Experimentation in Teams

Sofia Moroni

April 15, 2014

Abstract

This paper examines learning, communication and dynamic moral hazard in teams. I consider a relationship between a principal and a team of agents who work on a risky project. Agents have private information about their own productivity or ability and choose unobservable effort; their beliefs about the feasibility of the project evolve privately as they exert effort. The principal must provide incentives for agents to exert effort, but must also incentivize the proper amount of information sharing among the agents. We characterize optimal contracts for teams in the presence of adverse selection and dynamic moral hazard.

1 Research Questions

This paper examines learning, communication and dynamic moral hazard in teams. I consider a relationship between a principal and a team of agents who work on a risky project. Agents have private information about their own productivity or ability and choose unobservable effort; their beliefs about the feasibility of the project evolve privately as they exert effort. The principal must provide incentives for agents to exert effort, but must also incentivize the proper amount of information sharing among the agents.

We consider an exponential bandit model in the style of Keller, Rady and Cripps (2005) in which workers experiment in an arm that yields a payoff to the principal at a stochastic time. The principal does not observe the amount of effort that the agents exert in the arm and has to create incentives through outcome contingent contracts. If, additionally, the agents have private information about their capabilities to exert effort and the cost they incur during effort provision the principal offers contracts that depend on type reports and must satisfy incentive compatibility.

From the principal's point of view, the agents' outcomes are correlated because of the uncertain feasibility of the project. In the absence of communication concerns and adverse selection, I show that the principal would like to sell to "sell the risky arm" to each player, paying him if and only if he achieves the first breakthrough. Such schemes create a high level of competition between agents, and hence may hinder information transmission among agents about the state of the project. There is a tradeoff between incentives and information sharing between the players when the optimal scheme rewards outcomes in which agents do better than their peers. This paper characterizes optimal schemes when the principal faces this tradeoff.

2 Model

There are n agents attempting to perform a task and a principal who owns the production of the agents. The task may be "good" or "bad". Only a good task can ever be completed. The probability that the task is good is $\bar{p}$ which is commonly known by all participants. The agents exert a privately observed effort over time $t \in \mathbb{R}^+$. Effort is costly. Agent i exerts effort $u_{i,t} \in [0, \bar{u}_i]$ at time t at cost $\alpha_i u_{i,t}$.

If the task is good and agent i exerts effort $u_{i,t}$ at time t the task is completed with a breakthrough by player i with instantaneous probability $\lambda_i u_{i,t}$ where λ_i denotes the arrival rate of player i.

Once a breakthrough occurs the game ends. A breakthrough has a value of v for the principal. As long as no breakthrough has occurred the agents do not incur any benefits from the project. All players discount the future at rate r.

The principal offers contracts to each agent i that specify a wage schedule $\tilde{w}_{i,t}$ that is contingent on the history (i.e whether the task has been completed, by whom and at what time) in order to induce the agents to exert effort.

Assume the realized wages paid to each player i consists of flow payoff $\tilde{w}^f_{i,t}$ and lump sums $\tilde{w}^l_{i,t_k}$ at times $\{t_k\}_{k\in I}$ I for $i \in \{1,\dots,n\}$. If the task is completed at time t the payoff to the principal is:

$$r(ve^{-rt} - \int_0^\infty \tilde{w}^f_{i,s} e^{-rs} ds - \sum_{k\in I} \tilde{w}^l_{i,t_k})$$

If the task is never completed, the payoff to the principal is

$$-r(\int_0^\infty \tilde{w}^f_{i,s} e^{-rs} ds + \sum_{k\in I} \tilde{w}^l_{i,t_k})$$

The agents' payoff from exerting effort $u_{i,t}$ is

$$r\int_0^\infty e^{-rt}(\tilde{w}^f_{i,s} - \alpha_i u_{i,s})ds + \sum_{k\in I} \tilde{w}^l_{i,t_k}$$

If an agent rejects the contract offered by the principal he gets a payoff of zero.

Information We will assume throughout that the effort is unknown. In the first part of the paper we consider the case in which only the effort is unobserved, that is the pure moral hazard case. In the second part, we analyze the case in which either the cost of effort, α_i, maximum effort, $\bar{u}_i$, or the rate of arrival of each player, λ^i, is unobserved, thus, combining moral hazard and adverse selection. All players observe the task completion as soon as it occurs.

Related literature

The papers that are closest to ours is Bonatti and Hörner (2009,2011) and Halac, Kartik, Liu (2013). The former analyzes a bandit model in which multiple players exert joint effort on an arm with moral hazard. In the working paper version they consider a setting with a principal and two agents. In contrast to my paper, in their setting individual outcomes cannot be identified, so that the principal can only reward team outcomes. The latter paper analyzes a dynamic principal-agent relationship in which the agent is privately informed about his type and has unobservable actions. The authors characterize optimal contracts under adverse selection and dynamic moral hazard. Because there is only one agent there is no concern with team incentives or information sharing between agents.

Outside of experimentation models, there is a large literature on team incentives with and without a principal. Holmstrom (1982) is the foundational paper in this literature; many of the intuitions here, such as the optimality of relative performance measures when agents' outcomes are correlated, can be found in that paper.

Results I first extend a result of Halac, Kartik and Liu to our setting, showing that it is sufficient to consider contracts with a relatively simple form. In particular for any contract there is an equivalent "bonus contract" consisting of (possibly negative) transfers at time zero and time-dependent transfers at the time that an agent makes a breakthrough.

Under pure moral hazard, the principal can implement the first best by "selling each agent his own arm". That, is each agent makes a transfer to the principal at the beginning of the game equal to the expected value of his own arm and receives v if he completes the task. Because of risk neutrality there are many contracts that give the principal the first best payoffs. For example, the first best can be attained by a contract that makes a transfer to the agents at the beginning of the game and then charges them flow penalties as long as they do not complete the task.[1]

I next characterize the optimal contract under pure moral hazard and limited liability; the principal offers contracts with a time-dependent bonus on the first breakthrough to the agent who achieved the breakthrough.

I then turn to the case of moral hazard and adverse selection, both with and without limited liability.

References

[1] Alessandro Bonatti and Johannes Horner. Collaborating. 2009. Cowles Foundation Discussion Paper No. 1695.

[2] Alessandro Bonatti and Johannes Horner. Collaborating. *The American Economic Review*, 101(2):632–663, 2011.

[3] Marina C Halac, Qingmin Liu, and Navin Kartik. Optimal contracts for experimentation. 2013. Department of Economics, Columbia University.

[4] Bengt Holmstrom. Moral hazard in teams. *The Bell Journal of Economics*, pages 324–340, 1982.

[5] Godfrey Keller, Sven Rady, and Martin Cripps. Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68, 2005.

[1] Halac, Kartik and Liu (2013) find a similar result in a model with one agent in discrete time.