

INTRODUCTORY ESSAY TO “GROVES SCHEME ON RESTRICTED DOMAINS”

Bengt Holmstrom

Like many of Bob’s students, I began my career in operations research. Before going to Stanford in 1974 I spent two years as an operations research analyst at a Finnish conglomerate (Ahlstrom Oy). At Stanford, I first enrolled in the master’s program in operations research, thinking of returning to my old job after a year. Fortunately, those plans changed when I learned of and got accepted into the Ph.D. program at the GSB.

The first economics class I ever took was Bob’s legendary Multi-Person Decision Theory class. I didn’t understand very much of the material (it would have been prudent to go for something a little easier at first), but it became quickly clear that I was listening to an exceptional mind. Bob wasn’t polished in a strictly pedagogical sense, but the things he chose to emphasize were fascinating. He spent a lot of time on problem formulation: how to transform a question into an economically relevant, mathematically tractable model. “Formulation is 90% of the analysis” was one of his favorite phrases, though it took me a long time to fully appreciate the significance of this lesson.

Early in the class I asked Bob what I should read if I wanted to understand incentives better. My interest in incentives had been raised by the problems I had encountered when trying to implement a large-scale corporate planning model at Ahlstrom. It was evident that the information that divisions supplied was distorted as managers tried to game the model. Bob suggested that I look at Groves’ recent paper “Incentives in Teams,” which, he said, dealt with the exact problem I was talking about. I rushed to read it - and was gravely disappointed. I remember saying to myself that if this is what economists have to say about incentives then I better work on something else. To someone with an operations research background, the model looked hopelessly unrealistic.

The paper in this volume shows that despite this first, unhappy encounter with the economics of incentives, I eventually came around to appreciate both the economic approach and Groves’ scheme (but not before I had wasted a semester working on integer programming). The paper is very much a product of Bob’s wonderful, informal research group that I participated in as a thesis writer. Among the regular participants at that time were Takao Kobayashi, who worked with Bob on mechanism design, Froystein Gjesdal, who worked with Joel Demski on moral hazard, William Thomson, who worked with Mordecai Kurtz

on social choice; and Barry Weingast and Linda Cohen, who worked with Roger Noll on political science (all three visiting from Cal Tech).

The group met weekly. The meetings were almost entirely unstructured. We would often come to a session not knowing who would speak and on what topic. Whoever had something to say would go to the black board and sketch out a model, a problem or just bring up an interesting issue worth thinking about. We never ran out of topics to discuss, perhaps because the pressure to deliver new ideas was so high. It was part fun, part scary and above all tremendously educating and inspiring.

My paper on Groves' scheme got started with a simple question: do there exist incentive schemes, based on joint output alone, that implement efficient outcomes in a model of joint production ("moral hazard in teams")? Assuming differentiability, the answer is readily seen to be negative. When I presented this little discovery in the discussion group, I was surprised to see Bob get quite enthusiastic about it. He was convinced the idea could be leveraged further. Inspired by this, I went on to ask the same implementation question in the Groves' model using the same differentiable approach. I wanted a method for deriving Groves' scheme (something Groves' paper didn't do). The main insight was to note that the condition for an efficient allocation could be recast as the condition that the social planner wants truthful revelation (equation (6) in the paper). With this minor reformulation, Groves' scheme popped right out as the unique solution to a trivial differential equation. Validating Bob's thesis, the reformulation was 90% of the paper. The remaining 10% is a technical exercise showing that uniqueness obtains under quite weak differentiability conditions.

As an advisor, Bob was extremely generous with his time. He told me at the outset that he liked to put students on a regular schedule: discipline was important for doing good work. I met with him almost weekly. It was a great incentive scheme. I didn't dare to go to his office empty-handed and usually managed to produce something that at least could be discussed (often late the night before). Bob would rarely take a strong stand on issues, but I learned to read his delicate, invaluable signs of approval or disapproval. His comments were relatively few, but deep and insightful. It usually took me a long while to figure out their true meaning — in some cases I've figured it out only years later. His ability to interpret the literature and see connections between my results and earlier work were invaluable. And just as importantly, his whole vision of where the field was headed and the role that incentive theory and information would come to play in economics were a strong encouragement for a young, aspiring theorist.

It is only in retrospect that I have come to appreciate how unique my education at Stanford really was. It's hardly an exaggeration to say that modern game theory would not be what it is today without the foresight and enormous inspiration that Bob provided not just to his immediate students, but a whole generation of young scholars around the world.