American Economic Review 2008, 98:3, 577–585
http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.3.577

But Who Will Guard the Guardians?

By Leonid Hurwicz*

In posing the famous question,[1] the Roman author, Juvenal, was suggesting that wives cannot be trusted, and keeping them under guard is not a solution—because the guards cannot be trusted either.

Half a millennium or so earlier, Plato,[2] in *The Republic,* expressed a more optimistic view regarding the guardians or rulers of the city-state, namely that one should be able to trust them to behave properly; that it was absurd to suppose that they should require oversight.

Socrates, referring to an earlier statement[3] that "drunkenness is most unbefitting guardians," says: "From intoxication we said that they must abstain. For a guardian is surely the last person in the world to whom it is allowable to get drunk and not know where on earth he is." To which Glaucon, Socrates' interlocutor, replies: "Yes, it would be absurd[4] that a guardian should need a guard." Instead of Juvenal's later pessimism, indeed cynicism, Plato—through Glaucon—expresses the optimistic view that one should be able to trust the city's guardians and rulers to behave properly; that they should require oversight is an absurdity.

A casual perusal of daily newspapers indicates that there is nothing absurd about the present day "guardians"—by which I mean leaders and officials of political, economic, and social entities—needing, and indeed getting, a great deal of oversight.[5] The question is rather as to the extent oversight is, or even can be, effective. Publicly expressed attitudes toward government, law enforcement, and union and corporate leadership are often more reminiscent of Juvenal than Plato.

Parallel questions in economics have been analyzed using the idea of the institution,[6] in the sense of institutional arrangements or "rules of the game" (rather than organizations, which are referred to as "artificial players"). This economic analysis focuses on implementation (in particular, the rule of law and its enforcement) as essential to the functioning of institutions.[7]

† This article is a revised version of the lecture Leonid Hurwicz delivered in Stockholm, Sweden, on December 8, 2007, when he received the Bank of Sweden Prize in Economic Sciences in Memory of Alfred Nobel.

* Department of Economics, University of Minnesota, 1035 Heller Hall, 271-19th Avenue South, Minneapolis, MN 55455 (e-mail: hurwicz@umn.edu.) Thanks are due to Eric Maskin, Roger Myerson, and members of my family, especially Michael and Maxim Hurwicz, for their editorial suggestions. Remaining imperfections are mine.

[1] "Sed quis custodiet ipsos custodes?" (D. Juvenal 1895, Liber secundus, Satura VI, 325, lines 347–8).

[2] 1930, Book III, XII, 403E, volume I, p. 264 (Greek) and p. 265 (English).

[3] Ibid., pp. 246–7, no. 398E.

[4] In B. Jowett's 1908 translation, "ridiculous" replaces "absurd." (Jowett's rendering seems better because the etymology of the Greek word used by Glaucon appears to be "laughable"—as is that of "ridiculous.") But the recent translation by A.D. Lindsay (1992) again uses "absurd."

[5] Without explicitly mentioning the classical precedents, the *Wall Street Journal* of May 8, 1998 (B1–B2) carries a story under the headline "Guardians May Need Someone to Watch Over Them." It speaks of court-appointed guardians and conservators who dishonestly dissipate their wards' assets, and of the difficulties the courts have in exercising their oversight responsibilities. In turn, watchdog groups and legislative task forces try to improve the performance of the judicial branch.

[6] In the sense of institutional arrangements ("rules of the game") rather than entities such as various types of organizations ("artificial players").

[7] Elinor Ostrom, James Walker, and Roy Gardner (1992) point out that emphasis on the importance of enforcement is found in Hobbes, although they disagree with his stress on the need for an external enforcing authority (the sovereign). Their stress on intra-group mutual enforcement. Andrew Schotter (1981, 11) makes

Some of the most basic contemporary policy issues involve choice of institutions: markets versus central planning, the scope and structure of social insurance (unemployment, old age, health), "property rights" as solutions to problems posed by externalities, world free trade, and the degree of economic integration of independent nations are obvious examples. The question is whether the role of institutions can be captured by appropriate analytical tools and incorporated into economic models—metaphorically, whether institutions can be introduced into models as variables, even as unknowns, rather than as fixed parts of the landscape (as is, for instance, perfect competition in so many mainstream models). If we can construct such models, then we can also consider incorporating implementation devices, with their limitations and potentials.

Much economic analysis is based on the perfectly competitive model, which implies strong assumptions concerning two things: 1) the information available to "agents" (individuals, firms, etc.) engaging in economic activity, and 2) the existence of implementation mechanisms such as enforcing contracts and preventing collusion. Similarly, conclusions concerning the effects of alternative forms of taxation, subsidies, or social insurance depend in an essential way on implementation mechanisms supplying information concerning obligations and entitlements, entities facilitating financial flows, as well as enforcement of payments and disclosure of relevant information. Until recently, much more attention has been paid to information requirements (and uncertainty when precise information is not available) than to implementation. Yet if implementation is impossible or prohibitively costly, even the most attractive mechanism remains a utopia.

In the 1960s and 1970s, researchers studying informationally decentralized and game-modeled mechanisms addressed those issues.

I. Message Exchange (Non-Game-Theoretic) Processes

We now proceed to the process of message exchange, an analysis that does not involve game theory.

A process is defined as informationally decentralized when each economic unit (such as a consumer or firm) initially has information only about itself (its preferences, technology, or resources) but not about other units. The process then requires an exchange of signals (called messages) to attain objectives such as efficiency of the system as a whole. (The goal, for instance, could be Pareto optimality or efficiency, in which no one can be made better off without someone else being made worse off.) By quantifying the information in the signals—for example, how many units or what the dimension is of the spaces used—in the context of rigorously formulated models, it has been shown that it may be impossible in certain cases to guarantee finite dimensional message space. An example is economies with detrimental externalities, such as pollution or infinite time horizon. This impossibility applies to all informationally decentralized mechanisms, even if one assumes that all participants are truthful and abide by prescribed rules: there are cases where no rules can guarantee the attainment of the desired objectives (for instance, efficiency). Thus, enforceability of rules is not always an issue.

But what if we don't assume that participants are always truthful? Samuelson raised this question in the mid-1950s in classic articles on public goods, and in particular on the so-called Lindahl solution, in which individuals pay for a public good according to their marginal willingness to pay: that is, how eager they are to pay to achieve a supply of some specified public good. Samuelson pointed out that the Lindahl mechanism did not take care of the free rider problem, in which individuals gain an advantage by understating their interest in public goods. Samuelson

self-policing or external policing authority an integralpart of the definition of a social institution. He recognizes the possibility of intra-group enforcement through a supergame model (165, note 8).

went on to conjecture that the same problem would arise in any decentralized mechanism allocating public goods. Could these two points be formalized and rigorously justified?

II. Game-Theoretic Framework

If we use a game theoretic framework, how does this go?

A natural framework for such analysis turned out to be Nash's theory of noncooperative games and the concept of Nash equilibria, in which no player gains by changing only his or her own strategy unilaterally. Viewing the economic process as a noncooperative game in which the strategy of each participant is a statement (not necessarily truthful) about his/her preferences, and the "payoff" is the utility of the outcome for that participant (defined in Samuelson's analysis by the Lindahl formula), Samuelson's first point can be formalized as the proposition that truthful revelation of preferences is not a Nash equilibrium in the Lindahl game. That is, if all other players have told the truth, and it is now your turn, you will in general gain by misrepresenting your preferences. So, Samuelson's first point turned out to be correct: truth is not a Nash equilibrium.

However, Samuelson's second point, that illegal behavior would be advantageous in any decentralized system, allows consideration of different criteria for judging system performance (for example, Pareto optimality or degree of income inequality) and different kinds of decentralized mechanisms, including one proposed by Theodore Groves and John O. Ledyard (1977), which admitted as strategies any signals ("messages"), not necessarily just revelations of one's own preferences.

In fact, in the simplest version of the Groves-Ledyard mechanism, a participant is just told, "Give me any number within a certain set of limits." The participant is not answering any question, and not necessarily revealing anything.

Groves and Ledyard constructed a nonrevelation (which should really be understood as "not necessarily revelation") model yielding Pareto-optimal (though not Lindahl) Nash equilibrium outcomes in economies with three or more participants. Later, in Hurwicz (1979a), it was shown that even Lindahl outcomes could be obtained as Nash equilibria of suitably designed nonrevelation games. If one regards this broader (not necessarily revelation) class of mechanisms as decentralized—and that's a very important "if"—Samuelson's second (impossibility) point cannot be accepted. Hence, optimism replaces pessimism. However, is there still a problem of enforcement?

III. Are Nash Equilibria Self-Enforcing?

I believe the answer is, "Yes, there is a problem of enforcement."

Now we come to a very important and closely related issue. We are asking whether a given Nash equilibrium in a specified game is enforceable or not. And one occasionally hears the claim that there can be no enforcement problem with Nash equilibria because allegedly Nash equilibria are self-enforcing. I want to stress that I am denying this claim, but I want to give the other side an opportunity to give their arguments. Their argument is that, by definition, in a Nash equilibrium, no player can profit by a unilateral departure from his or her equilibrium strategy. Furthermore, collusions are infeasible in a noncooperative game. Hence, the argument goes, there is no need (or possibility, really) for enforcement.

But this argument assumes, first, that players will not or cannot cheat, and that they will consider only strategies prescribed by the mechanism governing the system, what we call the "legal" strategies. In addition, it assumes machinery (involving not only enforcement but also information processing and financial procedures) to ensure that specified outcomes will in fact be produced.

These assumptions are accepted as given in the usual argument about Nash equilibria. However, if you question these assumptions, then there is no doubt that enforcement is a valid question to pose, even in Nash equilibria.

IV. Formalizing Rules of a Game as Game-Forms (Mechanisms)

We now proceed to another important issue, namely, formalizing rules of a game as "game forms" or, synonymously, "mechanisms."

A game is defined by the players' strategy domains and payoff functions. A mechanism or game form (rules of the economic game), shares with the game the strategy domains, but differs in that it involves an outcome function. Let me stress that I am distinguishing between a game form and a game: they have strategy domains in common. In the case of a game, what is added to that are payoff functions. In contrast, a mechanism involves only physical outcomes. Each may be viewed as a pair where the first component of each pair is a strategy domain associated with various individuals. The second component of a *game form* is what we call an *outcome function* associating a physical outcome with a given point in the domain. So, for example, the outcome of an election is the player who got elected. On the other hand, the "payoff" component in the concept of a *game* is the listing of how players feel about the outcome as measured by the numerical value of the utility using the strategy chosen. For example, Mr. Bush is the outcome of the 2000 year election. The "payoff" is how players feel about the outcome. The "outcome" is the physical description of what happened, depending on what kinds of strategies were chosen. And it may happen that strategies that were chosen by some agents were not legal strategies, because somebody decided to risk being illegal.

We assume that each player has a utility function, associating a real numerical value of the outcome for that participant with any conceivable outcome of the game.

In a model devoted to the analysis of institutional arrangements, it is essential to separate that which belongs to the data ("environment")—thus preferences, endowments, and technologies—from that which is subject to human manipulation, in our model the strategy domains and the outcome functions. It is natural to refer to a "strategy-outcome pair," that is, the pair consisting of strategy choices and the outcome corresponding to them. It is natural to refer to these pairs as the *rules of a game* since the strategy domain defines the legal moves, and the outcome function defines their consequences.

On the other hand, in game theory jargon, the "strategies-outcomes" pair is called a *game-form* (hyphenated); in economics, it's called a *mechanism*. The mechanism is not affected by changes in preferences, but the payoff functions are. The game-form can be changed directly by legislation or other human actions; the payoff functions only indirectly through changes in the game-form. The outcome is an objective fact, but the payoff combines that fact with the participant's preferences for the outcome. The preferences may be the positive and negative.

A need for enforcement implies the possibility of behavior that violates the rules of the game. The point is that, if there were no possibility of violation, then you wouldn't need enforcement. We refer to such strategies as illegal, as distinct from legal, of course.

To represent this, we introduce the set of all feasible actions, legal and illegal, and we call that the *true* strategy domain. Thus the set of true strategies consists of two feasible and disjoint subsets, that of legal strategies and that of illegal strategies. That means that this person we have in mind is turning over in his or her mind, "Should I use a legal strategy, or maybe I would be better off to risk the illegal?" And that is what we call true strategy (as distinct from talking only about legal strategies). Some very ethical or law-abiding person might never consider anything illegal. But we are talking about people who might risk it. And if that's the case, and they may have chosen an illegal strategy, then you are interested to know what the true outcome function is. The

decision was illegal, but it does produce consequences. We consider participants to be playing the "*true* game" whose game-form may contain both legal and illegal strategies, although it is the legal game we want them to play.

V. Successful Enforcement and Implementation

Now, what is successful enforcement? You might say, "Doesn't everybody know what successful enforcement means?" But, for example, whether or not Samuelson's second conjecture is true may determine how success is measured. So it's important to specify this.

This framework makes it possible to formalize the notions of enforcement and implementation.

A. *Successful Enforcement*

To say that the *legal game rules are being successfully enforced* means that the outcomes of the true game ensure that illegal strategies are less attractive than legal strategies. A strong formulation of successful enforcement might require that, for every player, every illegal strategy is *dominated* by (that is, is less attractive than) some legal strategy. A "weak" domination would require only that a player at least be no worse off by staying within the law.

In fact, however, if everyone else is acting illegally, a normally law-abiding player may not find it advantageous to remain law-abiding. It seems, therefore, more reasonable to adopt a somewhat weaker concept of successful enforcement.

Now we study further the implementation idea.

B. *Implementation*

Consider what is involved in making effective an institution such as Social Security. The desired mechanism—say as defined by legislation—specifies the class of persons receiving payments, say in relation to previous income and other variables. It may also specify the sources of funding. There is a need to formulate a *modus operandi*, verify whether specific applicants are entitled to receive payments and, if so, at what level, and how to collect the required funds. Typically, special agencies are created to accomplish such tasks, both informational and related to enforcement. It is the complex of such activities and arrangements that I think of as the effort to implement the legislation.

I make this remark here about what I think of as implementation because, for example, one of the other people who, along with myself, received the Nobel prize for mechanism design uses a narrower definition of implementation. You see, among other things, when I am talking about implementation—or sometimes I refer to this as "genuine implementation"—it means that we have the money to run the institutions, we have the information to run the institutions, and there is legislation authorizing this. Some people, as I say, have a definition that doesn't ask for quite that much, but this would require a more elaborate model than the one I am using myself elsewhere in this paper.

And I conclude this section by saying that *implementation is successful* if the equilibrium outcomes correspond to those of the desired game, i.e., those envisioned by the legislation.

Expressed in this framework, a Nash equilibrium is not self-enforcing because, while it is unprofitable to move to alternative legal strategies, it may be profitable, in the absence of enforcement, to move to illegal strategies.

Similarly, Nash equilibria are not self-implementing because implementing actions are required to ensure that the true outcome is the same as the legal outcome.

VI. Back to Juvenal

Let us now come back to the initial problem, the need for guarding the guardians. Juvenal's cynical question suggests either that there is no way to guard the guardians, or that, in addition to having "guardians of the first level" (those guarding the wives), one must also, according to him, presumably have "guardians of the second level" to guard the guardians of the first-level. But the need to guard second-level guardians conjures the image of an infinite regress of guardians. Since an infinity of guardians is not usually available, this seems to preclude enforcement! Even if some guardians have multiple clients (and that means that you don't need an infinity of them), the capacity for guarding may be exhausted, precluding successful enforcement.

So, from Juvenal one gets a very pessimistic view of human nature. But the question is, "Is it quite so bad?" Let's look at this question without using high-powered statistical tools. Economists (some of them at least) refer to that as "casual empiricism"—that is, just what one observes without being high-powered mathematically, statistically, or whatever.

A. *Empirical—Not So Bad*

This "casual empiricism" suggests that there are many situations where rules are substantially (if not perfectly) implemented and/or enforced. But many factors are in the picture determining whether implementation of rules is possible.

B. *Physical*

We shall now discuss different kinds of situations with respect to enforcement. But let me make a point: at some place there may be a remark such as "enforcement (or generally in implementation)." Why is that? Well, for example, part of an implementation program may be not taxing people, but giving subsidies, for instance, to poor people to keep children healthy. So that is also part of implementation. But you wouldn't call giving extra money to somebody enforcement. So that illustrates why implementation is a more general concept than enforcement. But both are very important.

First, though perhaps least important, implementation may depend on purely *physical or mechanical* factors that do not require human guardians. Examples: devices in parking lots that slash the tires of those who use unauthorized exits; or punishment by placing the culprit on an isolated island without a boat, and too far to swim ashore. That is from the customs of some tribes, at least in the past, in Washington state, where they punished people in that particular way. The point is that this relies on physical or mechanical factors and not on any police or anything of that sort.

C. *Intervenors*

Second, and more importantly, there may be guardians (individual or collective) who are in sympathy with the rules, for instance, whose ethical standards rule out corrupt behavior. Such persons or groups, whom we will call *intervenors* (that's my own term—I don't think anybody else uses it), may have the ability (through power, financial assets, personal charisma, or status combined with the population's respect), as well as the inclination, to discourage improper behavior of lower-level guardians. In this case, the rule has a chance to be successfully enforced. Well-functioning societies try to choose judges and rulers from among such individuals. So my point here is that, just because there are people who violate rules, behave illegally, and so on, it does not mean that everybody is like that. On the contrary, if you watch certain judges, you

observe that they honestly try to implement what they believe the constitution says and just put it into effect. And that's why I invented a word, a term, to describe people of that kind. And I think maybe they are rare, but not nonexistent.

D. *Elective Office Helps*

Juvenal's pessimistic question suggests that there are no intervenors. Even in that case, there are systems conducive to successful enforcement. Consider, for example, a law designed to protect citizens from harmful or dangerous behavior of certain individuals. Suppose that those charged with enforcement of the law (first-level guardians) are corrupt or otherwise ineffective, and so are their supervisors (second-level guardians). If the latter hold elective office, citizens who are voters (and of course, we assume free voting) may function as third-level guardians (as well as being guarded by first-level guardians). These voters have both an incentive and the power to intervene by throwing the supervisors out of office. This gives the supervisors an incentive to make sure that first-level guardians discharge their duties properly. Effective enforcement may be the result. So that's really one role of democracy in a system.

I anticipate that some people don't believe in the existence of intervenors. I happen to. But, even if you don't believe in intervenors, you still don't necessarily have to be as cynical as Juvenal.

As a special case, one can imagine a situation with a finite number of guardians, say guards 1, 2, and 3. Guard 1 watches the treasure. Guard 2 watches guard 1. Guard 3 watches guard 2. But then guard 1 also watches guard 3. If "watching" implies being successful in creating truthful behavior, then the "circle" represented by such a system is successful. In this circle, any guard can also be an intervenor.

The concept of intervenors can be used to compare Glaucon's views on the one hand with Juvenal's on the other. In fact, Glaucon assumed that guardians are intervenors.

E. *Without Intervenors*

I start with the general statement that rational, self-interested behavior may result in desired legal outcomes, even without intervenors.

In particular, there is no infinite regress where the electorate, directly or indirectly, can make a difference. For instance, suppose voters remove a corrupt official from office. The next, newly elected official, out of rational self-interest, might refrain from corrupt practices, given that the people had demonstrated their willingness to remove him or her from office based on such practices. This does not require an infinite regress of dishonest guardians, merely an aroused public, exhibiting rational self-interest.

As an example of success without intervenors, members of an agency were accepting bribes, and giving a percentage to their supervisors, who in turn passed a percentage on to their supervisors. Higher judicial officials, however, set up a sting operation that resulted in the arrest and prosecution of the corrupt members of the agency. The officials were not elective. However, the sting operation advanced the careers of the officials, many of whom had higher aspirations. It also protected them from possible blame for the corruption. Thus, it is reasonable to believe that the officials arranged the sting operation out of rational self-interest. In any case, there was not an infinite regression of dishonest guardians.

I am not saying that rational self-interested behavior will always result in noncorrupt behavior. However, the example of the officials does, I think, provide an illustration of rational self-interested behavior leading to truthful and uncorrupted behavior within a finite (not infinite) regression of guardians.

Similarly, a judge may want to advance to the post of attorney-general, and the mayor may aspire to being governor or president. Given such aspirations, and fear of negative consequences, we cannot rule out legal (as distinct from illegal) behavior, even in the absence of intervenors.

When implementation is possible, it can be modeled in terms of the relationship of the true game to the desired mechanism. The equilibria of the true game depend on what is feasible, on the actual consequences, and on the preferences of the players. Thus, to achieve equilibria of the true game (not only the legal part of it, but including the illegal), the intervenor's utility function should have preferences that rule out his/her own illegal behavior. On the other hand, strategies available to the intervenor should include powers and assets enabling him/ her to influence others in the desired direction. These features, including strategies available to the intervenor and the consequences of their use, are contained in the true game-form mechanisms.

Depending on the participants in the game (i.e., composition of the society) and their preferences/values, a given desired mechanism may or may not be implementable within that society. The history of prohibition, namely, that it was never effectively enforced, may illustrate the negative case. Clearly, implementability depends on the nature of the goal and attitudes toward likely outcomes.

VII. Ethics

Legal behavior is not incompatible with rational, self-interested behavior. This means that, under the assumption of rational self-interest, Juvenal's claims about corrupt behavior being unavoidable cannot be accepted. And this may be true even without counting on intervenors.

What are our conclusions? Those conclusions are divided in two groups: those that have to do with the claim by Juvenal (or similar claims) on the one hand, and then claims or conjectures by Samuelson. But genuine implementability may exist, even without the presence of intervenors.

So now, in the first group: bad outcomes with an infinite regression are logically possible, but enforcement (or more generally, implementation) is not always impossible either, despite Juvenal's claim or question. For one thing, intervenors may exist. Also, particularly in a society where voters' preferences are freely expressed, citizens may act as top-level guardians. By "top level," we mean above all other guardians. Finally, it may be possible to create systems in which rational self-interest, due to a variety of rewards and punishments, leads, or may lead, to legal and ethical behavior.

The validity of Samuelson's second conjecture (that is, that there may be no successful implementation for public goods under decentralization) remains, in my opinion, an open question. The example of voters removing a corrupt official from office suggests that successful decentralization is possible. But the answer may depend on how we define decentralization and what class of mechanisms is considered relevant for alternatives to Lindahl. On the other hand, Samuelson's first point is clearly valid; namely, that the Lindahl approach does not address situations where there may be a Nash equilibrium, but with illegalities or free riding. I am not aware whether more recent literature addresses such issues.

REFERENCES

Calsamiglia, Xavier. 1977. "Decentralized Resource Allocation and Increasing Returns." *Journal of Economic Theory*, 14(2): 263–83.

Groves, Theodore, and John O. Ledyard. 1977. "Optimal Allocation of Public Goods: A Solution to the 'Free Rider' Problem." *Econometrica*, 45(4): 783–809.

Hurwicz, Leonid. 1972. "On Informationally Decentralized Systems." In *Decision and Organization*, ed. C. B. McGuire and Roy Radner, 297–336. Amsterdam: North-Holland.

Hurwicz, Leonid. 1977. "On the Dimensional Requirements of Informationally Decentralized Pareto-Satisfactory Processes." In *Studies in Resource Allocation Processes*, ed. Kenneth J. Arrow and Leonid Hurwicz, 413–24. Cambridge, MA: Cambridge University Press.

Hurwicz, Leonid. 1979a. "Outcome Functions Yielding Walrasian and Lindahl Allocations at Nash Equilibrium Points." *Review of Economic Studies*, 46(2): 217–25.

Hurwicz, Leonid. 1979b. "Balanced Outcome Functions Yielding Walrasian and Lindahl Allocations at Nash Equilibrium Points for Two or More Agents." In *General Equilibrium, Growth and Trade*, ed. Jerry R. Green and Jose Alexandre Sheinkman. New York: Academic Press.

Hurwicz, Leonid. 1993. "Implementation and Enforcement." In *Political Economy, Institutions, Competition, and Representation*, ed. William A. Barnett, Melvin J. Hinrich, and Norman J. Schofield, 51–59. Cambridge, MA: Cambridge University Press.

Hurwicz, Leonid. 1996. "Institutions as Families of Game Forms." *The Japanese Economic Review*, 47(2): 113–32.

Hurwicz, Leonid. 1997. "Feasible Balanced Outcome Functions Yielding Constrained Walrasian and Lindahl Allocations at Nash Equilibrium Points in Economies with Two Agents when the Designer Knows the Feasible Set." Paper presented at the Pennsylvania State University Decentralization Conference, University Park, PA.

Juvenal, D. 1895. *Junii Juvenalis Saturarum Libri V.* Mit Erklärenden Anmerkungen von Ludwig Friedlaender, Erster Band. Leipzig: Verlag von S. Hirzel.

Maskin, Eric S. 1965. "The Theory of Implementation in Nash Equilibrium: A Survey." In *Social Goals and Social Organization, Essays in Memory of Elisha Pazner*, ed. Leonid Hurwicz, David Schmeidler and Hugo Sonnenschein, 173–204. Cambridge, MA: Cambridge University Press.

Maskin, Eric S. 1977. "Nash Equilibrium and Welfare Optimality." Unpublished.

Ostrom, Elinor, James Walker, and Roy Gardner. 1992. "Covenants With and Without a Sword: Self-Governance is Possible." *American Political Science Review*, 86(2): 404–17.

Plato. 1930. *The Republic.* [ΠΟΛΙΤΕΙΑ]. Trans. Paul Shorey. New York: G. P. Putnam's Sons.

Plato. 1992. *The Republic*. Trans. A. D. Lindsay. London: Everyman's Library.

Plato. 1908. *The Republic*. Trans. B. Jowett. Oxford: Clarendon Press.

Schotter, Andrew. 1981. *The Economic Theory of Social Institutions*. Cambridge, MA: Cambridge University Press.

Von Hayek, Friedrich August. 1945. "The Use of Knowledge in Society." *American Economic Review*, 35(4): 519–30.

This article has been cited by:

1. Myerson Roger B.. 2009. Learning from Schelling's Strategy of ConflictLearning from Schelling's Strategy of Conflict. *Journal of Economic Literature* **47**:4, 1109-1125. [Abstract] [View PDF article] [PDF with links]
2. Michael J. Prietula, Daniel Conway. 2009. The evolution of metanorms: quis custodiet ipsos custodes?. *Computational and Mathematical Organization Theory* **15**:3, 147-168. [CrossRef]