

Kalai, Ehud; Ledyard, John O.

Working Paper

Repeated Implementation

Discussion Paper, No. 1205

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Kalai, Ehud; Ledyard, John O. (1997) : Repeated Implementation, Discussion Paper, No. 1205, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/221561>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Discussion Paper No. 1205R

REPEATED IMPLEMENTATION^{*}

by

Ehud Kalai^{*}

and

John O. Ledyard^{**}

August 1996

Revised March 1998

JEL Classifications:

C7 - Game Theory and Bargaining Theory
D7 - Analysis of Collective Decision Making

^{*}Department of Managerial Economics and Decision Sciences, J. L. Kellogg Graduate School of Management, Northwestern University, 2001 Sheridan Road, Evanston, Illinois 60208. Tel: (847) 491-7017, E-mail: kalai@nwu.edu

^{**}Division of Humanities and Social Sciences, California Institute of Technology, 1200 California Blvd., Pasadena, CA 91125. Tel: (818) 395-8482, Fax: (818) 405-9841, E-mail: jledyard@caltech.edu

^{*}The authors wish to thank Andreas Blume, Tim Feddersen, Alvaro Sandroni, and a referee of this journal for helpful suggestions. Kalai's research is partly supported by NSF Economics Grant No. SBR-955421. Ledyard's research is partly supported by the New Millennium Project Office of the Jet Propulsion Laboratory of NASA.

REPEATED IMPLEMENTATION⁺

by

Ehud Kalai^{*}

and

John O. Ledyard^{**}

August 1996

rev. March 1998

JEL Classifications:

C7 - Game Theory and Bargaining Theory
D7 - Analysis of Collective Decision Making

^{*}Department of Managerial Economics and Decision Sciences, J. L. Kellogg Graduate School of Management, Northwestern University, 2001 Sheridan Road, Evanston, Illinois 60208. Tel: (847) 491-7017, E-mail: kalai@nwu.edu

^{**}Division of Humanities and Social Sciences, California Institute of Technology, 1200 California Blvd., Pasadena, CA 91125. Tel: (818) 395-8482, Fax: (818) 405-9841, E-mail: jledyard@caltech.edu

⁺The authors wish to thank Andreas Blume, Tim Feddersen, Alvaro Sandroni, and a referee of this journal for helpful suggestions. Kalai's research is partly supported by NSF Economics Grant No. SBR-955421. Ledyard's research is partly supported by the New Millenium Project Office of the Jet Propulsion Laboratory of NASA.

Abstract

In the traditional static implementation literature it is often impossible for implementors to enforce their optimal outcomes. And when restricting the choice to dominant-strategy implementation, only the dictatorial choices of one of the participants are implementable.

Repeated implementation problems are drastically different. This paper provides a strong implementation “folk theorem”: for patient implementors, every outcome function they care about is dominant-strategy implementable.

1. Introduction

Moving from static one shot environments to dynamic repeated ones can enrich the implementation literature in several ways. This paper highlights these possibilities by offering a formal detailed analysis of a simple example that points to pronounced improvements.

In the one shot problem an implementor has to select a single social alternative, from a set of feasible ones, that will be optimal relative to his own preferences and the unknown preferences of the members of a group of participants. Since the participants' preferences are not known to him, and since they may act strategically in supplying needed information, the implementor has to find a method, or a game, whose strategic equilibrium outcome has the desired optimality. It is well known that in many cases this is an impossible task.¹ This has often led researchers to forego the goal of dominant strategy implementation and turn to Bayesian, Nash, or other types of implementation.²

But, for many economic applications, imposing a one shot selection procedure may be unduly restrictive. The problem of selecting a socially optimal path of economic growth allows a move to dynamic implementation, where economic choices can be made repeatedly and updated as feasible sets and information become available. An implementor can also increase the set of feasible choices in a repeated static problem by offering inter-temporal solutions. For example, in a one shot procedure designed to select a chairman of an economics department, the dean may be restricted to the choice of a micro-economist or a macro- economist. But in a repeated

¹ The following papers, and their references, describe some major studies: Arrow (1963), Barbera and Jackson (1995), Dasgupta, Hammond, and Maskin (1979), Gibbard (1973), Hurwicz (1972), Hurwicz (1990), Ledyard (1977), Moore (1992) and Satterthwaite (1975).

² See Groves and Ledyard (1977), Jackson (1991), Myerson (1985), and Postelwaite and Schmeidler (1987).

selection problem, she can consider solutions that alternate between micro and macro economists in consecutive periods.

Even from a purely technical point of view, the differences between a one shot implementation and a dynamic one can be substantial.³ One major factor lies with the ability to learn people's preferences by observing their past choices. While this is impossible in one shot implementation it is essentially unavoidable in repeated interaction. The learning can be done by the implementor as well as by the players observing each other and the implementor's past choices.

It is, however, not clear a priori whether repetition and learning makes implementation easier or harder. Consider for example the problem of determining socially optimal levels of individual contributions to the problem of cleaning air pollution.⁴ In an attempt to free ride on the contribution of others, agents who wish to be charged a little are likely to exaggerate down the importance of this undertaking. However, if too many agents do so, society may end up with polluted air which does not reflect the true preferences of its members. If the decision on optimal contributions were done repeatedly, it is not clear whether such inefficiency could persist. After some time society's members would learn that the air did not get cleaned, and the free riders would have an informative reason to offer contributions. But now sophisticated rational agents, knowing that others will eventually learn, may have an even stronger incentive to

³ This is to be expected, given the folk theorems of infinitely, see for example Rubinstein (1979), and finitely, see Friedman (1985) and Benoit and Krishna (1985), repeated games. More closely related, however, are well known works on repeated games with incomplete information and reputation, for example Aumann and Maschler (1967), Kreps et al (1982), Neyman (1985), Fudenberg and Maskin (1986), and Fudenberg and Levine (1989). The feature of these papers, that early reputation is built to be enjoyed throughout the duration of the game, takes on a different form when an uninformed, but powerful, implementor is brought into the picture.

⁴ An interesting take on this problem can be found in Rangel(1997) where he looks at public goods in an overlapping generations model.

"teach" their opponents that they do not care with the hope that the opponents will give in first.

As the above example suggests, closely related to the issue of learning is the issue of time preference. It is quite likely that the levels of patience of the agents and the implementor will have a significant influence on the outcome of the interaction. Of particular interest to us is the difference in the patience rates between the social implementor and the economic agents. Our intuition is as follows. (1) Without learning by the planner, all implementation must be done as, or is equivalent to, one shot implementation even in problems that admit intertemporal solutions such as the growth model or sequential candidate choice. (2) If agents also do not learn or if they are extremely impatient, then we are constrained to Bayesian implementation. (3) If agents do learn, the results of Kalai and Lehrer (1993) and others tell us they will eventually play a Nash equilibrium in the true environment. So even if the planner does not learn, if the planner is patient enough then Nash implementation may be possible. (4) Somewhat surprisingly, we will show below that if the planner learns and is more patient than the agents, then we can achieve dominant strategy implementation, making it irrelevant whether the agents learn or not. (5) What we don't yet have a feel for is what happens if all learn but some agents are more patient than the planner, although an infinitely impatient planner is probably restricted to the equivalent of a sequence of one shot implementations.

In this paper we prove a result that illustrates how significant these issues may be. Our purpose in presenting it is not to supply a general model, but rather to illustrate the power of the phenomenon. For that reason, we keep this example as elementary and simple as possible. In particular, we show how patience on the part of the implementor can help her in overcoming the information extraction problem. When we take it to an extreme, where she is only interested in

doing things right in the long run, as compared to the economic agents who perform time discounting, we see that dominant strategy implementation becomes limitless. This is in striking contrast to the Gibbard (1973) and Satterthwaite (1975) theorem that shows that in a one shot problem, dominant strategy implementation can only be used to implement dictatorial choices.

We are able to get this strong implementation result, using a strictly dominant strategies equilibrium, in a simple direct revelation game involving a single type-declaration move at an initial stage. Using the revelation principle, this saves us from having to write a general model with a discussion of a variety of equilibrium options. But less severe assumptions and a search for necessary conditions, as opposed to our sufficient one, will undoubtedly require significantly more modeling involving notions of Nash and/or Bayesian implementation. The revelation principle will still hold for Bayesian implementation and the type of revelation game we use in the next section could continue to be a major tool for the analysis of implementation in more complex environments.

Finally, we have made the analysis quite simple by restricting the number of players' types to be finite and by considering stationary environments. This paper is meant to illustrate a fundamental contrast between repeated and one shot implementation. Attempts by us to obtain similar results under the most general conditions would only cloud the issues by turning this into a learning, rather than an implementation, paper. We note, however, that the inference used in our simple mechanism is a special case of Bayesian learning, and we refer the interested reader to papers of Jordan (1991) and Kalai and Lehrer (1993) for more powerful generalizations.

2. Patient Implementation by Dominant Strategies

The Environment

A , with generic elements a, b , denotes a set of social alternatives.

N , with generic elements i, j , denotes a finite set of economic agents.

Θ_i , with generic elements θ_i , denotes a finite set of player i types.

$\Theta = \times_i \Theta_i$, with generic elements θ , denotes the set of type profiles.

$u_i(\theta_i, a)$ denotes the utility of a player i from the alternative a when he is of type θ_i .

A^∞ is the set of infinite sequences of social alternatives. For each

$$a^\infty \in A^\infty, \text{ we write } a^\infty = (a^1, a^2, \dots, a^t, \dots).$$

δ_i , a real number between 0 and 1, denotes the discount parameter of agent i .

$$u_i(\theta_i, a^\infty) = \sum_{t=1}^{\infty} \delta_i^{t-1} u_i(\theta_i, a^t)$$

Repeated Implementation

An outcome function o assigns to each type profile θ an outcome sequence $o(\theta) \in A^\infty$.

Since, for our purpose, revelation games and dominant strategy equilibria are sufficient, we restrict the notions of implementations to these cases.

A revelation game form g is a function assigning an outcome sequence to each (reported) type profile θ , i.e., for every θ $g(\theta) \in A^\infty$.

The revelation game g dominance-implements the outcome function o if for each type profile θ ,

1. for each agent i , θ_i is a dominant strategy in the game induced by g and by θ , and
2. $g(\theta) = o(\theta)$.

In this paper, we restrict attention to perfectly patient planners. That is, we assume throughout that if there are two outcome functions $g(\theta)$ and $h(\theta)$ and a time T such that $g(\theta)^t = h(\theta)^t$ for all $t > T$ then the planner is indifferent between g and h . Such a planner cares only about eventual implementation. It is sufficient, for example, that the planner's utility function be $V = \lim_{T \rightarrow \infty} (1/T) \sum_{t=1}^T v(a^t)$. With this in mind, we can relax the concept of *dominance-implementation*.

The revelation game g patiently dominance-implements the outcome function o if for each type profile θ ,

1. each θ_i is a dominant strategy as above, and
2. for some time T , $g(\theta)^t = o(\theta)^t$ for all $t > T$.

If we can find a game g that patiently dominance-implements the outcome function o we say that o is patiently dominance implementable.

As we will see below, a patient planner can rely on revealed preference to extract all relevant information from the players without worrying about incentive compatibility. This, however, requires a simple condition that says that the different types of a player are sufficiently different.

We say that types are separable in the environment if for every player i and for every distinct types θ_i and $\bar{\theta}_i$ there is a pair of social alternatives a and b in A with $u_i(\theta_i, a) < u_i(\theta_i, b)$ and $u_i(\bar{\theta}_i, a) > u_i(\bar{\theta}_i, b)$.

An example of violation of separability is when a player i has two types, θ_i and $\bar{\theta}_i$, with type θ_i being completely indifferent between all the social alternatives and type $\bar{\theta}_i$ not being completely indifferent (a more detailed discussion of this condition is offered in the sequel). The power of the assumption that types are separable is given by the following theorem.

Theorem: In a types separable environment with bounded utility functions, every outcome function is patiently dominance implementable.

Proof: Given the outcome function o , we will construct the function g with the following two properties:

(Uniform) eventual coincidence with o . For some time T for all type profiles $\theta \in \Theta$ and all $t > T$, $g(\theta)^t = o(\theta)^t$.

Strict dominance of truthful revelation. For all type profiles $\theta \in \Theta$ and for every player i $u_i(\theta_i, g(\theta_1, \dots, \hat{\theta}_i, \dots, \theta_n))$ is uniquely maximized at $\hat{\theta}_i = \theta_i$.

We will select positive integers D_1, D_2, \dots, D_n and C with the following properties. The first D_1 periods will represent player 1's tenure as a dictator. This means that all social choices during this period, $g(\theta)^t$ for $t = 1, 2, \dots, D_1$, will depend entirely on θ_1 and be independent of θ_j with $j \neq 1$. Similarly, the next D_2 periods will represent player 2's tenure as dictator, and so on. Thus, each player in succession will receive a number of periods when only his preferences dictate the social outcomes in a manner described below.

Following this sequential dictatorship phase of length $D_1 + D_2 + \dots + D_n$, the subsequent C periods represent a cooling down phase. The social choices during these periods will be

independent of θ . For example, a constant alternative a , which is the same for all θ , will be repeated C times.

We let W be the waiting time, denoting the length of the initial phase consisting of all the dictatorships and the cooling down period, $W = D_1 + D_2 + \dots + D_n + C$. Following this initial segment we define the period of eventually "doing the right thing" by $g(\theta)^t = o(\theta)^t$ for $t = W + 1, W + 2, \dots$. Notice that the first desired property, eventual coincidence of g and o , is thus satisfied if θ is revealed correctly.

The selection of social choices during the dictatorship phase of player i is described as follows. Let the length of the dictatorship phase, D_i , be the number of unordered pairs of distinct types of player i and let p^1, p^2, \dots, p^{D_i} be a fixed enumeration of these pairs. For each p^q let (a^q, b^q) be an ordered pair of social alternatives separating the pair of types of p^q , and let $L_i = D_1 + D_2 + \dots + D_{i-1}$ denote the last period prior to player i 's dictatorship. For $q = 1, 2, \dots, D_i$ define $g(\theta)^{L_i+q}$ to be the preferred choice between a^q and b^q , according to θ_i if $u_i(\theta_i, a^q) \neq u_i(\theta_i, b^q)$. Otherwise, let $g(\theta)^{L_i+q} = a^q$.

If player i 's true type is θ_i , then reporting $\bar{\theta}_i \neq \theta_i$ will strictly decrease his total payoff during his dictatorship. During all the periods $L_i + q$ in which he is really indifferent between a^q and b^q he is also indifferent between the social choices resulting from reporting θ_i and reporting $\bar{\theta}_i$. In periods of no indifference he can only lose by switching from θ_i to $\bar{\theta}_i$. And at least in one period, the one corresponding to the pair $\{\theta_i, \bar{\theta}_i\}$, he strictly loses by reporting $\bar{\theta}_i$.

Since player i does not affect the social outcomes during the dictatorships of the other players, player i also strictly maximizes his own utility during the entire sequential dictatorships

phase by reporting his true type.

The purpose of the cooling down period is to keep this strict preference for true revelation for the entire game. By counting on the discounting of the players, we can make the duration of the cooling down time, C , sufficiently long that all the discounted utility received after the cooling down period is meaningless in comparison to the strict preferences established during a person's tenure as a dictator. Since the number of player types is finite, and since θ_i was the unique maximizer of a player's utility during his dictatorship, C can be chosen to be finite. QED

While the theorem and proof above were selected mainly to make the presentation of the message simple, we foresee possible improvements in both the statement of the theorem and the constructed mechanism. For this purpose the following comments may be useful.

3. On the Condition of Type Separability

As usual in Bayesian settings, different types of a player may stand for different preferences, but they may also stand for different information. For example, a player i may be of two types, one type possessing private information that the state of the environment is in a set A , with the other type possessing the information not in A . Assume further that, despite the different information, the two types have identical preferences over the social alternatives. In this case the types are not separable, and our mechanism would fail if we wish to implement different social alternatives for these two types.

Another failure of separability occurs when two types of a player are identical in all their ordinal selections among pairs of socially feasible alternatives, but the intensity of their

preferences is different. Consider, for example, the following allocation problem. In each period there is one unit of a perishable indivisible item to be allocated to either player 1 or player 2. Formally, the social choices in each period are $(1,0)$ or $(0,1)$, indicating the player that receives 1 unit and the player that receives 0 units of the item. Each player may be of two types: one that highly values the item, θ_i , and another type, $\bar{\theta}_i$, that has a low, even if positive, value for the item. Formally, let $u_1(\theta_1, (1,0)) = 2$, $u_1(\bar{\theta}_1, (1,0)) = 1$, and $u_1(\theta_1, (0,1)) = u_1(\bar{\theta}_1, (0,1)) = 0$, with similar types θ_2 and $\bar{\theta}_2$ defined for player 2.

The implementor desires to implement the following outcome function. If one of the players is of the high type and the other is of a low type, the item should be given repeatedly to the high-type player. If both players are of the same type, the item should be given to player 1 on even periods and to player 2 on odd periods.

It is easy to see that this outcome function cannot be implemented and that the type-separation condition fails. If the implementor could introduce another social alternative to create separation it would work. For example, if he could offer some temporary private good whose utility to the low types was greater than 1 and to the high types was smaller than 2, he could separate their initial types by giving them dictatorial choices involving the separating temporary good early on, and later doing the "right thing."

4. Simpler Randomizing Mechanisms

If the implementor can randomize in making choices, simpler mechanisms work. The following random-dictator revelation mechanism is an example.

In an initial period, each player i announces a type θ_i . Then, in the first step, the

implementor randomly chooses a player i and one of this player's separating pairs of alternatives (from the earlier constructions) $p^q = (a^q, b^q)$. The selected social alternative is a^q or b^q according to the strict preference of the reported θ_i (make it a^q if θ_i is indifferent between the two). To obtain the strict dominance result it is enough now to generate a long cooling off phase as before, and follow it for all periods t afterwards by selecting $o(\theta)^t$.

A second and, perhaps, better alternative for the cooling off phase is to use the same alternative a^q or b^q , chosen for the first period, and to make it the repeated social selection for many periods. This strengthens the incentives of impatient players to reveal correctly, resulting in a shorter cooling off phase.

5. On the Patience of the Implementor

First conceptually, the assumption of patient implementor may be more suitable to implementation problems within organizations. For example the dean of a school may be interested in the long run performance (or the evaluation by others of how he contributed to the long run performance) of the school, while the faculty members may be taking a shorter time horizon in their individual optimization problems. But in social implementation, for example in constitutional design, the assumption of patient implementor with impatient participants is more questionable and should be a subject for further study. For example assuming time consistency on the part of the implementor may lead to behavior compatible with a higher level of patience.

From a technical viewpoint, our construction was very wasteful and in specific applications one may assume much less than perfect patience, depending on the nature of the outcome function and the extend of diversity among types. For example, in the extreme case that

the participants are completely myopic the random-dictator mechanism described above strictly-dominance implements any outcome function from period 2 on, without the need for a cooling off phase.

REFERENCES

- Aumann, R. and M. Maschler. 1967. "Repeated games with incomplete information: A survey of recent results," in *Report to the U.S. Arms Control and Disarmament Agency* ST-116, Washington, D.C., III: 287-403.
- Arrow, K. J. 1963. *Social Choice and Individual Values*, 2nd ed. New York: Wiley.
- Barbera, Salvador and Matthew O. Jackson. 1995. "Strategy-Proof Exchange," *Econometrica*, Vol. 63, No. 1, 51-87.
- Benoit, J.-P. and V. Krishna. 1985. "Finitely repeated games," *Econometrica* 53: 905-922.
- Dasgupta, P., P. Hammond, and E. Maskin. 1979. "The Implementation of Social Choice Rules: Some Results on Incentive Compatibility." *Review of Economic Studies* 46:185-216.
- Friedman, J. W. 1985. "Cooperative equilibria in finite horizon non-cooperative supergames," *Journal of Economic Theory* 35: 390-398.
- Fudenberg, D. and E. Maskin. 1986. "The folk theorem in repeated games with discounting and with incomplete information," *Econometrica* 54: 533-554.
- Fudenberg, D. and D. Levine. 1989. "Reputation and equilibrium selection in games with a patient player," *Econometrica* 57: 759-779.
- Gibbard, A. 1973. "Manipulation of voting schemes: A general result", *Econometrica* 41: 587-601.
- Groves, T. and J. Ledyard. 1977. "Optimal Allocation of Public Goods: A Solution to the 'Free Rider' Problem", *Econometrica*, 45, 783-809.
- Hurwicz, L. 1972. "On informationally decentralized systems," in C.B. McGuire and R. Radner (eds.), *Decision and Organization*. Amsterdam: North-Holland, pp.297-336.
- Hurwicz, L., AND M. Walker 1990. "On the Generic Nonoptimality of Dominant-Strategy Allocation Mechanisms: A General Theorem that Includes Pure Exchange Economies," *Econometrica*, 58, 683-704.
- Jackson, M. 1991. "Bayesian Implementation", *Econometrica*, 59, 461-477.
- Jordan, J. S. 1991. "Bayesian Learning in Normal Form Games," *Games and Economic Behavior*, 3, 60-81.

Kalai, Ehud and E. Lehrer. 1993. "Rational Learning Leads to Nash Equilibrium," *Econometrica*, Vol. 61, No. 5 1019-1045.

Kreps, D, P. Milgrom, J. Roberts and R. Wilson. 1982. "Rational cooperation in finitely repeated prisoner's dilemma," *Journal of Economic Theory* 27: 245-252.

Ledyard, John O. 1977. "Incentive Compatible Behavior in Core-Selecting Organizations", *Econometrica*, Vol. 45, No. 7 1607-1621.

Moore, J. 1992. "Implementations in Environments with Complete Information," in *Advances in Economic Theory: The Proceedings of the Congress of the Econometric Society*, ed. By J.-J. Laffont. Cambridge: Cambridge University Press.

Myerson, R. B. 1985. "Bayesian Equilibrium and Incentive Compatibility." In L. Hurwicz, D. Schmeidler, and H. Sonnenschein, eds., *Social Goals and Social Organization*, pp. 229-259. Cambridge: Cambridge University Press.

Neyman, A. 1985. "Bounded complexity justifies cooperation in the finitely repeated prisoner's dilemma," *Economic Letters* 19: 227-229.

Postlewaite, A. and D. Schmeidler. 1987. "Differential Information and Strategic Behavior in Economic Environments: A General Equilibrium Approach". In T. Groves, R. Radner and S. Reiter (Eds.), *Information, Incentives and Economic Mechanisms - Essays in Honor of Leonid Hurwicz*. University of Minnesota Press.

Rangel, Antonio. 1997. "Intergenerational Public Goods". Manuscript.

Rubinstein, A. 1979. "Equilibrium in supergames with overtaking criterion," *Journal of Economic Theory* 21: 1-9.

Satterthwaite, M. 1975. "Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions", *Journal of Economic Theory* 10: 187-217.

Discussion Paper No. 1205R

REPEATED IMPLEMENTATION⁺

by

Ehud Kalai^{*}

and

John O. Ledyard^{}**

August 1996

rev. March 1998

JEL Classifications:

**C7 - Game Theory and Bargaining Theory
D7 - Analysis of Collective Decision Making**

^{*}Department of Managerial Economics and Decision Sciences, J. L. Kellogg Graduate School of Management, Northwestern University, 2001 Sheridan Road, Evanston, Illinois 60208. Tel: (847) 491-7017, E-mail: kalai@nwu.edu

^{**}Division of Humanities and Social Sciences, California Institute of Technology, 1200 California Blvd., Pasadena, CA 91125. Tel: (818) 395-8482, Fax: (818) 405-9841, E-mail: jledyard@caltech.edu

^{*}The authors wish to thank Andreas Blume, Tim Feddersen, Alvaro Sandroni, and a referee of this journal for helpful suggestions. Kalai's research is partly supported by NSF Economics Grant No. SBR-955421. Ledyard's research is partly supported by the New Millenium Project Office of the Jet Propulsion Laboratory of NASA.