Randomization in Dynamic Principal-Agent Problems

Dr Holger M. Müller

Dynamic principal-agent models are concerned with the intertemporal structure of the contractual relationship between two parties, a principal and an agent. For instance, the principal could be a boss and the agent could be a worker. Here, we assume that the boss wants the worker to expend a consistently high level of effort. Another example is the relationship between two firms, a buyer and a supplier. Suppose the buyer orders 10,000 widgets and wants the supplier to exert the same level of care in the production of every single widget. Thirdly, the principal could be a landlord and the agent could be a farmer. In this case, we assume that the landlord wants the farmer to use the same amount of water and fertilizer for every square inch of the cornfield.

What makes the principal-agent problem intricate is the existence of asymmetric information. In particular, we assume that the principal can only observe the agent's output, but not how much effort he actually expended. In the landlord-farmer example, this means for instance that the landlord can observe the quality of the corn, but not how much water and fertilizer the farmer actually used. One implication of this is that the farmer may be lazy. If the crop is bad, the farmer simply asserts that he followed the principal's instructions and puts the blame on bad weather. Due to asymmetric information, the landlord cannot verify whether the farmer tells the truth. Hence the two parties cannot write a contract that stipulates a certain amount of water and fertilizer. In economics, this problem is known as *moral hazard.*

As a consequence, the principal must provide the agent with (monetary) incentives to exert a consistently high level of effort. One example of a bad incentive is fixed wages. If the agent gets a fixed wage, he will be lazy since more effort is not rewarded by a higher income. From this it follows that the agent's pay must depend on output. But how? A famous contribution in the literature suggests that linear incentive schemes are optimal. The reason for this is that linear schemes provide the same marginal reward for every additio-

nal unit of output. Consequently, the agent will expend a constant level of effort. In addition, the effort intensity can be controlled through the slope of the incentive scheme. A popular example of linear incentive schemes are piece rates. If a factory worker gets $K = 5$ dollars per produced widget, he has incentives to spend the same amount of time and effort on each widget. Furthermore, he will work harder than if he got only $K = 3$ dollars per widget.

In my dissertation I show that in dynamic principal-agent problems, linear incentive schemes may not be optimal for the principal. In particular, a random spot check in conjunction with a penalty scheme has the same incentive effects but performs better. Intuitively, random spot checks, like linear incentive schemes, provide constant incentives. Moreover, the effort intensity can be controlled through the size of the penalty. In the landlord farmer example, this means that the landlord randomly selects a small part of the cornfield and measures the quality of the corn. If the quality exceeds a certain cutoff, the farmer is rewarded, otherwise he is penalized. Since the farmer does not know in advance which part of the field is going to be selected, he will use the same amount of water and fertilizer everywhere. And clearly, the farmer will use more water and fertilizer per square inch if the penalty is high than if it was low. Random spot checks are prevalent in many real-world situations. For instance, they occur in the form of random quality inspections by buyers, random inspections by bosses or supervisors, random checks in the military, and random inspections of arms production sites in Iraq, Libya, etc by the UN.