When does it pay to get informed? (extended abstract)

Ayça Kaya*

When an agent is initially employed by a principal, the technology determining the outcome of this match may be uncertain to both parties. For instance, if the agent is hired as a sales person, the principal and the agent may be uninformed about the exact market conditions that the agent is going to operate in. If the principal has recently adapted a new technology of production, both of the parties may be uncertain about the potential of this technology. Or simply, both parties may be uncertain about the innate competence of the agent or about how good a match the agent is for the specific job that he is employed for. In many cases, the principal possesses the means to find out more about this technology—he may conduct market research or certain other tests to assess the technology, the agent or the match quality. I address the question of whether and, if yes, after which histories the principal chooses to get informed. It is conceivable to think that, if the assessment tools that the principal possesses are cheap enough or, at the extreme, costless, the principal should at least weakly prefer to acquire information after any history. In this paper I present a model that establishes that this intuition is not necessarily correct: the principal may choose to never get informed even if information acquisition is costless.

I consider an infinitely repeated principal-agent problem where both parties are risk neutral. The environment is that of moral hazard, so that the principal needs to give incentives to the agent to induce effort. The principal is not able to commit to long term contracts, however he can commit to a bonus pay at the end of each period in case the period outcome is favorable. There is uncertainty about the distribution of output as a function of the effort exerted by the agent. In particular, I assume that the agent can choose to work or shirk. If the agent shirks, the outcome is failure with probability 1. If the agent works, the outcome is success with probability α . The uncertainty is about α which can take on one of two values: $\alpha_H > \alpha_L$. At the beginning of each period the principal chooses whether to find out the exact value of the parameter α . If he chooses, he is able to acquire this information costlessly. After making this decision, the principal either

*ayca-kaya@uiowa.edu

terminates the agent or makes a take-it-or-leave-it contract offer to the agent. If the relationship is dissolved, the principal matches with a new agent. I consider two versions of the model: in the first the unknown parameter α denotes a characteristic of the principal, such as the specifics of a newly adapted technology that will stay with the principal even if the relationship is dissolved. In the second, α denotes a match-specific characteristic or a characteristic of the agent, so that after dissolving the relationship, the principal can get a new draw of α .

Conventional wisdom suggests that, if information acquisition is costless, the principal can do weakly better when he gets informed because he can always mimic the choices of an uninformed principal and be at least as well off. The glitch in this argument is that it does not take into account the agent's reaction to the knowledge that the principal has become informed.¹ One can argue that the agent should anticipate a "separating equilibrium" once the principal becomes informed. After all, the principal that receives good news about the technology is the one who wants to communicate this information. Therefore, it is reasonable to assume that the agent will expect the principal to signal once he gets informed and will interpret no signaling as bad news. Hence, the decision of the principal to get informed is akin to taking a lottery which pays either the separating equilibrium payoffs of an efficient principal or the separating equilibrium payoff of an inefficient principal. On the other hand, if the principal stays uninformed, he continues to get payoffs identical to those under a pooling equilibrium. Clearly, whether taking the lottery is worthwhile or not depends on the probabilities assigned by the principal to each of the potential continuation paths, which in turn depends on the history of outcomes. Moreover, these probabilities also determine the value of remaining uninformed, as they also determine the beliefs of the agent and in turn how much rent the principal can extract from the agent in a pooling equilibrium.

I characterize the decision of the principal about whether to get informed as a function of the beliefs about the unknown parameter α . In order to perform this analysis, I first characterize the payoff to the principal after getting informed as the ex-ante expected value of the payoffs from a separating perfect Bayesian equilibrium of the repeated contract offering game. I show that, in this context, the repetition of the one-shot optimal signaling contract is optimal. Secondly, I remark that, due to the risk neutrality assumption, the highest payoff for the uninformed principal can be achieved by charging the agent an upfront payment at the beginning of each period and making him the residual claimant. I also note that in this setup the principal has no incentive to end the

¹For all of the arguments that follow, it is not necessary to assume that the agent can observe the decision of the principal to get informed. In an equilibrium setting, since both the agent and the principal can observe the history of outcomes, the agent anticipates the principal's decision to get informed.

relationship if his decision to get informed is observable by all potential future employees.²

If the beliefs attach high enough probability to α_H , it is well known that the principal is better off in the “best” pooling equilibrium than in the “best” separating equilibrium, even conditional on being the good type. Moreover, pooling payoff is higher the higher the belief attached to the good type is. Therefore, it is unambiguously better to stay uninformed when beliefs attach high probability to the good type of the principal. On the other hand, if the beliefs attach most probability to the bad type of the principal, the value of a separating equilibrium conditional on being the good type is higher than that of the pooling equilibrium. However, of course, the lottery in this case is one that puts large probability on the bad continuation value—namely that of the bad type principal in a separating equilibrium. In this case, and also for non-extreme beliefs, the decision of the principal becomes non-trivial.

I show that in the first version of the model, where α represents a characteristic of the principal, the principal never chooses to get informed. The intuition is simple: given the risk neutrality of the two parties, if the principal gets informed and finds out that he is the bad type, he can extract all the surplus of the relationship conditional on being the bad type. If the principal finds out that he is the good type, he needs to leave rents to the agent to be able to signal that he is the good type, and hence cannot appropriate the whole surplus. Therefore, the average payout of the lottery of getting informed is strictly less than the expected value of the surplus of the match. However, the principal can always extract the expected surplus by staying uninformed indefinitely. Hence, in this version, it is not optimal for the principal to get informed after any history.

In the second version of the model—namely the version in which α is a characteristic pertaining to the agent or is a match-specific parameter—if the principal gets informed and finds out that α is low, he can terminate the agent and essentially re-start the game. This renders getting informed more valuable. Note that the principal that finds out that α is high may still need to engage in costly signaling. This is because, if by simply not terminating, the principal can convince the agent that the news is good, and if the payoff from restarting the game is not very high (e.g. if the prior on α attaches large probability to α_L), the principal with bad news may wish to continue the relationship instead of terminating. I conjecture that in this version, the decision of the principal takes the form of a cut-off rule whereby the principal gets informed when the beliefs start attaching low enough probability to α_H , i.e. after histories that involve many failures.

²Again, for this argument to work, it is not necessary for the potential agents to observe that the principal gets informed. It is sufficient that they observe the history of outcomes for this principal, in which case they can anticipate when the principal gets informed.