Games and Economic Behavior 82 (2013) 424–443

Contents lists available at SciVerse ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Mechanism design with a restricted action space ☆

Liad Blumrosen [a,b,∗], Michal Feldman [c,b,1]

[a] Department of Economics, The Hebrew University of Jerusalem, Mount Scopus, Jerusalem, 91905, Israel
[b] Center for the Study of Rationality, The Hebrew University of Jerusalem, Givat Ram, Jerusalem, 91904, Israel
[c] School of Business Administration, The Hebrew University of Jerusalem, Mount Scopus, Jerusalem, 91905, Israel

ARTICLE INFO

Article history:
Received 24 February 2009
Available online 18 April 2013

JEL classification:
D83
D82

Keywords:
Mechanism design
Message space
Implementation
Communication complexity
Expressiveness

ABSTRACT

While traditional mechanism design typically assumes isomorphism between the type space of the players and their action space, behavioral, technical or regulatory factors can severely restrict the set of actions that are actually available to players. We study single-parameter mechanism-design problems in environments with restricted action spaces. In our first main result, we provide sufficient conditions under which the information-theoretically optimal solution can be implemented in equilibrium. Our second main result shows that for a wide family of social-choice rules the optimal mechanisms with k actions incur an expected loss of $O(\frac{1}{k^2})$ compared to the optimal mechanisms with unrestricted action space. We also fully characterize the optimal mechanisms in some simple environments and, finally, we apply our general results to signaling games, public-good models and project planning.

1. Introduction

In standard mechanism-design settings, social planners wish to implement some social-choice rule that chooses an alternative based on the private information of the players. Since social planners cannot observe the private information of the players (their *types*), they design mechanisms that make decisions by observing the *actions* of the players. Each player determines his action in the mechanism according to his type in order to maximize his own utility. The challenge of the social planner is to elicit information that will allow him to implement system-wide goals although such goals may conflict with the objectives of the individual players.

Much of the literature on mechanism design restricts attention to *direct revelation* mechanisms, in which the action space of the players is identical to their type space. This focus is owing to the *revelation principle* (Myerson, 1981; Green and Laffont, 1977; Dasgupta et al., 1979), which asserts that every mechanism can be transformed into an equivalent incentive-compatible direct-revelation mechanism that implements the same social-choice function.

Nonetheless, in most practical settings, direct-revelation mechanisms are not viable since the number of actions available to the players is significantly smaller than their preference space. The most straightforward example is posted-price

☆ The authors are grateful to Itzhak Gilboa, Zvika Neeman, Noam Nisan, Ady Pauzner, Yuval Salant and Ilya Segal for valuable discussions, and to anonymous referees for their thoughtful comments and suggestions. We also thank the participants of the SITE workshop on Market Design at Stanford, and of other seminars at Berkeley, Harvard, Cornell and the Hebrew University, for helpful comments. An extended abstract of this paper with preliminary results appeared in the Proceedings of the ACM Conference on Electronic Commerce, Ann Arbor, MI, 2006. The first author was supported by the Israel Science Foundation grant number 230/10, by the Google Inter-university center for Electronic Markets and Auctions, and by the Maurice Falk Institute for Economic Research in Israel.

∗ Corresponding author at: Department of Economics, The Hebrew University of Jerusalem, Mount Scopus, Jerusalem, 91905, Israel.
E-mail addresses: blumrosen@huji.ac.il (L. Blumrosen), michal.feldman@huji.ac.il (M. Feldman).

[1] Current address: School of Business Administration, The Hebrew University of Jerusalem, Mount Scopus, Jerusalem, 91905, Israel.

0899-8256/$ – see front matter
http://dx.doi.org/10.1016/j.geb.2013.03.005

mechanisms for a single object; in such mechanisms users have only two available actions (to buy the object or not), and the seller clearly lacks the information that would enable her to implement the efficient allocation rule. Another example is the *signaling* model for the labor market by Spence (1973), where employees send signals about their skills to potential employers by the education level they acquire. Although there is a continuum of skill levels, it is unreasonable to expect more than a few education levels in practice (e.g., PhD, M.A., and B.A.). The screening model by Rothschild and Stiglitz (1976) is another example where one might expect a small number of actions; consider an insurance firm that wishes to sell different types of policies to different drivers based on their privately known caution levels. In this model, drivers may have a continuum of possible caution levels, but insurance companies offer only a small number of policies (e.g., a small number of deductible amounts in case of a claim) since it is probably infeasible to market and sell more then a few types of policies. More complicated rules for generating policies may be feasible, but they are rarely used in practice.

Mechanisms with a small, manageable set of choices are widespread in practice, and the main reason for this phenomenon is probably their simplicity. This claim is also supported experimentally, e.g., by Iyengar and Lepper (2000), who showed that a *choice overload* can hamper the willingness of the players to participate in the game, and can degrade their performance in a given transaction. Iyengar et al. compared decision making under a small set of choices and under larger choice sets (not unusually large) and showed that such phenomena are significant even when the number of possible actions is increased from 6 to around 24 or 30. In fact, in many real-life mechanisms the players are required to map their complex preferences into discrete, often dichotomic, decisions. For instance, many mechanisms avoid negotiations and simply post prices for packages or services, and the players are left to decide whether they buy or not under the posted prices. In other settings, players decide whether they participate in or abstain from some transaction, vote for or against some issue in a referendum, and many other similar examples.

Additionally, there are clear evidences for the rare practical use of direct-revelation mechanisms, most prominently VCG mechanisms. One major reason for this fact relates to the *price discovery* process; players usually do not know their exact types and the discovery process may be prohibitively costly (hiring consultants, etc.) or even intractable to compute (see, e.g., Larson and Sandholm, 2001). A well-designed mechanism with limited actions will guide the attention of the players to the information that is most relevant for the decision making. Another critical flaw of direct-revelation mechanisms is that players are typically unwilling to reveal their exact types, even if it is beneficial for them in the short run, worrying that this might harm them in future transactions. A small action space allows the players to preserve some degree of privacy. Papers by Rothkopf et al. (1990) and Ausubel and Milgrom (2006) provide more details on why VCG mechanisms are indeed rare.

Mechanisms with a small action space were studied in several earlier papers in the context of specific models. Wilson (1989) measured the effect of discrete "priority classes" of buyers on the efficiency of electricity markets and found that a few priority classes can realize most of the efficiency gains. In a related work, McAfee (2002) showed that in matching and rationing problems at least half of the social value created by optimal complex schemes can be obtained using very coarse action schemes. Dow (1991) considered a simplified decision problem of a single agent searching for a low price with a limited memory; the memory restrictions force the player to divide the set of possible histories into a limited number of categories. It turns out that the optimal partition of the history is obtained, as in our paper, by dividing the range of prices into disjoint intervals. Compared to the above work, our paper incorporates incentive issues in general multi-player domains and also characterizes the exact effect of the expressiveness level allowed in the system. A similar result was obtained in a different setting, studied in Bergemann and Pesendorfer (2007). There, a revenue-maximizing seller faces a set of bidders, who do not know their private types, and he needs to determine the accuracy level by which they learn their types. On the one hand, more information increases efficiency and thus the seller's revenue, but on the other hand, it increases the information rent of the bidders, thus decreases the seller's revenue. Once again, partitioning the information range into disjoint intervals is shown to maximize the seller's revenue. In a recent paper, Milgrom (2010) highlights the need to study restrictions on the message spaces in mechanisms (what results in "simplified" mechanisms), and applies his results in internet advertising auctions and spectrum auctions. Milgrom (2010) focuses on equilibrium selection problems that we do not touch in this paper. The work of Blumrosen et al. (2007) and of Kos (2008, 2012) is the closest in spirit to our work. Blumrosen et al. studied single-item auctions with severely-restricted action space, and showed that nearly-optimal social welfare can be achieved even with very strict limitations on the action space. The two papers by Kos studied similar questions and some extensions and characterized the informationally-optimal mechanisms in multi-player environments. An earlier work that follows a similar spirit is the paper by Harstad and Rothkopf (1994) who analyzed discrete bid levels in English auctions.

We next present our framework and results.

1.1. Our framework

We consider a Bayesian model with players who have one-dimensional private types, independently distributed on real intervals, and a social planner who wishes to implement a *social-choice function* c that maps every profile of types to an alternative. We stress that although we explore the properties of Bayesian–Nash implementation in this paper, all the mechanisms that we construct have the even stronger dominant-strategy equilibrium. Due to the limited expressiveness that is implied by the restricted action space, for some realizations of the players' types the decision of the social planner will unavoidably be incompatible with the social-choice function c. In order to quantify how well mechanisms with bounded action space can approximate the original social-choice function, we assume that the social-choice function is derived from a