).

To test the relevance of informational injustice (Implication 2b) as opposed to rational updating (Implication 2a), we interact the indicators for the different-information tournament ("Different info") and for receiving a large sample in the last two columns of Table 3a. The underlying rationale is that the gap in the perceived variance of the luck component, as approximated by *Perceived range*, between L- and S-players should be smaller in the different-information tournament if the players update their beliefs rationally. Similarly, the gap should be larger in the different-information tournament if the informational-injustice channel operates through the perceived variance of the luck component. However, as can be seen in the last two columns of Table 3a, the interaction effect Large sample × Different info is not significant. Neither Implication 2a where agents rationally update their beliefs about their relative informedness and adjust effort accordingly, nor Implication 2b where the informationally disadvantaged develop a bleak outlook on the effort-reward relationship as a response

observations irrespective of whether we control for risk aversion or not.

[10]Note that all subjects were informed that they received their random bonus from the same hat/distribution as their counterparts (cf. instructions for "Stage 2" in Appendix B), so any decrease in the S-players' effort is unlikely to be a response to the suspicion that the game were rigged against them. Against this background, a particularly strong effect of receiving more information in the different-information tournament would underscore the importance of information for beliefs, and its link to subsequent performance.

to perceived informational injustice, leading to a larger performance gap in the different-information than in the identical-information tournament, is supported.

[Insert Table 4 about here]

In order to explore whether there is any differential impact of the different- vs. identical-information treatment operating directly through the sample sizes rather than indirectly through *Perceived range*, we also consider the reduced-form estimation (i.e., regressing scores on the large-sample indicator and the remaining variables included in Table 3a) in Table 4 . Columns 1 to 3 demonstrate that our results regarding Implication 1 are robust to the inclusion of multiple controls: L-players outperformed S-players overall. Column 4 shows that while L-players outperformed S-players in the different-information tournament (the sum of the coefficients of Large sample and Large sample $\times$ Different info is significant at the 2% level), the performance difference – albeit positive – does not significantly exceed that in the identical-information tournament. Thus, rational updating in the absence of negative inequity aversion is unlikely to explain our findings, as the performance gap is not smaller in the different-information than in the identical-information tournament. Either the two channels of influence discussed in Propositions 2a and 2b do not matter, or they cancel each other out.

5 Concluding Remarks

This paper explores the impact of noise in people's perceptions of the effort-reward relationship on their performance in a tournament setting, and demonstrates how informational differences can translate into differences in performance. In our laboratory experiment, we implement a new mechanism to manipulate beliefs about the role of luck for tournament outcomes by varying the amount of information people received on the latter, building on the simple statistical idea that smaller samples are noisier than larger samples. We show that receiving more information on the role of luck improves the understanding of the effort-reward relationship, and leads to significantly better performance. This has broader implications, and could help explain how beliefs about one's initial conditions may influence one's future labor market outcomes.

Consider our findings in the context of a well-known labor market phenomenon, namely the (rigidity of the) underrepresentation of women in top management positions. Women only hold a small fraction of leadership positions in the corporate world (Bertrand and Hallock 2001). At the Fortune 500 companies in 2010, 2.4 percent of the CEOs, 14.4 percent of the executive officers, and

15.7 percent of the board members were female.[11] Most notably, women are also consistently more likely to attribute success to luck rather than individual effort (Fisman and O'Neill 2009).

Our paper suggests why this might be the case, and hints at a potential mechanism underlying the persistence of gender gaps at the top: when people (have to) source career-relevant information on the effort-reward relationship from similar others, women being in the minority in top management positions are at a disadvantage because the size of the group of similar others determines how precise the information received is. As a consequence, women might end up overestimating the importance of luck in the effort-reward relationship and, thus, put forth less effort in the workplace. This in turn affects their likelihood of success under performance pay schemes and eventual promotion in an organization.

The theoretical framework in this paper fits gender imbalances in organizations quite nicely, as gender gaps are most pronounced in senior positions characterized by competitive work environments where managers are involved in promotion tournaments with substantial uncertainty about how effort translates into rewards. As in our model, promotion tournaments involve unique prize schemes, e.g., wages are often defined for different career stages and hardly vary among individual employees within a given slate. Tournaments tend to be particularly harsh at the top of the wage distribution, given that the loser prize typically decreases across the wage distribution. An extreme example is the up-or-out system implemented by firms in very competitive industries – e.g., consulting, investment banking, or legal practices – and in academia such that candidates below a certain percentile in the performance ranking are dismissed (corresponding to a loser prize of $W_2 = 0$ in our model, which is reminiscent of Proposition 2a and the discussion in Section 2.2.1, where we have shown that an upper bound on W_2 is a sufficient condition for a performance gap), while the remaining employees are promoted.

For our model to apply in this context, some aspects of career-relevant information must be gender-specific.[12] Informal accounts suggest this is the case. The scarcity of senior colleagues of the same sex puts female junior managers at a disadvantage: junior women "have inadequate information about acceptable (or successful) modes of behavior..." (Blau et al. 2005, p. 177). Similarly, Ibarra (1992, p. 67) argued that "organizational demography" constrains women's available set of comparable others to learn from: "Women and minorities usually have a much smaller set of 'similar others' with whom to develop professional relationships based on identity-group homophily." Such

[11] http://www.catalyst.org/publication/132/us-women-in-business

[12] In addition, the career relevance of such information must be independent of whether the firm discriminates in any form against any group, or whether returns to information vary between groups, which seems plausible.

networks matter: examining the effectiveness of same-sex networks in a professional service firm where only a small minority of women held senior management positions, Ibarra (1993) found that men reaped greater benefits from their larger same-sex networks than women.

The patterns of behavior for the informationally advantaged and disadvantaged groups found in our experiment are compatible with other laboratory and field observations based on women and men. Gneezy et al. (2003) as well as Booth and Nolen (2009) present experimental evidence suggesting that the gender performance gap is particularly pronounced in mixed or male-dominated competitive environments as compared to same-sex competitions. Reminiscent of our findings, gender differences in performance were also driven by women – or, in our case, the informationally disadvantaged group – adjusting their behavior to the different environments: women performed better in same-sex than in mixed-sex competitions, while men's performance was not affected by the gender composition (Gneezy et al. 2003). A similar pattern has been found in performance evaluations in an organization where women were in the minority, namely among officers in the Israeli military: women were evaluated more positively the larger their relative share in a group was, whereas men's evaluations were invariant to the gender balance (Pazy and Oron 2001). More generally, our findings relate to earlier work in sociology and political science, "critical mass theory," suggesting the importance of relative group size for economic success (Kanter 1977).

Clearly, the gender balance in organizations may affect women's and men's productivity through a multitude of channels. For instance, a larger share of women in an organization might be correlated with a larger share of women in the talent pool of organizationally relevant professions, thereby increasing the firm's economic benefits of adjusting working conditions to women's needs (see Bertrand et al. 2010 for a discussion). In addition, an increased proportion of women in counter-stereotypical positions may also affect implicit biases, changing women's and men's beliefs about career trajectories (Beaman et al. 2009). More generally, differences in the evaluations of women and men based on the gender composition of the group are also compatible with statistical discrimination and information asymmetries where the employer is worse informed about the productivity of the minority group than of the majority group (Coate and Loury 1993), or where the minority group has invisible abilities (Milgrom and Oster 1987).

Our paper suggests an additional mechanism through which differences in performance, pay, and representation in leadership positions can emerge – informational differences due to the relative size of one's group. Organizational demography may thus be an important determinant of the productivity, promotion likelihood, and pay outcomes of an organization's employees.

References

[1] Alesina, A., E. Glaeser, and B. Sacerdote (2001), "Why Doesn't the United States Have a European-Style Welfare State?" *Brookings Papers on Economic Activity*, 2, 187-254

[2] Beaman, L., R. Chattopadhyay, E. Duflo, R. Pande, and P. Topalova (2009), "Powerful Women: Does Exposure Reduce Bias?" *Quarterly Journal of Economics*, 124(4), 1497-1540.

[3] Bénabou, R. J. and J. Tirole (2006), "Belief in a Just World and Redistributive Politics," *Quarterly Journal of Economics*, 121(2), 699-746.

[4] Bertrand, M., C. Goldin, and L. F. Katz (2010), "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors," *American Economic Journal: Applied Economics*, 2(1), 228-255.

[5] Bertrand, M. and K. F. Hallock (2001), "The Gender Gap in Top Corporate Jobs," *Industrial and Labor Relations Review*, 55(1), 3-21.

[6] Bertrand, M., E. F. P. Luttmer, and S. Mullainathan (2000), "Network Effects and Welfare Cultures," *Quarterly Journal of Economics*, 115(3), 1019-1055.

[7] Blau, F. D., M. A. Ferber and A. E. Winkler (2005), *The Economics of Women, Men, and Work*, Upper Saddle River, NJ: Prentice-Hall, 5th edition.

[8] Booth, A. L. and P. J. Nolen (2009), "Gender Differences in Risk Behavior: Does Nurture Matter?" *IZA Discussion Paper No. 4026.*

[9] Bull, C., A. Schotter, and K. Weigelt (1987), "Tournaments and Piece Rates: An Experimental Study," *Journal of Political Economy*, 95(1), 1-33.

[10] Celen, B. and K. Hyndman (forthcoming), "Social Learning Through Endogenous Information Acquisition: An Experiment," *Management Science.*

[11] Coate, S. and G. C. Loury (1993), "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" *American Economic Review*, 83(5), 1220-1240.

[12] Edin, P., P. Fredriksson, and O. Aslund (2003), "Ethnic Enclaves and the Economic Success of Immigrants – Evidence from a Natural Experiment," *Quarterly Journal of Economics*, 118(1), 329-357.

[13] Fehr, E. and K. M. Schmidt (1999), "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, 114(3), 817-868.

[14] Figlio, D. N., S. Hamersma, and J. Roth (2011), "Information Shocks and Social Networks," *NBER Working Paper No. 16930.*

[15] Fisman, R. and M. O'Neill (2009), "Gender Differences in Beliefs on the Returns to Effort: Evidence from the World Values Survey," *Journal of Human Resources*, 44(4), 858-870.

[16] Freeman, R. B. and A. M. Gelber (2010), "Prize Structure and Information in Tournaments: Experimental Evidence," *American Economic Journal: Applied Economics*, 2(1), 149-164.

[17] Gneezy, U., M. Niederle, and A. Rustichini (2003), "Performance in Competitive Environments: Gender Differences," *Quarterly Journal of Economics*, 118(3), 1049-1074.

[18] Holmes, M. A. and S. O'Connell (2007), "Leaks in the Pipeline: Why Do Women Remain Curiously Absent from the Ranks of Academia?" *Nature*, 446(7133), 346.

[19] Holt, C. A. and S. K. Laury (2002), "Risk Aversion and Incentive Effects," *American Economic Review*, 92(5), 1644-1655.

[20] Ibarra, H. (1992), "Homophily and Differential Returns: Sex Differences in Network Structure and Access in an Advertising Firm," *Administrative Science Quarterly*, 37(3), 422-447.

[21] Ibarra, H. (1993), "Personal Networks of Women and Minorities in Management: A Conceptual Framework," *Academy of Management Review*, 18(1), 56-87.

[22] Jackson, C. K. and H. S. Schneider (2010), "Do Social Connections Reduce Moral Hazard? Evidence from the New York City Taxi Industry," *NBER Working Paper No. 16279*.

[23] Kanter, R. M. (1977), *Men and Women of the Corporation*, New York: Basic Books.

[24] Lazear, E. P. and S. Rosen (1981), "Rank-Order Tournaments as Optimum Labor Contracts," *Journal of Political Economy*, 89(5), 841-864.

[25] Lundberg, S. J. and R. Startz (1983), "Private Discrimination and Social Intervention in Competitive Labor Market," *American Economic Review*, 73(3), 340-347.

[26] Milgrom, P. R. and S. Oster (1987), "Job Discrimination, Market Forces, and the Invisibility Hypothesis," *Quarterly Journal of Economics*, 102(3), 453-476.

[27] Munshi, K. (2003), "Networks in the Modern Economy: Mexican Migrants in the U.S. Labor Market," *Quarterly Journal of Economics*, 118(2), 549-599.

[28] Orrison, A., A. Schotter, and K. Weigelt (2004), "Multiperson Tournaments: An Experimental Examination," *Management Science*, 50(2), 268-279.

[29] Pazy, A. and I. Oron (2001), "Sex Proportion and Performance Evaluation among High-Ranking Military Officers," *Journal of Organizational Behavior*, 22(6), 689-702.

[30] Ragins, B. R. (1999), "Gender and Mentoring Relationships: A Review and Research Agenda for the Next Decade," in G. N. Powell (ed.), Handbook of Gender and Work, 347-370, Thousand Oaks, CA: Sage Publications.

Tables

Table 1: Differences in Mean Scores

Treatment	Rounds	Large sample ($N = 102$)	Small sample ($N = 104$)	Difference
All	All	10.637	9.632	1.005***
		[3.83]	[3.87]	[0.27]
All	1 & 2	10.134	9.203	0.931***
		[3.59]	[3.62]	[0.36]
All	3 & 4	11.146	10.063	1.083***
		[4.00]	[4.08]	[0.40]
All	2 & 4	11.280	9.966	1.314***
		[3.96]	[4.05]	[0.60]
Treatment	Rounds	Large sample ($N = 46$)	Small sample ($N = 48$)	Difference
Identical info	All	10.793	10.138	0.656
		[3.85]	[3.90]	[0.40]
Identical info	1 & 2	10.231	9.800	0.431
		[3.72]	[3.66]	[0.54]
Identical info	3 & 4	11.375	10.479	0.896
		[3.92]	[4.11]	[0.60]
Identical info	2 & 4	11.231	10.505	0.726
		[4.13]	[4.05]	[0.60]
Treatment	Rounds	Large sample ($N = 56$)	Small sample ($N = 56$)	Difference
Different info	All	10.509	9.205	1.304***
		[3.82]	[3.81]	[0.36]
Different info	1 & 2	10.055	8.696	1.358***
		[3.50]	[3.52]	[0.47]
Different info	3 & 4	10.964	9.714	1.249**
		[4.08]	[4.04]	[0.55]
Different info	2 & 4	11.321	9.509	1.812***
		[3.84]	[4.00]	[0.53]

Notes (Tables 1 and 2): In the first two columns, standard deviations are in parentheses. The third column indicates the results of a two-sided difference-in-means test (with standard errors in parentheses) where */**/*** denote significance at the 10%/5%/1% level, respectively.

Table 2: Differences in Mean Scores by Rounds

Treatment	Round	Large sample ($N = 102$)	Small sample ($N = 104$)	Difference
All	1	9.554	9.068	0.486
		[3.57]	[3.48]	[0.27]
All	2	10.720	9.337	1.383***
		[3.54]	[3.75]	[0.51]
All	3	10.439	9.524	0.915*
		[3.57]	[3.86]	[0.53]
All	4	11.840	10.602	1.238**
		[4.29]	[4.25]	[0.60]
Treatment	Round	Large sample ($N = 46$)	Small sample ($N = 48$)	Difference
Identical info	1	9.756	9.851	-0.096
		[3.64]	[3.62]	[0.76]
Identical info	2	10.696	9.750	0.946
		[3.78]	[3.73]	[0.78]
Identical info	3	10.953	9.681	1.273*
		[3.30]	[3.85]	[0.76]
Identical info	4	11.778	11.277	0.501
		[4.43]	[4.26]	[0.91]
Treatment	Round	Large sample ($N = 56$)	Small sample ($N = 56$)	Difference
Different info	1	9.393	8.411	0.982
		[3.54]	[3.25]	[0.64]
Different info	2	10.741	8.982	1.759**
		[3.35]	[3.77]	[0.68]
Different info	3	10.036	9.393	0.644
		[3.75]	[3.89]	[0.73]
Different info	4	11.891	10.036	1.855**
		[4.22]	[4.20]	[0.80]

Table 3a: Determinants of Perceived Variance of Bonus Component (First Stage)

	Dependent variable: 1−*Perceived range*			
Large sample	0.087***	0.104***	0.103**	0.107**
	[0.03]	[0.03]	[0.04]	[0.05]
Diff. info	0.026	0.002	0.041	0.004
	[0.03]	[0.03]	[0.04]	[0.04]
Large sample × Diff. info			-0.031	-0.005
			[0.06]	[0.07]
Rounds 2 & 4	-0.011	-0.010	-0.011	-0.010
	[0.01]	[0.01]	[0.01]	[0.01]
Rounds 3 & 4	-0.124***	-0.120***	-0.124***	-0.120***
	[0.01]	[0.01]	[0.01]	[0.01]
Female	-0.011	-0.041	-0.010	-0.041
	[0.03]	[0.03]	[0.03]	[0.03]
Student	-0.026	-0.076**	-0.024	-0.075**
	[0.04]	[0.03]	[0.04]	[0.03]
Economic-background FE	Yes	Yes	Yes	Yes
Risk-aversion FE	No	Yes	No	Yes
# of observations	768	622	768	622
First-stage F-statistic	9.47	10.45		

Notes: */**/*** denote significance at the 10%/5%/1% level, respectively. In the linear regressions, standard errors are given in parentheses, and are clustered at the pair level. Self-reported economic background is scaled from 1 (very poor) to 6 (very rich), and included in the regressions as six dummy variables. Risk aversion is measured on a scale from 1 (very risk loving) to 10 (very risk averse), and included in the regressions as ten dummy variables. The different-information tournament, under which informational differences are public information, is labeled as "Diff. info."

Table 3b: Determinants of Task Performance (Second Stage)

	Dependent variable: *Score*	
1−*Perceived range*	11.142*	11.419**
(endogenous)	[6.22]	[5.88]
Diff. info	-0.845	-0.525
	[0.57]	[0.52]
Rounds 2 & 4	1.032***	1.118***
	[0.23]	[0.24]
Rounds 3 & 4	2.254***	2.122***
	[0.84]	[0.78]
Female	0.508	0.860
	[0.50]	[0.55]
Student	2.051***	2.468***
	[0.61]	[0.73]
Economic-background FE	Yes	Yes
Risk-aversion FE	No	Yes
# of observations	768	622
Durbin-Wu-Hausman test (p-value)	0.001	0.000

Notes: */**/*** denote significance at the 10%/5%/1% level, respectively. In the linear regressions, standard errors are given in parentheses, and are clustered at the pair level. Self-reported economic background is scaled from 1 (very poor) to 6 (very rich), and included in the regressions as six dummy variables. Risk aversion is measured on a scale from 1 (very risk loving) to 10 (very risk averse), and included in the regressions as ten dummy variables. The different-information tournament, under which informational differences are public information, is labeled as "Diff. info." The corresponding first-stage regressions are reported in the first two columns of Table 3a.

Table 4: Determinants of Task Performance (Reduced Form)

	Dependent variable: *Score*			
Large sample	1.005**	1.007**	1.117**	0.498
	[0.45]	[0.45]	[0.51]	[0.68]
Diff. info			-0.541	-1.179*
			[0.48]	[0.67]
Large sample × Diff. info				1.217
				[0.92]
Rounds 2 & 4		0.972***	1.080***	1.080***
		[0.19]	[0.21]	[0.21]
Rounds 3 & 4		0.932***	0.771***	0.767***
		[0.16]	[0.19]	[0.19]
Female			0.326	0.290
			[0.49]	[0.49]
Student			1.696***	1.602***
			[0.56]	[0.56]
Constant	9.632***	8.679***		
	[0.34]	[0.34]		
Economic-background FE	No	No	Yes	Yes
Risk-aversion FE	No	No	Yes	Yes
# of observations	812	812	637	637

Notes: */**/*** denote significance at the 10%/5%/1% level, respectively. In the linear regressions, standard errors are given in parentheses, and are clustered at the pair level. Self-reported economic background is scaled from 1 (very poor) to 6 (very rich), and included in the regressions as six dummy variables. Risk aversion is measured on a scale from 1 (very risk loving) to 10 (very risk averse), and included in the regressions as ten dummy variables. The different-information tournament, under which informational differences are public information, is labeled as “Diff. info.”

Figures

Figure 1: Histogram of Scores (Pooled)

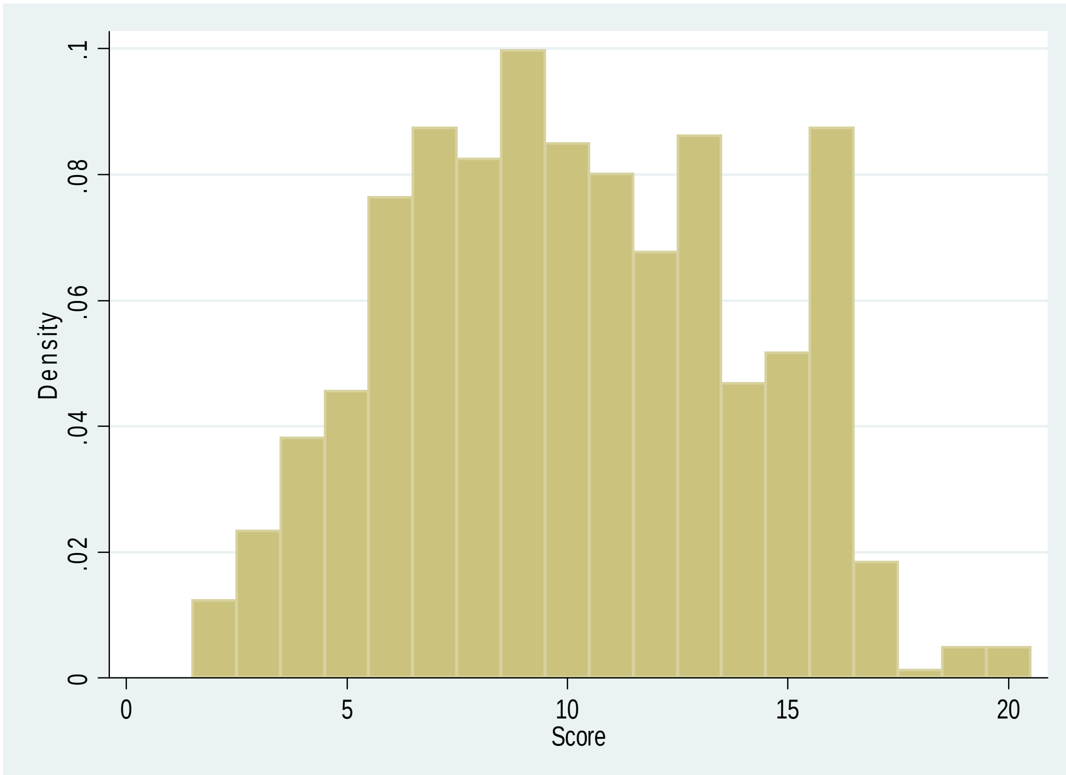

Appendix A

Proofs

Proof of Proposition 1 (follows from Lazear and Rosen 1981) The optimum investment in effort μ_i^* $(i = L,\ S)$ will be a function of ΔW, the prize spread, and the variance of the net dose of bad luck ($\varepsilon_{L,S} - \varepsilon_{S,S}$ and $\varepsilon_{S,L} - \varepsilon_{L,L}$ for S and L, respectively), which equals $2\sigma_S^2$ for S and $2\sigma_L^2$ for L. Under risk neutrality, S's and L's expected utilities are:

$$W_2 + prob_S\left(q_S > q_L\right)\Delta W - C(\mu_S) \tag{A.1}$$

and

$$W_2 + prob_L\left(q_L > q_S\right)\Delta W - C(\mu_L). \tag{A.2}$$

Then, one yields the following FOCs for S and L, respectively:

$$\Delta W g\left(\mu_S - \mu_L\right) = C'(\mu_S) \tag{A.3}$$

and

$$\Delta W h\left(\mu_L - \mu_S\right) = C'(\mu_L), \tag{A.4}$$

where $g\left(\mu_S - \mu_L\right)$ is the probability density function of a normal distribution with zero mean and variance $2\sigma_S^2$, and $h\left(\mu_L - \mu_S\right)$ is the probability density function of a normal distribution with zero mean and variance $2\sigma_L^2$.

Given normality of the error terms in the output equations, one knows that $g(\cdot)$ and $h(\cdot)$ are symmetric. Also, each agent believes that both contestants are homogenous, so the beliefs of S and L reflect the standard Nash-Cournot case s.t.:

$$g(0) < h(0). \tag{A.5}$$

Turning to the principal, she has the same beliefs as the majority group to which agent L belongs, leading the principal to assume the same optimization problem for all agents:

$$\Delta W h(0) = C'(\mu). \tag{A.6}$$

Now that the firm makes zero profit in expectation, it holds that:

$$2V\mu^* = W_1 + W_2, \tag{A.7}$$

where V is the marginal social return and μ^* is the solution to (A.6).

Also, the firm has to make the right choice with regard to the prize spread even to achieve zero profit in expectation. Hence, the principal chooses ΔW such that (given her beliefs):

$$\Delta W^* = \arg\max\left\{\tfrac{1}{2}W_1 + \tfrac{1}{2}W_2 - C(\mu^*)\right\} = \arg\max\left\{V\mu^* - C(\mu^*)\right\}. \tag{A.8}$$

Recall that μ^* depends on ΔW, so the corresponding FOC is:

$$(V - C'(\mu^*))\frac{\partial \mu^*}{\partial \Delta W} = 0 \Leftrightarrow V = C'(\mu^*). \tag{A.9}$$

Combining (A.6), (A.7), and (A.9), one yields:

$$\Delta W^* = \frac{V}{h(0)} = \frac{W_1+W_2}{2\mu^* h(0)}. \tag{A.10}$$

Given that $C(\cdot)$ is quadratic, we can finally conclude from (A.3), (A.4), and (A.10) in conjunction with (A.5) that:

$$C'(\mu_S^*) = \frac{(W_1+W_2)g(0)}{2\mu^* h(0)} < \frac{(W_1+W_2)h(0)}{2\mu^* h(0)} = C'(\mu_L^*) = C'(\mu^*) = V. \tag{A.11}$$

Thus, we have that $\mu_S^* < \mu_L^* = \mu^*$. ■

Proof of Proposition 2a Take as given the best-response effort choices μ_L^* and μ_S^*. First, we show that the upper bound on W_2 is equivalent to an upper bound on μ^* and, thus, on $\mu_L^* - \mu_S^*$:

$$\begin{aligned} &W_2 < \sigma_L\left(4\sigma_L \ln\frac{\sigma_S}{\sigma_L} - \sqrt{\pi}V\right) \Leftrightarrow \sigma_L\sqrt{\pi}V + W_2 < -4\sigma_L^2 \ln\frac{\sigma_L}{\sigma_S} \\ &\Leftrightarrow \sqrt{\sigma_L\sqrt{\pi}V + W_2} < \sqrt{-4\sigma_L^2\ln\frac{\sigma_L}{\sigma_S}} \Leftrightarrow \sqrt{\frac{V}{2h(0)} + W_2} \\ &< \sqrt{-4\sigma_L^2\ln\frac{\sigma_L}{\sigma_S}} \Leftrightarrow \sqrt{\frac{\Delta W^*}{2} + W_2} < \sqrt{-4\sigma_L^2\ln\frac{\sigma_L}{\sigma_S}}, \end{aligned} \tag{A.12}$$

as the principal chooses the prize spread optimally (cf. (A.10)). Now, given a quadratic cost function, one knows from (A.9) and (A.10) that $\mu^* = \sqrt{\frac{W_1+W_2}{2}}$. So, the above inequality translates to $\mu^* < \sqrt{-4\sigma_L^2\ln\frac{\sigma_L}{\sigma_S}}$, from which one can infer that $\mu_L^* - \mu_S^* < \mu_L^* < \sqrt{-4\sigma_L^2\ln\frac{\sigma_L}{\sigma_S}}$ because $\mu_L^* < \mu^*$ due to (A.11) and the symmetry of $h(\cdot)$ around zero.

In the next step, we show that the latter condition is equivalent to $g(0) < h\left(\mu_L^* - \mu_S^*\right) \Rightarrow \mu_S^* < \mu_L^* < \mu^*$ as dictated by the logic of (A.11). Note that the alternative assumption solely impacts the decision problem of agent L:

$$\begin{aligned} &\mu_L^* - \mu_S^* < \sqrt{-4\sigma_L^2\ln\frac{\sigma_L}{\sigma_S}} \Leftrightarrow 4\sigma_L^2\ln\frac{\sigma_L}{\sigma_S} < -\left(\mu_L^* - \mu_S^*\right)^2 \\ &\Leftrightarrow \left(\frac{\sigma_L}{\sigma_S}\right)^{4\sigma_L^2} < \exp -\left(\mu_L^* - \mu_S^*\right)^2 \Leftrightarrow \frac{\sigma_L}{\sigma_S} < \exp\frac{-\left(\mu_L^*-\mu_S^*\right)^2}{4\sigma_L^2} \\ &\Leftrightarrow \frac{1}{\sqrt{4\pi\sigma_S^2}} < \frac{1}{\sqrt{4\pi\sigma_L^2}}\exp\frac{-\left(\mu_L^*-\mu_S^*\right)^2}{4\sigma_L^2} \Leftrightarrow g(0) < h\left(\mu_L^* - \mu_S^*\right). \end{aligned} \tag{A.13}$$

From this, it follows that $\mu_S^* < \mu_L^* < \mu^*$.

The result that this time both agents invest inefficiently is due to the fact that L knows the correct distribution of the luck component, but also incorporates S's belief in her best response. However, the actual performance gap has *decreased* compared to the baseline case in Proposition 1, as the distance between $h(x)$ and $g(x)$ is maximized for $x = 0$.

The above proof holds true for any $\sigma_S^2 > \sigma_L^2$ and, thus, also for $\widetilde{\sigma}_S^2 < \sigma_S^2$. Denote by $\widetilde{g}(\cdot)$ the probability density function of a normal distribution with some lower variance $\widetilde{\sigma}_S^2$ s.t. $\sigma_L^2 < \widetilde{\sigma}_S^2 < \sigma_S^2$, and the S-players' equilibrium investment in effort becomes μ_S^{**}. Now, if $W_2 < \sigma_L\left(4\sigma_L\ln\frac{\widetilde{\sigma}_S}{\sigma_L} - \sqrt{\pi}V\right)$

and L-players are unaware of this learning process on the part of the S-players, the performance gap is characterized by the difference $h\left(\mu_L^* - \mu_S^*\right) - \widetilde{g}(0) > 0$ where $\widetilde{g}(0) > g(0)$ and $\mu_S^* < \mu_S^{**}$, so the performance gap remains but becomes smaller.

With full updating, S-players will update their beliefs s.t. $\widetilde{\sigma}_S^2 = \sigma_L^2$, and L-players will be aware of this, so that $\mu_S^* = \mu_L^* = \mu^*$. ■

Proof of Proposition 2b L's beliefs reflect the standard Nash-Cournot case, whereas S – given her negative inequity aversion – feels discouraged and powerless, as reflected by a *higher* perceived variance of the luck component. Thus, the performance gap is characterized by the difference $h(0) - \overline{g}(\mu_L^* - \overline{\mu}_S^*)$ where $\overline{g}(\cdot)$ is the probability density function of a normal distribution with variance $\overline{\sigma}_S^2$ and $\overline{\mu}_S^*$ is S's optimal effort given μ_L^* and $\overline{g}(\cdot)$. Hence, as long as L knows that S is worse informed (but does not incorporate S's inequity aversion), there will be a performance gap because $h(0) - \overline{g}(\mu_L^* - \overline{\mu}_S^*) > 0$.

Finally, to see that the performance gap exceeds the one in Proposition 1:
$\frac{1}{\sqrt{4\pi\sigma_S^2}} > \frac{1}{\sqrt{4\pi\overline{\sigma}_S^2}} \exp \frac{-\left(\mu_L^* - \overline{\mu}_S^*\right)^2}{4\overline{\sigma}_S^2} \Rightarrow h(0) - \overline{g}(\mu_L^* - \overline{\mu}_S^*) > h(0) - g(0)$. ■

Appendix B

Letter Matrix (Example: Nations of the World)

B	T	U	W	T	T	B	P	M	S	K	L	L	L	T	W	Q	N	B	V	ALGERIA
O	O	M	E	X	J	A	E	B	K	I	D	A	M	A	R	B	H	W	Y	BELGIUM
F	F	H	N	G	N	G	Y	O	M	U	H	J	I	U	T	K	U	B	W	CANADA
W	C	P	I	A	Y	L	S	T	A	K	C	S	A	R	I	V	Y	B	P	EGYPT
S	Z	E	M	U	T	P	E	V	I	R	T	K	Q	P	E	G	I	M	X	FINLAND
R	A	A	A	L	S	B	T	U	U	A	Z	Q	D	A	A	G	L	A	Q	GREECE
D	N	A	L	A	E	Z	W	E	N	I	X	W	D	V	O	N	L	E	F	HONDURAS
S	B	V	V	T	W	U	A	I	R	N	X	N	J	U	W	R	I	A	B	INDONESIA
N	G	Y	K	H	Y	D	F	T	B	E	A	R	C	B	V	V	X	G	H	JAPAN
S	F	T	V	A	R	Y	B	R	L	W	Y	A	O	I	I	N	T	M	Y	KOREA
A	G	D	D	I	M	D	G	U	R	A	U	S	E	D	E	C	J	A	U	LATVIA
R	U	W	P	L	B	G	Z	X	G	Q	M	T	F	M	A	J	I	X	W	MALTA
U	D	B	N	A	C	O	T	S	Z	U	N	P	E	N	U	S	D	F	H	NEW ZEALAND
D	H	N	X	N	O	I	D	B	H	A	R	Y	A	P	E	K	I	I	U	PANAMA
N	O	C	P	D	O	M	X	V	M	G	P	D	G	N	O	N	V	R	K	RWANDA
O	Q	L	L	J	M	H	I	B	R	A	A	H	O	R	L	Q	T	T	R	SINGAPORE
H	F	B	V	G	R	E	E	C	E	M	V	D	E	A	J	Q	B	L	P	THAILAND
A	B	I	T	L	K	P	H	C	W	B	N	A	N	P	J	P	G	U	O	UKRAINE
C	X	Q	U	G	I	Z	F	V	J	I	C	D	O	Z	K	N	T	I	V	VIETNAM
B	W	Y	L	D	I	S	I	N	G	A	P	O	R	E	V	I	D	K	M	YEMEN

Experimental Instructions

You are participating in a study in which you will earn some money. The amount will depend on how well you do in a task plus a bonus (described below). At the end of the study, your earnings (1 token = $1) will be added to a show-up fee, and you will be paid in cash.

Main task We will show you matrices containing letters. Some letters appear in random order and some form words. Words can be found by combining letters next to each other horizontally (moving from left to right or from right to left), vertically (moving from top to bottom or from bottom to top), or diagonally (moving from left to right or from right to left). A list of all words contained in a given matrix is displayed next to each matrix. You will be shown a matrix for 3 minutes. You are asked to find as many words from the list as possible. Your point score for the task is calculated as follows:

- For every correct word marked in the matrix, 10 points are added to your score.
- Words that are not marked receive no points.

You are randomly matched with another person present. You and your counterpart see the same letter matrix, and you are both asked to find as many words as possible.

Your *final point score* depends on your point score from the task plus a bonus corresponding to a number between X and Y (you will be informed of the values of X and Y in each round). The number will be randomly drawn from a hat containing 18 balls.

In order for you to get a better sense of the likely value of the bonus, we will first draw 3 balls from the hat at random and inform you of their values. We will then put the balls back into the hat.

In order for your counterpart to get a better sense of the value of the bonus, we will draw 12 balls from the same hat at random and inform your counterpart of their values.

Your *final point score* equals: TASK POINT SCORE + BONUS NUMBER

For example (choosing numbers outside of the range of possibilities in this study), if you identified 1,000 correct words and your counterpart identified 900 correct words, the value of your bonus (randomly drawn number) was 2,000 and the value of your counterpart's bonus (randomly drawn number) was 1,500, then your final point score would be: 10,000 + 2,000 = 12,000. Your counterpart's final point score would be: 9,000 + 1,500 = 10,500.

Calculation of payout The person getting the higher final point score in your pair will receive 10 tokens. The person with the lower final point score will receive 2 tokens.

How the study is conducted It is conducted in five stages.

Stage 1 You will be informed of the lowest and the highest possible bonus (X and Y) contained in a given hat and the value of the balls randomly drawn from the hat. We will ask you to give us your best guess of the average value of all balls in the hat.

Stage 2 We will draw a ball at random for you and then put the ball back into the hat. This ball is your bonus, i.e., the value of the ball counts towards your final score. From the same hat, we will also draw a bonus at random for your counterpart and then return the ball into the hat. You will be informed of your bonus after completing the word-find task.

Stage 3 You will complete the word-find task.

Stage 4 We will then calculate your scores and inform you of your final score (i.e., your score for the task plus the value of your bonus) and your counterpart's final score.

Stage 5 The person getting the higher final score in your pair will receive 10 tokens, the other person will receive 2 tokens.

The study is conducted anonymously and without communication between you and your counterpart. Participants will be identified only by code numbers.

The exercise is *repeated four times.*

- You will remain in the same pair for all four rounds.
- We will always draw 3 balls for you and 12 balls for your counterpart.
- The composition of the hat, that is the average value of the bonus (all balls), will be different in each round. At the beginning of each round, you will be informed of the highest and the lowest possible value of the bonus (X and Y).
- The prize for the winner, 10 tokens, and the prize for the other person, 2 tokens, will remain the same in each round. At the end of each round, we will determine the winner and inform you whether you won 10 or 2 tokens.

Specific instructions for how to mark the words Once we start, you will see a letter matrix on your screen. You can highlight the words you find by marking them with your mouse. Your task is to mark as many correct words as possible. We will practice this in a trial round.

If you have any questions, please press the help button now. Once we have addressed all questions, we will start with the trial round. You will have another chance to ask questions after you have familiarized yourself with the methodology in the trial round.