Games and Economic Behavior 94 (2015) 169–181

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Optimal private good allocation: The case for a balanced budget

Moritz Drexl, Andreas Kleiner *

Bonn Graduate School of Economics, University of Bonn, Kaiserstr. 1, 53113 Bonn, Germany

ARTICLE INFO

Article history:
Received 11 December 2012
Available online 28 October 2015

JEL classification:
D82
D45

Keywords:
Mechanism design
Bilateral trade
Myerson–Satterthwaite theorem
Budget balance

ABSTRACT

In an independent private value auction environment, we are interested in strategy-proof mechanisms that maximize the agents' residual surplus, that is, the utility derived from the physical allocation minus transfers accruing to an external entity. We find that, under the assumption of an increasing hazard rate of type distributions, an optimal deterministic mechanism never extracts any net payments from the agents, that is, it will be budget-balanced. Specifically, optimal mechanisms have a simple "posted price" or "option" form. In the bilateral trade environment, we obtain optimality of posted price mechanisms without any assumption on type distributions.

1. Introduction

Most parts of the mechanism design literature studying welfare maximization problems focus on mechanisms implementing the efficient allocation. However, in general it is not possible to implement the efficient allocation in dominant strategies using budget-balanced mechanisms (Green and Laffont, 1979). Given this result, we study how to choose among different mechanisms that cannot attain both, allocative efficiency and budget-balancedness. Since we are concerned with welfare maximization, the social planner's objective function should consist of the agents' aggregate utility and therefore include aggregate transfers. In other words, one seeks to find mechanisms that maximize what we call residual surplus. This is the surplus, or utility, the agents derive from the chosen physical allocation, reduced by the amount of transfers that are lost to an external agency (this is often called "money burning").

A common approach is to implement the efficient allocation via Groves mechanisms and to redistribute as much money to the agents as possible without distorting incentives (Cavallo, 2006; Guo and Conitzer, 2009, 2010; Moulin, 2009). This approach aims at characterizing the optimal mechanism for allocating private goods that implements the *efficient allocation* in *dominant strategies*, is *individually rational* and *never creates a budget deficit* (ex-post).[1] However, if mechanisms that allocate inefficiently yield higher residual surplus (Guo and Conitzer, 2014) it is not clear why one should use a mechanism that allocates efficiently.

Consequently, we drop the requirement that mechanisms allocate efficiently. Instead, we take an optimal mechanism design approach and consider mechanisms that are comparable to the ones considered before in that they are *strategy-proof*, *deterministic*, *never run a deficit* and satisfy *ex-post participation constraints*. We analyze which mechanism maximizes residual

* Corresponding author.
E-mail addresses: drexl@uni-bonn.de (M. Drexl), akleiner@uni-bonn.de (A. Kleiner).

[1] In our setting, the best Groves mechanism is implemented by a second-price auction with two bidders.

http://dx.doi.org/10.1016/j.geb.2015.10.009
0899-8256/

surplus when an indivisible good is auctioned among two agents with independent private values that are distributed according to prior type distributions. We show that under an increasing hazard rate assumption on type distributions, the optimal mechanism will never waste any payments, thereby deviating distinctly from the efficient allocation (Theorem 1). In fact, our proof method reveals that all mechanisms that allocate efficiently are worse than the simple mechanism where the object is always given to the same agent (the one with the higher expected valuation; Corollary 1), showing that our general mechanism design approach has clear advantages over the previous approach to search for the optimal Groves mechanism. We show that the optimal mechanism is either a "posted price" or an "option" mechanism: The object is assigned to one of the agents unless both agents agree to trade at a prespecified price (posted price mechanism) or unless the second agent uses his option to buy the object at a fixed price from the first agent (option mechanism). Therefore, the optimal mechanisms do not invoke money burning and are of a particularly simple form. Moreover, numerical simulations indicate that these simple mechanisms obtain a large share of first-best welfare (92 per cent on average in our simulations). In the bilateral trade setting, we establish optimality of posted price mechanisms without any restrictions on type distributions (Theorem 2). This provides an argument for the focus on budget-balanced mechanisms (see Myerson and Satterthwaite, 1983; Hagerty and Rogerson, 1987).

The requirement that a mechanism does not produce a budget deficit ex-post is considerably stronger than the requirement that this holds in expectation. However, in many situations it is reasonable that a budget breaker is infeasible and therefore ex-post constraints need to be obeyed. This includes situations where there is no insurance or where agents have restricted access to capital markets. Also, hidden information issues towards a third party cannot always be resolved, and autarkic mechanisms that can be implemented without explicit intervention by a third party might be preferable (e.g., when mechanisms are used to model bargaining situations; Myerson and Satterthwaite, 1983; Hagerty and Rogerson, 1987). If all these considerations do not apply and mechanisms that create no deficit in expectation can be implemented (for example, because the designer has unlimited liability), then one can achieve the first-best solution (see Section 5). Similarly, we show that one can achieve the first-best if mechanisms are only required to be Bayesian incentive compatible (Proposition 1). In contrast to these two constraints, which are the main driving forces behind our results, we argue that the participation constraint and the restriction to deterministic mechanisms are not essential to the spirit of our results (Section 5).

Our work is part of a small literature that searches for mechanisms maximizing residual surplus when the first-best is not achievable. Miller (2012) studies a model of firms colluding in a Bertrand oligopoly. A mechanism used by a cartel to allocate market shares should maximize residual surplus. Miller shows that under general conditions it is never optimal to allocate market shares efficiently and gives numerical evidence that for some type distributions it is optimal to give up efficiency in order to obtain a balanced budget. However, other examples indicate that this observation does not hold for all distributions. Athey and Miller (2007) study residual surplus maximization in a repeated bilateral trade setting and obtain numerical results suggesting that for many type distributions the optimal mechanism is a posted price mechanism. Closely related to our paper is independent work by Shao and Zhou (2012), who obtain the characterization of our Theorem 1 when restricting to symmetric distributions of types and allowing mechanisms to violate individual rationality.

The result that the efficient allocation is never optimal contrasts with the literature cited above that restricts attention to efficient rules (Cavallo, 2006; Guo and Conitzer, 2009, 2010; Moulin, 2009). Recently, Athanasiou (2013) and Sprumont (2013) relax this requirement. Similarly to our work, they focus on mechanisms that are deterministic, strategy-proof, ex-post individually rational and create no deficit ex-post. However, they require mechanisms in addition to be anonymous, which immediately implies that whenever the object is allocated, it is allocated to the agent that values it the most (weak assignment efficiency). This restricts the set of mechanisms severely and excludes the mechanisms that turn out to be optimal in our analysis.

The restriction to efficient allocation rules has also been relaxed in a series of papers that study specific mechanisms in a multi-unit setting. Faltings (2005) and Moulin (2009) propose simple mechanisms where one agent is designated as a residual claimant and is allocated one unit (or no unit, respectively) independent of his type. The remaining units are auctioned among the other agents and the residual claimant receives all payments accruing in the auction. Faltings uses numerical examples to argue that his mechanism often outperforms the VCG mechanism. Moreover, Moulin shows that his mechanism provides a higher worst-case welfare guarantee than any VCG mechanism given that there are sufficiently many objects and agents. In our setting with two agents and one object, these mechanisms always allocate the object to a fixed agent and therefore correspond to a degenerate option mechanism. Our Corollary 1 supports Faltings' numerical results in the two agent setting by showing that under regular prior distributions his mechanism indeed outperforms the VCG mechanism. Building on the ideas of Faltings and Moulin, Guo and Conitzer (2014) provide worst-case welfare guarantees for two specific classes of mechanisms that allocate inefficiently: *Burning allocation mechanisms* burn a (random) number of units and assign the remaining units efficiently. *Partitioning mechanisms* partition units and agents randomly into two groups, allocate the objects in each partition efficiently to the agents in the corresponding partition and distribute the payments to agents in the other partition. Similarly, de Clippel et al. (2014) propose a deterministic mechanism where the burning of items is contingent on the reports of the agents; they provide worst-case welfare guarantees that converge to 0.88 asymptotically as the number of agents grows. Our work differs from these papers by evaluating mechanisms according to a Bayesian prior, restricting ourselves to the two agent setting and using a general optimal mechanism design perspective.

Another related strand of the literature studies the expected residual surplus of Bayesian incentive compatible mechanisms when it is not possible to redistribute any payments among the agents (Hartline and Roughgarden, 2008; Chakravarty and Kaplan, 2013; Condorelli, 2012). This implies that methods similar to those in Myerson (1981) can be applied. It is