ARTICLE IN PRESS
YGAME:2500

Games and Economic Behavior ••• (••••) •••–•••

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Generalized Groves–Ledyard mechanisms

Paul J. Healy *, Ritesh Jain

Department of Economics, The Ohio State University, 1945 North High Street, Columbus, OH 43210, USA

ARTICLE INFO

Article history:
Received 2 March 2015
Available online xxxx

JEL classification:
D02
D47
D61
D62
D71

Keywords:
Nash implementation
Public goods
Demand revelation

ABSTRACT

Groves and Ledyard (1977) construct a mechanism for public goods procurement that can be viewed as a direct-revelation Groves mechanism in which agents announce a parameter of a quadratic approximation of their true preferences. The mechanism's Nash equilibrium outcomes are efficient. The budget is balanced because Groves mechanisms are balanced for the announced quadratic preferences. Tian (1996) subsequently discovered a richer set of budget-balancing preferences. We replicate the Groves–Ledyard construction using this expanded set of preferences, and uncover a new set of complex mechanisms that generalize the original Groves–Ledyard mechanism. The original mechanism, however, remains the most appealing in terms of both simplicity and stability.

1. Introduction

Groves and Ledyard (1977) provided one of the first decentralized economic mechanisms to solve the free rider problem for general economies. Specifically, they devised a government (or, mechanism) such that self-interested equilibrium behavior by all parties always leads to a Pareto optimal allocation. This work has been cited widely, with many researchers building off its original insights.

What has not been appreciated, however, is the manner in which Groves and Ledyard (1977) actually construct their mechanism. The final mechanism looks simple: players announce a single number, and taxes are based on a proportional share of the cost and a quadratic penalty. In fact, the mechanism has a more complex foundation: Groves and Ledyard (1977) present it as being derived from a Groves mechanism (Groves, 1973) in which agents announce entire quasilinear utility functions and are taxed (or rewarded) based on the Marshallian surplus calculated from the announced preferences of all other agents in the economy. But agents may not actually have quasilinear preferences, so these announcements represent approximations of their true preferences. The space of admissible announcements is parameterized with a single parameter, which greatly simplifies agents' messages. In equilibrium, agents with non-quasilinear preferences announce quasilinear preferences that best approximate their true willingness to pay at the Pareto optimal allocation, and that allocation is selected by the mechanism.[1]

What is crucial to the Groves-cum-Groves–Ledyard construction is that the set of approximating preferences that the agents are allowed to announce always generate a budget-balanced allocation in the Groves mechanism. Budget balance

* Corresponding author.
E-mail addresses: healy.52@osu.edu (P.J. Healy), jain.224@osu.edu (R. Jain).

[1] Groves (1973) mechanisms are usually only discussed in settings where all agents actually have quasilinear preferences, and therefore have a dominant strategy of truthful revelation. Groves and Ledyard (1977) are novel in that they consider the Groves mechanism for more general preferences.

http://dx.doi.org/10.1016/j.geb.2016.01.005
0899-8256/

Please cite this article in press as: Healy, P.J., Jain, R. Generalized Groves–Ledyard mechanisms. Games Econ. Behav. (2016), http://dx.doi.org/10.1016/j.geb.2016.01.005

does not always obtain with general quasilinear preferences; Groves and Loeb (1975) showed that balance is achieved if we further restrict agents to quadratic valuation functions. Following this insight, Groves and Ledyard (1977) only allow agents to announce quadratic (approximate) preferences, guaranteeing that the final allocation will balance the budget. Agents need only to announce the intercept of their (approximate) valuation function, and are charged according to the balanced Groves mechanism. Algebraic manipulation reveals that these outcome and tax functions reduce to the familiar Groves–Ledyard mechanism with its quadratic form.

For two decades it was believed that quadratic preferences are the only ones for which a Groves mechanism could be budget balanced.[2] This belief was overturned when Tian (1996) discovered a much richer class of preferences for which Groves mechanisms can be balanced. And, like the quadratic preferences, these are parameterized by a single parameter.

In this paper we replicate the Groves–Ledyard procedure of turning a Groves mechanism into a one-dimensional, budget-balanced, Pareto efficient mechanism, but we do so on Tian (1996)'s larger domain of preferences. In doing so, we uncover a broad family of complex mechanisms that, in theory, could also be used to solve the free-rider problem. We also gain a deeper understanding of the incentive properties of Groves mechanisms applied to general preferences. We see that getting the Samuelson condition (equating the sum of marginal rates of substitution to marginal costs) is a relatively trivial matter, so that the real difficulties in the design problem lie in achieving budget balance and equilibrium existence. We lean on Tian (1996) to accomplish the former; for the latter, we devise a simple trick of allowing agents to announce transfers among their partners, providing just enough richness to the range of the mechanisms (without distorting incentives) to guarantee equilibrium existence in these mechanisms.

Finally, we study the stability properties of these Generalized Groves–Ledyard mechanisms, since one beneficial property of the original Groves–Ledyard mechanism is that it becomes dynamically stable under appropriate parameter values (Chen and Plott, 1996; Page and Tassier, 2004; Healy, 2006; Healy and Mathevet, 2012). It is clear that stability will become difficult to obtain in the generalized versions of the mechanism. We show that in one generalized version the best response functions' slopes depend crucially on the messages sent. With very extreme messages the best response slopes become steep enough that stability is violated. We argue that this will be a pervasive problem in any of the generalized versions of the mechanism. This suggests that the original Groves–Ledyard mechanism is not only the simplest among the generalized versions, but it is likely the only one that can be made globally stable.

2. Setup

2.1. The economic environment

We present here a simplified version of the Arrow–Debreu setting studied by Groves and Ledyard (1977). Specifically, there is one private good and one public good, I consumers, F firms, and a mechanism (that can be thought of as a government). The mechanism receives messages from the consumers and uses this information to procure the public good from the firms. This purchase is financed by (message-dependent) the taxes collected from the consumers.

The price of the private good is normalized to one, and the price of the public good is given by $p \in \mathbb{R}$. Quantities are given by $x \in \mathbb{R}$ for the private good and $y \in \mathbb{R}$ for the public good. Each consumer i has a consumption set $\mathcal{X}_i \subseteq \mathbb{R}^2$, a preference relation $\succeq_i$ on $\mathcal{X}_i$ that is representable by a differentiable utility function u_i, and an initial endowment of private goods $\omega_i \in \mathbb{R}$. Firms are characterized by a production set $\mathcal{Z}_f \in \mathbb{R}^2$ and a vector of profit shares $\theta_f = (\theta_f^1, \ldots, \theta_f^I)$. Production vectors are given by $z_f = (z_f^x, z_f^y)$, with negative components representing inputs. Firm profits are distributed to consumers according to θ_f. An *economy* is therefore represented by $e = ((\mathcal{X}_i, \succeq_i, \omega_i)_{i=1}^I, (\mathcal{Z}_f, \theta_f)_{f=1}^F)$.

We think of there being a set of admissible economies $\mathcal{E}$, where each $e \in \mathcal{E}$ differs only in the preference profiles of the consumers. We assume that, for every $e \in \mathcal{E}$, each u_i is continuously differentiable, quasi-concave, and strictly increasing in the private good. In some cases we may discuss further restrictions on $\mathcal{E}$, such as requiring that all preferences are quasilinear ($u_i(x_i, y) = v_i(y) + x_i$ for some concave function v_i), or even quadratic-quasilinear ($v_i(y)$ is concave and quadratic).

An *allocation* is a vector $((x_i)_{i=1}^I, y, (z_f)_{f=1}^F)$ with $x_i \in \mathbb{R}$ for each consumer i, $y \in \mathbb{R}$, and $z_f \in \mathbb{R}^2$ for each firm f. An allocation is *feasible* if (i) $(x_i, y) \in \mathcal{X}_i$ for each i, (ii) $z_f \in \mathcal{Z}_f$ for each j, and (iii) $(\sum_i (x_i - \omega_i), y) \leq \sum_f z_f$. An allocation $((x_i)_{i=1}^I, y, (z_f)_{f=1}^F)$ is *Pareto optimal* if it is feasible and there is no other feasible allocation $((x_i')_{i=1}^I, y', (z_f')_{f=1}^F)$ such that $(x_i', y') \succeq_i (x_i, y)$ for all i and $(x_i', y') \succ_i (x_i, y)$ for some i.

2.2. Mechanisms and competitive equilibrium

A *mechanism* (or *government*) is simply tuple which specifies a message space $M = \times_i M_i$ (one for each consumer), an allocation rule $y(m, p)$ selecting a public goods level for each message profile m and (public good) price p, and a list

[2] This was conjectured by Laffont and Maskin (1980), who provide general conditions for budget balance.