Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 172 (2017) 313–347

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Approximate revenue maximization with multiple items ☆

Sergiu Hart [a,*], Noam Nisan [b,c]

[a] The Hebrew University of Jerusalem, Federmann Center for the Study of Rationality, Department of Economics, and Institute of Mathematics, Israel

[b] The Hebrew University of Jerusalem, Federmann Center for the Study of Rationality, and School of Computer Science and Engineering, Israel

[c] Microsoft Research, Israel

Received 27 March 2017; final version received 15 August 2017; accepted 9 September 2017
Available online 20 September 2017

Abstract

Maximizing the revenue from selling *more than one* good (or item) to a single buyer is a notoriously difficult problem, in stark contrast to the one-good case. For two goods, we show that simple "one-dimensional" mechanisms, such as selling the goods separately, *guarantee at least* 73% of the optimal revenue when the valuations of the two goods are independent and identically distributed, and at least 50% when they are independent.

For the case of $k > 2$ independent goods, we show that selling them separately guarantees at least a $c/\log^2 k$ fraction of the optimal revenue; and, for independent and identically distributed goods, we show that selling them as one bundle guarantees at least a $c/\log k$ fraction of the optimal revenue.

Additional results compare the revenues from the two simple mechanisms of selling the goods separately and bundled, identify situations where bundling is optimal, and extend the analysis to multiple buyers.

☆ Previous versions: February 2012; April 2012 (arXiv:1204.1846 and Center for Rationality DP-606); May 2014; March 2017. Research partially supported by European Research Council Advanced Investigator grant 249159 (Hart) and by Israel Science Foundation grants 230/10 and 1435/14 (Nisan). We thank Motty Perry and Phil Reny for introducing us to the subject and for many helpful discussions, and the referees for their very careful reading and useful comments. A presentation that covers some of this work is available at `http://www.ma.huji.ac.il/hart/abs/2good-p.html`.

* Corresponding author.
E-mail addresses: hart@huji.ac.il (S. Hart), noam@cs.huji.ac.il (N. Nisan).
URLs: http://www.ma.huji.ac.il/hart (S. Hart), http://www.cs.huji.ac.il/~noam (N. Nisan).

http://dx.doi.org/10.1016/j.jet.2017.09.001

JEL classification: C7; D42; D44; D82

Keywords: Revenue maximization; Mechanism design; Auctions; Multiple goods; Bundling; Separate selling; Simple mechanisms

1. Introduction

Suppose that a seller has one good (or "item") to sell to a single buyer whose willingness to pay (or "value") for the good is x. While x is known to the buyer, it is unknown to the seller, who knows only its distribution (given by a cumulative distribution function F). If the seller offers to sell the good for a price p then the probability that the buyer will buy is $1 - F(p)$, and the seller's revenue will be $p \cdot (1 - F(p))$. The seller will choose a price p^* that maximizes this expression.

This problem is the classic monopolist-pricing problem. Looking at it from an auction point of view, one may ask whether there are mechanisms for selling the good that yield a higher revenue. Such mechanisms could be indirect, could offer different prices for different probabilities of getting the good, and so on. Yet, the characterization of optimal mechanisms of Myerson (1981) (see also Riley and Samuelson, 1981 and Riley and Zeckhauser, 1983) concludes that the take-it-or-leave-it offer at the above price p^* yields the optimal revenue among *all* mechanisms. Even more, Myerson's result also applies when there are multiple buyers, in which case p^* would be the reserve price in a second-price auction.

Now suppose that the seller has two (different) goods that he wants to sell to a single buyer. Furthermore, consider the simplest case where the buyer's values for the two goods are independently and identically distributed according to the distribution F ("i.i.d.-F" for short), and where, furthermore, his valuation is additive: if the value of the first good is y and that of the second is z, then the value of the *bundle* consisting of both goods is[1] $y + z$. It would seem that since the two goods are completely independent of each other, then the best one should be able to do is to sell each of them separately in the optimal way, and thus extract exactly twice the revenue one would make from a single good. Yet this turns out to be false.

Example 1. Consider the one-good distribution F taking values 1 and 2, each with probability $1/2$. Let us first look at selling a single good optimally: the seller can either choose to price it at 1, selling always[2] and getting a revenue of 1, or choose to price the good at 2, selling it with probability $1/2$, again obtaining an expected revenue of 1, and so the optimal revenue from a single good is 1. Now consider the following mechanism for selling both goods: bundle them together, and sell the bundle for price 3. The probability that the sum of the buyer's values for the two goods is at least 3 is $3/4$, and so the revenue is $3 \cdot 3/4 = 2.25$—larger than the revenue of 2 that is obtained by selling them separately.

[1] Our buyer's demand is thus *not* limited to one good (as is the case in some of the existing literature; see "unit-demand" in Section 1.1).

[2] Since we maximize revenue we can assume without loss of generality that ties are broken by the buyer in a way that maximizes the seller's revenue. This "seller-favorable" property can always be achieved by appropriate small perturbations of the mechanism; for instance, by the seller giving a small fixed proportional discount on all payments. See Hart and Reny (2015a, Section 1.2, and Remark (a) after Corollary 18).