OPERATIONS RESEARCH (ONLINE FORUM COMMENTARY)
Vol. 55, No. 2, March–April 2007, online commentary
ISSN 0030-364X | EISSN 1526-5463 | 07

INFORMS

Comment on "Thirteen Reasons Why the Vickrey-Clarke-Groves Process is Not Practical" by Michael Rothkopf

Paul J. Healy
The Ohio State University, healy.52@osu.edu, http://healy.econ.ohio-state.edu

Whenever a research topic which has developed in one field spreads to new disciplines, there is the potential benefit of new perspectives and new tools opening unexplored avenues, and there is the potential danger of re-learning the old lessons of the originating discipline. The spread of auction theory and, more generally, mechanism design to computer science, operations research, and related fields is certainly no exception.

On one hand, important practical issues such as the computability of equilibria have long been ignored by the economics profession. On the other hand, computer scientists and operations researchers have focused almost exclusively on the dominant strategy mechanisms developed in the 1970s with no attention to the vast and elegant implementation literature of the 1980s and 1990s, much of which was developed with the focus of finding continually better mechanisms. Fortunately, cross-disciplinary discussions such as this can help both sides reap the benefits of progress without incurring the costs of repetition. My hope is that Michael Rothkopf's critique of the VCG mechanism will lead non-economists interested in mechanism design to explore the implementation literature and to use new tools to generate additional useful insights.[1]

Much of the renewed interest in mechanism design has focused narrowly on auction design, which is certainly due to the simultaneous rise in online auction participation and in research on e-commerce platforms. Similarly, Rothkopf's critiques are aimed only at the VCG mechanism's use in auction environments. The more general definition of the VCG mechanism, however, extends to a very broad class of social choice and allocation problems. Unfortunately, the impossibility results obtained in the more general environments imply that the VCG mechanism is necessarily inefficient; and it is this inefficiency, among other concerns, that led to the development of the implementation literature. Hopefully non-economists, seeing related deficiencies in dominant strategy mechanisms, will continue to follow this path.

The following describes the fundamental difficulty with the VCG mechanism. Suppose we want a mechanism for some very general setting, and we want our agents always to have a dominant strategy that results in a Pareto optimal outcome. Immediately we are faced with the impossibility results of Gibbard (1973) and Satterthwaite (1975), which tells us that if *every* possible preference is admissible, and if the mechanism must always have a dominant strategy, then the outcomes that must be chosen must be exactly those chosen by simply assigning one agent to be a dictator. The easiest escape from this dead-end is to get rid of the assumption that every possible preference is admissible. For example, we could proceed by assuming agents' preferences are of the quasilinear form $v_i(y) + m_i$ where y is the allocation decision and m_i represents holdings of some numeraire

commodity, such as cash, which can be used to tax and reward agents. Proceeding with our objective, we apply the Revelation Principle (Jackson 2003, see, for example,), which says that if agents have a dominant strategy in some mechanism then we can pick an equivalent mechanism where they are asked instead to report their preferences (v_i) and the mechanism designer then 'plays the dominant strategy' of the original mechanism on the players' behalf. Thus, we can restrict our search to 'preference-revealing' mechanisms. But then (Green and Laffont 1977, Theorem 3) tell us that the *only* mechanism in which truth-telling is a dominant strategy is the class of VCG mechanisms. Since it is known, however, that the VCG mechanisms cannot always achieve budget balance, we again reach a dead-end; we can guarantee a Pareto optimal allocation *decision* y, but we must sometimes collect a surplus of transfers to achieve it. Further restrictions on preferences can guarantee budget balance, but these restrictions are severe (Liu and Tian 1999, see).[2]

Because of these dead-ends, economists more-or-less abandoned the VCG mechanism (outside the auctions literature) through the 1980s and 1990s, working instead on mechanisms whose desirable outcomes obtain from Nash equilibrium behavior (or refinements of Nash equilibrium) instead of dominant strategy behavior. In fact, a special issue of the journal *Public Choice* (volume 29-2, Spring 1977) was dedicated to the VCG mechanism and its many shortcomings. For example, it isn't always budget balanced, it can generate highly inequitable (though efficient) outcomes, it can leave individual agents worse off than before the mechanism was run (even causing bankruptcy), it is not impervious to deviations by coalitions (in fact no dominant strategy mechanism is), and the requirement of quasilinear preferences ($v_i(y)+m_i$) is crucial. Although some of Rothkopf's critiques extend at least as far back as this 1977 discussion (especially revenue deficiency and collusion by coalitions), others have only recently been explored (such as the preparation and communication costs of bidding, which is a topic that economists are not well suited to handle).

The limitations of the VCG mechanism suggests to me two avenues: First, as I have already suggested, much could be gained by revisiting the implementation literature with a critical eye for these practical limitations. Second, once the theorists have identified a set of candidate mechanisms whose shortcomings are acceptably small by some measure, the question of which to use in a real-world setting is best answered through empirical investigation.

Suppose that, for a given allocation problem, a set of candidate mechanisms has been identified. Ideally, one would like to test each mechanism in the real-world setting to discover which realizes the 'best' outcomes, according to some metric. The obvious flaw with this approach is that such tests are rarely practical and highly risky. The more subtle flaw with real-world implementation is that, since preferences are not observable, the efficiency of the mechanism outcome is not measurable. The work-around to this problem is in laboratory testing, where cash payments are used to induce known preferences in real human subjects, who then participate in the mechanisms of interest. Here, the laboratory serves as a wind tunnel or 'testbed' in which we can examine mechanism properties in the face of real behavior with real incentives, but on a manageable scale with more control and observability.

Laboratory auction experiments are by far the most common mechanism-comparison studies, dating back at least as far as Coppinger et al. (1980). A survey of the results is far beyond the scope of this comment (Kagel 1995, see) for an excellent survey through 1995), but the general finding is that several theoretical predictions such as revenue equivalence fail to hold in reality. Though, to my knowledge, there is still no agreement about the underlying cause of these theoretical failures, the empirical result is immediately useful: an auctioneer interested in maximizing revenue can look to these data for advice. He can also look to these data for signs of some of Rothkopf's (and others') concerns about the VCG (second-price) auction, such as collusive bidding. If one of the concerns is not addressed by an existing experiment, a new experiment can be run to test the seriousness of the issue. For example, communication costs can be explicitly added to the experimental design, or subjects could be asked to participate in a sequence of related auctions rather than a single, isolated auction.

The following is a good example of how experiments can inform the mechanism design process. In the domain of public goods provision problems, experiments comparing mechanisms have been run by Chen and Plott (1996), Chen and Tang (1998), Attiyeh et al. (2000), Cason et al. (2003), and Healy (2006). One apparent lesson from the Chen & Plott, Chen & Tang, and Healy papers is that dynamics matter: if a mechanism's Nash equilibrium is dynamically stable (under the appropriately defined dynamic), then actual behavior converges to the efficient equilibrium as the mechanism is repeated. If the equilibrium is unstable, behavior cycles and outcomes remain inefficient. Thus, a mechanism designer should add stability to his list of desiderata if the mechanism is to be implemented repeatedly with a fixed set of agents.

A second lesson from the Attiyeh *et al.*, Cason *et al.*, and Healy studies is that, in the VCG mechanism, demand revelation occurs with surprisingly low frequencies. The design of the Healy experiment suggests a reason, which is first among Rothkopf's concerns: the truth-telling dominant strategy is (in most settings) only a weak equilibrium, meaning agents are indifferent between truth-telling and some other announcement. Subjects in the experiment who fail to reveal truthfully often submit announcements that are in fact best responses, meaning they experience no loss in earnings from these particular misrevelations. But the important question (usually) isn't whether or not people truthfully reveal their preferences; it's whether or not we realize efficient outcomes. In fact, these misrevelations have little impact on efficiency because, as theory predicts, the non-truthful best response strategies are only best responses if they lead to the same (efficient) outcome as truth-telling. Thus, we should be concerned that the weak equilibrium property of the VCG mechanism will have an effect on behavior, but perhaps we should not be overly concerned about its effect on outcomes.

It is my hope that researchers outside the field of economics will raise new concerns that can be explored using a similar methodology. By understanding how the economics literature has evolved through impossibility results, implementation theory, and now experimental testing of mechanisms, we can make the introduction and integration of computer science and operations research issues and techniques more efficient. Michael Rothkopf's paper should provide the impetus for theorists to move forward beyond the VCG mechanism, which will not only provide exciting new problems and solutions in mechanism design theory, but will also provide new and valuable laboratory tests to help transition from theory to reality.

Endnotes

[1]Excellent surveys and reviews of the implementation literature include Jackson (2001), Jackson (2003), Moore (1995), Palfrey (2002), Palfrey and Srivastava (1993), Hurwicz (1994), and Groves and Ledyard (1987). For those who would prefer only to test the water without opening an economics journal might begin with Dash et al. (2003).

[2]In an auction environment, the net transfers go to the seller (who is not considered a player of the mechanism), so the failure of budget balance simply means that the auctioneer earns positive revenues. As Rothkopf points out, when there are multiple buyers and sellers, sellers become agents in the mechanism and the budget balance failure becomes an issue.

References

Attiyeh, Greg, Robert Franciosi, R. Mark Isaac. 2000. Experiments with the pivot process for providing public goods. *Public Choice* **102** 95–114.

Cason, Tim, Tatsuyoshi Saijo, Tomas Sjostrom, Takehiko Yamato. 2003. Secure implementation experiments: Do strategy-proof mechanisms really work? California Institute of Technology Department of Humanities and Social Sciences Working Paper.

Chen, Yan, Charles R. Plott. 1996. The groves-ledyard mechanism: An experimental study of institutional design. *Journal of Public Economics* **59** 335–364.

Chen, Yan, Fang-Fang Tang. 1998. Learning and incentive-compatible mechanisms for public goods provision: An experimantal study. *Journal of Political Economy* **106** 633–662.

Coppinger, Vicki, Vernon L. Smith, John Titus. 1980. Incentives and behavior in english, dutch, and sealed-bid auctions. *Economic Inquiry* **18** 1–22.

Dash, Rajdeep K., Nicholas R. Jennings, David C. Parkes. 2003. Computational mechanism design: A call to arms. *Intelligent Systems, IEEE* **18** 40–47.

Gibbard, Allan. 1973. Manipulation of voting schemes: A general result. *Econometrica* **41** 587–601.

Green, Jerry, Jean-Jacque Laffont. 1977. Characterization of satisfactory mechanisms for the revelation of the preferences for public goods. *Econometrica* **45** 783–809.

Groves, Theodore, John O. Ledyard. 1987. Incentive compatibility since 1972. Theodore Groves, Roy Radner, Stanley Reiter, eds., *Information, Incentives, and Economic Mechanisms*. University of Minnesota Press, Minneapolis.

Healy, Paul J. 2006. Learning dynamics for mechanism design: An experimental comparison of public goods mechanisms. *Journal of Economic Theory* **129** 114–149.

Hurwicz, Leonid. 1994. Economic design, adjustment processes, mechanisms, and institutions. *Economic Design* **1** 1–14.

Jackson, Matthew O. 2001. A crash course in implementation theory. *Social Choice and Welfare* **18** 655–708.

Jackson, Matthew O. 2003. Mechanism theory. Ulrich Derigs, ed., *Encyclopedia of Life Support Systems*. EOLSS Publishers, Oxford.

Kagel, John H. 1995. Auctions: A survey of experimental research. John H. Kagel, Alvin E. Roth, eds., *The Handbook of Experimental Economics*. Princeton University Press, Princeton.

Liu, Liqun, Guoqiang Tian. 1999. A characterization of the existence of optimal dominant strategy mechanisms. *Review of Economic Design* **4** 205–218.

Moore, John. 1995. Implementation, contracts, and renegotiation in environments with complete information. *Advances in Economic Theory: Sixth World Congress*, vol. 1, chap. 5. Cambridge University Press, Cambridge.

Palfrey, Thomas R. 2002. Implementation theory. *Handbook of Game Theory*, vol. III, chap. 64. North Holland, Amsterdam.

Palfrey, Thomas R., Sanjay Srivastava. 1993. *Bayesian Implementation*. Harwood Academic Publishers, New York.

Satterthwaite, Mark. 1975. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *J. Econ. Theory* **10** 187–217.