Pigouvian Pricing and Stochastic Evolutionary Implementation

William H. Sandholm[*]
Department of Economics
University of Wisconsin
1180 Observatory Drive
Madison, WI 53706
whs@ssc.wisc.edu
http://www.ssc.wisc.edu/~whs

September 23, 2005

[*] I thank Bart Lipman, Nolan Miller, and Larry Samuelson for helpful discussions, and two anonymous referees and an anonymous Associate Editor for exceptionally helpful reports. Financial support from NSF Grant SES-0092145 is gratefully acknowledged.

Abstract

We study the implementation of efficient behavior in settings with externalities. A planner would like to ensure that a group of agents make socially optimal choices, but he only has limited information about the agents' preferences, and can only distinguish individual agents through the actions they choose. We describe the agents' behavior using a stochastic evolutionary model, assuming that their choice probabilities are given by the logit choice rule. We prove that there is a simple price scheme with the following property: regardless of the realization of preferences, a group of agents subjected to the price scheme will spend the vast majority of time in the long run behaving efficiently. The price scheme defines a game that may possess multiple equilibria, but we are able to obtain a unique and efficient selection from this set because of the stochastic nature of the agents' choice rule. We conclude by comparing the performance of our price scheme with that of VCG mechanisms. *Journal of Economic Literature* classification numbers: C72, C73, E32, D62, D82.

1. Introduction

Externalities are a basic source of economic inefficiency. Since the work of Pigou [2], it has been understood that this inefficiency can be mitigated if agents are charged prices that reflect the externalities they impose upon others. More precisely, if one charges agents for the externalities they create at the efficient state, one can ensure that this state constitutes an equilibrium.

While this approach to dealing with externalities is quite powerful, it has certain limitations. First, it requires that the social planner know the agents' preferences, so that he can determine the efficient state: without this knowledge, he does not know what prices to levy. Second, even if preferences are known, there may still be a problem of multiple equilibria: while the Pigouvian prices render efficient behavior an equilibrium, they often create other equilibria as well. This is almost certain to be true when externalities are positive, as is the case, for example, in contexts involving consumer technology choice (Katz and Shapiro [20]) or macroeconomic spillovers (Cooper [8]).

In this paper, we view externality pricing as an implementation problem. Each member of a population of agents must choose from the same set of actions. The utility function of each agent is composed of two terms. One term is common across agents, and captures the externalities they impose upon one another. The other term varies from agent to agent, and only depends on the agent's own choice.

A social planner would like to ensure efficient behavior, but he faces two constraints. First, he has no knowledge of the agents' idiosyncratic payoffs, and hence does not know what constitutes efficient behavior. Second, he is unable to distinguish individual agents except through their action choices, and so can only influence behavior using simple pricing schemes. This anonymity assumption can be viewed as a requirement that the mechanism be easy to administer even when the number of agents is large.

Rather than assume equilibrium play, we model the agents' behavior using an evolutionary approach. As time passes, each agent occasionally considers switching actions. At these moments, the agent's action choice is determined by the logit choice rule. Under this rule, an agent usually selects his current optimal action, but sometimes chooses other actions, with actions yielding lower payoffs being less likely choices.[1]

[1] The logit choice rule has been used to model behavior in a variety of economic contexts—see Anderson, de Palma, and Thisse [1], Durlauf [7], and the references therein.

Given this specification of the agents' behavior, we can describe the planner's problem in the following way. The planner would like to choose a price scheme with the property that regardless of the realization of their preferences or their initial behavior, the agents spend the vast majority of time in the long run behaving efficiently.

Our main result shows that the planner can achieve this goal by introducing a simple pricing scheme. Under this scheme, players are always made to experience the externalities that they *currently* create. Because the scheme focuses on current externalities and not those created at the efficient state, a planner can execute this scheme despite the fact that he is unable to determine the efficient state.

This variable price scheme always ensures that the efficient state is an equilibrium, though possibly only one among many. Were the agents' adjustment processes deterministic, their eventual behavior under the price scheme would generally depend on their initial behavior. Nevertheless, we show that if there are small probabilities of errant behavior, only the efficient strategy profile is played in a non-negligible fraction of periods after a long enough history of play.

The proof of this result is surprisingly simple. It relies on the notion of a potential game, introduced by Monderer and Shapley [21].[2] A potential game is a game that admits a potential function: a function on the space of strategy profiles that better-reply adjustment processes must ascend. Intuitively, players' incentives in these games are aligned in such a way that it is as if every player were attempting to maximize a common payoff function.

We show that our price scheme has the following critical property: regardless of the realization of types, the game that is created when the price scheme is imposed is a potential game; its potential function is precisely the realized aggregate payoff function. This ensures that the efficient strategy profiles (those which maximize aggregate payoffs) are Nash equilibria, although other strategy profiles may be Nash equilibria as well.

At this stage, the stochastic specification of behavior plays its role. Since the work of Foster and Young [13], Kandori, Mailath, and Rob [18] and Young [31], it has been known that introducing small probabilities of mistakes to an evolutionary process can generate a unique selection among multiple strict equilibria. To establish our implementation result, we appeal to a theorem of Blume [6], who shows that if players in a potential game make decisions according to the logit choice rule, they spend the

---

2 Early examples of potential games were used to model congestion (Beckmann, McGuire, and Winsten [3], Rosenthal [24]) and in the study of population genetics (Hofbauer and Sigmund [16]). For more recent results on potential games and further references, see Sandholm [25, 28].

majority of time in the long run playing the global maximizers of the potential function. Since our price scheme guarantees that potential always corresponds to aggregate payoffs, Blume's theorem enables us to establish the long run efficiency of play.

In other papers (Sandholm [26, 27]), we consider the use of price schemes to solve implementation problems in settings with negative externalities. These papers consider continuous population, deterministic models (unlike the finite population, stochastic model considered here), and provide conditions under which efficient behavior can be rendered globally stable under deterministic adjustment processes. The key condition needed to prove these results is that the total payoffs from externalities form a concave function of the population's aggregate behavior. This condition ensures that the equilibria generated by the price scheme are always unique, globally stable, and globally efficient. However, this concavity condition is particular to models of negative externalities,[3] and when it fails to hold, price schemes typically generate multiple equilibria. In this paper, we consider a general model which allows for negative externalities, positive externalities, or combinations of the two. We therefore dispense with concavity conditions, and instead rely on a stochastic specification of individual behavior. The stochastic model of choice enables us to select among the multiple equilibria that the price scheme generates, allowing us to render efficient behavior the unique long run prediction of play.

The next section describes our model of externalities and the behavior revision process. Section 3 presents the formal definition of the planner's problem and the statement of our main result. Section 4 provides its proof. A detailed comparison of the optimal price scheme and VCG mechanisms is offered in Section 5. Section 6 concludes with a discussion of the limitations of our results.

---

[3] As a simple illustration, consider a continuum of agents who choose from a set of n activities. Because of externalities, the payoff $\pi_a(z_a)$ to activity a is a function of the proportion z_a who choose it. Total payoffs are given by $\Pi(z) = \Sigma_a\, z_a\, \pi_a(z_a)$. In a model of negative externalities (e.g., traffic congestion), one might assume that the functions π_a are decreasing and concave; since $\Pi_{aa}(z) = z_a \pi_a''(z_a) + 2\pi_a'(z_a) \leq 0$ and $\Pi_{ab}(z) = 0$, it follows that Π is concave. In a model of positive externalities, one might assume instead that the functions π_a are increasing and convex (or, more generally, that they satisfy $z_a \pi_a''(z_a) + 2\pi_a'(z_a) \geq 0$), in which case Π is actually convex.

2. The Model

2.1 Basic Definitions

We study implementation problems based on N player normal form games. We let S^i denote the strategy set for player $i \in I = \{1, \ldots, N\}$, and let $S = \prod_i S^i$ denote the set of strategy profiles. Player i's utility function is given by $U^i : S \to \mathbf{R}$.

In our model, all players choose strategies from a common set: $S^i = A = \{0, 1, 2, \ldots, n\}$ for all players i. Typical strategies are denoted s^i, a, and b. Strategy 0 always represents an outside option; the remaining strategies are *active strategies*. For each strategy profile $s \in S$, we let $x(s)$ describe the number of players choosing each active strategy $a \in \{1, \ldots, n\}$; hence, $x_a(s) = \#\{i : s^i = a\}$. The vector $x(s)$ is an element of the set $X^N = \{x \in \mathbf{Z}^n_+ : \sum_a x_a \leq N\}$.

We consider games in which players' utility functions are the sum of two components. The first component, the *common payoff*, depends on a player's choice and on the population's aggregate behavior, and is the same for all players. This payoff captures the externalities that each individual imposes upon the others through his strategy choice. Common payoffs are described by a function u: $X^N \to \mathbf{R}^A$; $u_a(x)$ is the payoff to strategy a when the strategy distribution is x.[4] Observe that this model allows for cross-effects between strategies: the payoff to strategy a can depend not only on the number of players choosing strategy a, but also on the full distribution of players over all strategies in A. The common payoff to the outside option is always zero: $u_0(x) \equiv 0$. Apart from the definition of the outside option, we place no restrictions on the common payoff function.[5]

The second component of each player's utility function, the *idiosyncratic payoff*, only depends on the player's own choice, but varies from player to player. Player i's idiosyncratic payoffs are described by a vector $\theta^i \in \Theta^i = \mathbf{R}^A$; θ^i_a represents player i's bias toward or against strategy a. Summing the two components, we see that player i's utility function is given by

$$U^i(s^i, s^{-i}) = u_{s^i}(x(s)) + \theta^i_{s^i}.$$

---

[4] The value of $u_a(x)$ can be specified arbitrarily when $x_a = 0$.

[5] When we introduce our price schemes below, we will not allow the planner to price the outside option. Therefore, including the outside option is a way of introducing participation constraints to the model. Of course, all of our implementation results continue to hold if the outside option (and hence the participation constraints) are removed.

When we want to emphasize the dependence of player i's payoffs on his type realization, we write $U^i(s, \theta^i)$ for $U^i(s)$. We refer to a game in the class described here as a pair (u, θ), where u is a common payoff function and $\theta \in \Theta = \prod_i \Theta^i$ is a type profile.

As an application of this approach, consider the following model of consumer technology choice. Each consumer i is able to choose from a set of n technological standards (e.g., computer platforms, word processing formats), or can opt not to use the technology at all. If a consumer chooses a particular standard, he is able to interact (e.g., share files) with all others who have chosen the same standard. Because of these network externalities, the benefit of using each standard is increasing in the number of consumers who do so. Moreover, there maybe some degree of cross-compatibility between standards, so that consumers choosing standard a may gain if others choose standard b. These benefits are understood by the planner, and are captured by the common payoff function u.

Network externalities are not the only source of benefits from choosing a particular technological standard. Consumers may prefer one specific standard because it is particularly well suited to the tasks to which they will apply it, and the benefits derived from this fit between product and task may even dominate those obtained through network effects. Each consumer's idiosyncratic benefits are unknown to the planner, and are captured in the consumer's type θ^i.

A social planner would like to ensure that the consumers behave efficiently. Typically, efficient behavior will entail most players choosing a single standard, but will involve some consumers whose idiosyncratic payoffs are especially large choosing another standard or the outside option. However, *every* strategy profile is the most efficient one for *some* realization of types. Can the planner find a simple pricing mechanism which ensures the long run selection of the efficient strategy profile, regardless of the (unobserved) realizations of the consumer's types?

2.2 Evolutionary Dynamics

We address this question using a dynamic evolutionary approach, explicitly specifying the process through which players adjust their behavior in response to current payoffs.[6] We suppose that at the onset of play, each player draws a type θ^i, which is fixed throughout the course of play. At each discrete moment in time, one

6 Kandori, Mailath, and Rob [18] and Kandori and Rob [19] use evolutionary models to study consumer technology choice.

player is chosen at random and given an opportunity to change strategies. This evolutionary process is described by a discrete time Markov chain, $\{s_t\}_{t\in Z_+}$, which takes values in the space S of strategy profiles.[7]

To define this Markov chain, we must specify how players take advantage of their revision opportunities. One natural way to do this is to suppose that a player who receives such an opportunity switches to his current best response to the others' behavior. However, while keeping the probability of optimal responses high, we would like to allow the possibility that players sometimes choose suboptimally.

We accomplish this by modeling choice using the *logit choice function* L^ε. The logit choice function is parameterized by a scalar $\varepsilon \in (0, \infty)$, representing the level of noise in the players' decisions. Under this choice rule, a player facing a choice among $n + 1$ alternatives yielding payoffs of $\pi_0, \pi_1, \ldots, \pi_n$ chooses alternative a with probability

$$L_a^\varepsilon(\pi) = \frac{\exp(\varepsilon^{-1}\pi_a)}{\sum_b \exp(\varepsilon^{-1}\pi_b)}.$$

Alternatives yielding higher payoffs are chosen with higher probability, and if ε is very small, nearly all probability mass is placed on the best alternatives.[8] As we noted in the introduction, the logit choice function has been applied in a number of economic contexts, and in experimental work predictions based on the logit model have compared favorably with those based on pure optimization.

In each period of our evolutionary model, one of the N players is chosen at random to revise his choice according to the logit rule. This procedure generates the following transition probabilities for the Markov chain $\{s_t\}$:

$$P^\varepsilon\left(s_{t+1} = \hat{s} \,\middle|\, s_t = s\right) = \begin{cases} \dfrac{1}{N}\dfrac{\exp(\varepsilon^{-1}U^i(\hat{s}^i, s^{-i}))}{\sum_b \exp(\varepsilon^{-1}U^i(b, s^{-i}))} & \text{if } \hat{s}^i \neq s^i \text{ and } \hat{s}^{-i} = s^{-i} \text{ for some } i; \\ \displaystyle\sum_i \left(\frac{1}{N}\frac{\exp(\varepsilon^{-1}U^i(s))}{\sum_b \exp(\varepsilon^{-1}U^i(b, s^{-i}))}\right) & \text{if } \hat{s} = s; \\ 0 & \text{otherwise.} \end{cases}$$

7 Our results are easily extended to a continuous time model in which players' revision opportunities arrive via independent Poisson processes.

8 For the foundations of the logit choice model, see Anderson, de Palma, and Thisse [1].

The Markov chain $\{s_t\}$ defined by these transition probabilities is irreducible (i.e., any state can be reached from any other) and aperiodic.[9] It therefore admits a unique stationary distribution μ^ε. This distribution describes the long run behavior of the process in two distinct ways. First, it is the *long run distribution* of the process: μ^ε approximates the distribution of the random variable s_t for all large enough times t. More importantly, it is the *ergodic distribution* of the process: with probability one, μ^ε describes the long run time average of play. Both of these properties hold regardless of the initial state s_0. Thus, the measure μ^ε is an appropriate description of the population's long run behavior.[10]

We are especially interested in the long run behavior of the Markov chain when the noise level ε is small. To capture this, we define the *logit stochastically stable* strategy profiles to be those which retain positive weight under the stationary distribution μ^ε when ε is taken to zero. Formally, we let

$$LSS = \left\{ s \in S : \lim_{\varepsilon \downarrow 0} \mu^\varepsilon(s) > 0 \right\}.$$

When we want to emphasize the dependence of this set on the underlying game, we write $LSS(u, \theta)$ in place of LSS.

3. Stochastic Implementation

We now consider a social planner whose goal is to ensure efficient behavior. He faces two constraints. First, the planner has no information about the players' types θ^i, and hence does not know which strategy profiles are efficient. Second, the planner regards the players as *anonymous*. He can reward or punish players for choosing particular strategies, and is able to condition these rewards and punishments on the players' aggregate behavior, but he cannot condition them on the players' identities (or, of course, on the realization of types).

We therefore restrict the planner to a class of mechanisms we call *price schemes*. A price scheme is defined by a function p: $X^N \to \mathbf{R}^A$. The scalar $p_a(x)$ represents the payment that a player choosing strategy a must make to the planner when aggregate behavior is x.[11] We assume that the planner can neither price nor subsidize the outside

9 Irreducibility follows from the fact that the process can move from any strategy profile to any other within N periods. Aperiodicity is implied by the fact that there is always a positive probability that the strategy profile does not change.

10 For further background on Markov chains, see, e.g., Durrett [10].

11 Like $u_a(x)$, $p_a(x)$ can be specified arbitrarily when $x_a = 0$.

option: $p_0(x) \equiv 0$. Price schemes respect the planner's information constraint because they do not depend on the type profile θ. They respect anonymity because the prices associated with each action are the same for all players, and because these prices are only functions of aggregate behavior x.

If the price scheme p is imposed, the common payoff which players obtain from playing strategy a changes from $u_a(x)$ to $u_a(x) - p_a(x)$. Therefore, the game the players face if this scheme is executed is described by the pair $(u - p, \theta)$.

To define the planner's problem, we introduce the notion of a *social choice correspondence*. A social choice correspondence is a map, $\phi\colon \Theta \Rightarrow S$, which associates with each type profile θ a set of strategy profiles $\phi(\theta)$; these are the planner's preferred behaviors under type profile θ. We say that the price scheme p *stochastically implements* the social choice correspondence ϕ if for each type profile $\theta \in \Theta$, we have that

$$LSS(u - p, \theta) = \phi(\theta).$$

That is, p stochastically implements ϕ if regardless of the realization of types, the strategy profiles which are played often in the long run are precisely those which the planner prefers under this realization of types.

We are concerned with the implementation of efficient behavior.[12] Regardless of the type realization θ, the planner would like to ensure that players choose the strategy profiles which maximize aggregate utility. Define the *aggregate utility function* $\bar{U} : S \times \Theta \to \mathbf{R}$ by

$$\bar{U}(s,\theta) = \sum_a x_a(s) u_a(x(s)) + \sum_a \sum_{i:\, s^i = a} \theta^i_a .$$

The first term in the definition of $\bar{U}(s,\theta)$ captures the total common payoffs under strategy profile s, while the second term aggregates all idiosyncratic payoffs. The *efficient social choice correspondence*, $\phi^*\colon \Theta \Rightarrow S$, is then defined by

$$\phi^*(\theta) = \operatorname*{arg\,max}_{s \in S} \bar{U}(s,\theta).$$

Let e^a, $a \in \{1, \ldots, n\}$, denote the standard basis vectors in $\mathbf{R}^n$. Then $x - e^a$ is the state obtained from state x by excluding a player choosing strategy a. Our main result shows

12 Our analysis can be extended to allow the social planner to have alternate objectives.

that the efficient social choice correspondence can be stochastically implemented using the price scheme p^*, defined by

$$p_a^*(x) = -\left(\sum_b x_b (u_b(x) - u_b(x - e^a)) - (u_a(x) - u_a(x - e^a)) \right) \text{ for } a \in \{1, \dots, n\};$$
$$p_0^*(x) \equiv 0.$$

Theorem 1: *The price scheme p^* stochastically implements the efficient social choice correspondence ϕ^*. That is, p^* ensures that after any realization of types, aggregate behavior will be efficient during most periods of almost every long enough history of play.*

The price scheme p^* is a form of marginal externality pricing: the price player i pays for selecting strategy a is equal to the current impact of his choice on the other players' payoffs. More precisely, $p_a^*(x)$ is equal in magnitude to the net benefit accruing to i's opponents when i chooses strategy a rather than the outside option. However, unlike under standard Pigouvian pricing, players always pay for the externalities they *currently* create: when behavior is not efficient, the prices they pay correspond to those that they currently impose, not the ones they would impose at the efficient state. Indeed, the planner cannot charge the standard Pigouvian prices, as he does not know which state is efficient, and hence what these prices are. We shall see that by allowing the prices to vary as behavior changes, the planner is able to ensure that the efficient state is always an equilibrium, regardless of what this state turns out to be. Moreover, if players occasionally make suboptimal choices, then in the long run, the efficient state is the unique state played with non-negligible frequency after a long enough history of play.

4. Analysis

Our analysis relies on the notion of a potential game (Monderer and Shapley [21]). Formally, the normal form game U is a *potential game* if it admits a *potential function*: a function Π: $S \to \mathbf{R}$ with the property that

$$U^i(\hat{s}^i, s^{-i}) - U^i(s^i, s^{-i}) = \Pi(\hat{s}^i, s^{-i}) - \Pi(s^i, s^{-i})$$

for all players i, all strategies s^i and $\hat{s}^i$, and all opponents' strategy profiles s^{-i}. In words: whenever a player unilaterally deviates, the change in his payoffs is exactly matched by the change in potential. It follows immediately from this definition that the Nash equilibria of a potential game are precisely the local maximizers of potential on the set S, and that better-reply adjustment processes must converge to Nash equilibria.

The first step in our analysis shows that regardless of the realization of types θ, the price scheme p^* always creates a potential game whose potential function is the realized aggregate payoff function $\bar{U}(\cdot, \theta)$.

Proposition 2: *Fix the common payoffs u and a type profile θ. Then the game $(u - p^*, \theta)$ is a potential game whose potential function is $\bar{U}(\cdot, \theta)$.*

Proof: Fix a player i and two strategy profiles s and $\hat{s}$ with $s^i = a$, $\hat{s}^i = b$, and $s^{-i} = \hat{s}^{-i}$. Let $x^0 = x(s^{-i})$ be the distribution of the strategies of players besides i under these profiles, and let $x^a = x^0 + e^a$ and $x^b = x^0 + e^b$ be the complete strategy distributions under s and $\hat{s}$, respectively. (To allow player i to choose the outside option, we let e^0 denote the vector of zeros.)

For each action $d \in A$, common payoffs under the price scheme p^* are given by

$$u_d(x) - p_d^*(x) = \sum_c x_c(u_c(x) - u_c(x - e^d)) + u_d(x - e^d).$$

We can therefore compute that

$$\begin{aligned}
U^i(\hat{s},\theta^i) - U^i(s,\theta^i) &= \left(\sum_c x_c^b(u_c(x^b) - u_c(x^0)) + u_b(x^0) + \theta_b^i\right) \\
&\quad - \left(\sum_c x_c^a(u_c(x^a) - u_c(x^0)) + u_a(x^0) + \theta_a^i\right) \\
&= \sum_c x_c^0(u_c(x^b) - u_c(x^a)) + u_b(x^b) - u_a(x^a) + \theta_b^i - \theta_a^i \\
&= \left(\sum_c x_c^b u_c(x^b) - \sum_c \sum_{i:\hat{s}^i = c} \theta_c^i\right) - \left(\sum_c x_c^a u_c(x^a) - \sum_c \sum_{i:s^i = c} \theta_c^i\right) \\
&= \bar{U}(\hat{s},\theta) - \bar{U}(s,\theta). \ \blacksquare
\end{aligned}$$

Proposition 2 implies that regardless of the realization of types, any efficient strategy profile is a Nash equilibrium of the game induced by the price scheme p^*. But in general, there will be other Nash equilibria which correspond to "locally efficient" strategy profiles. In the technology choice example, these locally efficient profiles involve most players coordinating on a technology which is suboptimal given the players' preferences. If no errors are made during the strategy adjustment process, all of these equilibria are possible limit behaviors for some set of initial conditions.

Since the work of Foster and Young [13], Kandori, Mailath, and Rob [18], and Young [31], it has been understood that by introducing rare errors to a strategy adjustment process, one can obtain unique predictions about long run play. In the current context, allowing infrequent errors makes it possible in principle to establish the long run selection of the efficient state. However, the planner must ensure that regardless of the realization of types, the equilibrium which is selected is the one which is efficient given that realization of types.

To establish that the price scheme p^* accomplishes this goal, we rely on a result of Blume [6], who shows that in potential games, long run behavior under the logit choice rule can be described in a simple way.

Proposition 3 (Blume [6]): *Suppose that U is a potential game with potential function Π, and that behavior evolves under the logit choice rule with noise level ε. Then the stationary distribution of the process $\{s_t\}$ is given by*

$$\mu^\varepsilon(s) = \frac{\exp(\varepsilon^{-1}\Pi(s))}{\sum_{\hat{s}\in S}\exp(\varepsilon^{-1}\Pi(\hat{s}))}.$$

Consequently, $LSS(U) = \operatorname{argmax}_{s\in S} \Pi(s)$.

Intuitively, the potential function Π defines a "landscape" on the set of strategy profiles; the "peaks" of this landscape are locally stable Nash equilibria. Changes to better performing strategies lead up this landscape to Nash equilibria; errant choices lead down the landscape, causing occasional transitions between equilibria. Roughly speaking, Blume's theorem says that the highest peak of the potential function is more difficult to descend than any other, so much so that when choice errors are rare, the population is nearly always atop this peak. One proves this result by verifying that μ^ε is a reversible distribution for the Markov chain $\{s_t\}$ (i.e., that $\mu^\varepsilon(s)P^\varepsilon(s_{t+1}=\hat{s}|s_t=s)$ = $\mu^\varepsilon(\hat{s})P^\varepsilon(s_{t+1}=s|s_t=\hat{s})$ for all s, $\hat{s} \in S$), which immediately implies that μ^ε is the stationary distribution for $\{s_t\}$ (i.e., that $\sum_s \mu^\varepsilon(s)P^\varepsilon(s_{t+1}=\hat{s}|s_t=s) = \mu^\varepsilon(\hat{s})$ for all $s \in S$).

The proof of Theorem 1 follows easily from these two propositions. Fix the common payoff function u, and suppose that the realized type profile is θ. If the planner imposes the price scheme p^*, then Proposition 2 tells us that the realized game $(u - p^*, \theta)$ is a potential game with potential function $\bar{U}(\cdot,\theta)$. Hence, Proposition 3 implies that $LSS(u - p^*, \theta) = \text{argmax}_{s \in S}\, \bar{U}(s,\theta) = \phi^*(\theta)$, proving the theorem.[13]

5. Comparison with VCG Mechanisms

The implementation problem studied in this paper involves anonymous agents who possess private information. The price schemes we introduce to solve this problem are simple mechanisms that respect the agents' anonymity; the optimal price scheme p^* ensures that in the long run, agents nearly always choose the action profile that is socially optimal given their unobserved types.

Unlike price schemes, most mechanisms for eliciting hidden information are based on direct revelation. Revelation mechanisms require agents to send reports about their types to the planner; the planner then specifies a "social alternative" as a function of the profile of reports he receives. In the externality problem we study here, social alternatives take a special form: they are action profiles in some underlying game. Consequently, a planner using a revelation mechanism must not only elicit truthful reports; he must also ensure that each agent undertakes the action to which he is assigned. When agents are anonymous, this problem is not trivial: in Section 5.1, we show that together, the requirements of anonymity and obedience preclude dominant strategy implementation.

If non-anonymous mechanisms are allowed, obedience can be ensured using forcing contracts, and the implementation problem that remains is dominant strategy solvable using Vickrey-Clarke-Groves mechanisms. Both these mechanisms and our optimal price scheme can be described as instruments that make agents pay for the externalities

[13] It is well known that in 2 x 2 symmetric coordination games, the stochastic evolutionary models of Kandori, Mailath, and Rob [18] and Young [31] select the risk dominant equilibrium, which in general does not maximize the sum of the players' payoffs. While these selection theorems may appear to be at odds with our implementation theorem, there is actually no tension between the results. We apply Blume's [6] theorem to the game $(u - p^*, \theta)$, whose payoffs include the transfers p^*. As the results in [18] and [31] suggest, the equilibrium selected in the game $(u - p^*, \theta)$ is in general not the state that maximizes the sum of the players' payoffs *when the transfer payments are included in this sum*. However, when we speak of the efficiency of a state in our implementation model, we want to ignore transfers between the agents and the planner. We accomplish this by defining efficiency in terms of the payoffs of the original game (u, θ). To prove the implementation theorem, we show that the state that maximizes the sum of the players' payoffs in the original game (u, θ) is the equilibrium selected in the new game $(u - p^*, \theta)$. Since efficiency and equilibrium selection are defined with respect to different games, there is no conflict with the aforementioned results.

they create. Nevertheless, we show in Section 5.2 that the transfers required by the two sorts of mechanisms are not the same. We attribute these differences in transfers to differences in the choice sets offered to the agents under these mechanisms.

In Section 5.3, we argue that the choice of mechanism for an externality problem should depend on the details of the environment of application. In particular, in environments with many agents, executing a VCG mechanism can be costly: doing so requires the planner to collect type reports and to compute and announce an efficient assignment of agents to actions; when obedience is an issue, the planner must also monitor each agent's behavior and punish those who disobey. Since price schemes rely on indirect means of information elicitation and behavior control, they may be preferable to revelation-based mechanisms when the number of agents is large and the interaction recurs over time.

5.1 Anonymity and Obedience

It is implicit in the usual definition of a revelation mechanism that the planner has direct control over the social alternative. In the present context, where the social alternatives are action profiles in an underlying game, this assumption may not always be warranted: even when the planner has enough information to determine the optimal social alternative, he must still take measures to ensure that the players enact it.

If we assume away the problem of obedience, the hidden information problem that remains can be solved in dominant strategies using a VCG mechanism (see Section 5.2 below). Otherwise, the planner faces a dilemma: to ensure obedience, he must either use a mechanism that conditions on agents' names, or he must abandon the goal of implementation in dominant strategies.

Example 5.1: Suppose there are 10 players, each of whom chooses an action from S^i = {0, 1, 2}. Actions 1 and 2 are subject to negative externalities: payoffs are described by $u_1(x) = 10 - x_1$, $u_2(x) = 10 - x_2$, and $u_0 \equiv 0$. Finally, all players are of the same type, namely $\theta^i = (\theta^i_0, \theta^i_1, \theta^i_2) = (0, 0, 0)$.

Suppose that all of this information, including the information about the players' types, is known to the planner. Then there is no hidden information problem: to ensure efficiency, the planner simply needs to get five of the players to choose action 1 and five to choose action 2.

If the planner must treat the players anonymously, he can do nothing to provide different players with different incentives. The most he can do is recommend a set of efficient social alternatives: in this case, those in which five of the players choose each action. Every such action profile is a strict Nash equilibrium; therefore, none of them is a dominant strategy equilibrium. §

5.2 Comparison of Equilibrium Transfers

Suppose we abrogate the problem of obedience—for example, by allowing the planner to introduce player-specific forcing contracts. Then what remains is a standard hidden information problem with quasilinear preferences and private values. It is well known that such problems can be solved using *Vickrey-Clarke-Groves mechanisms* (Vickrey [30], Clarke [7], Groves [14]).

A VCG mechanism is a revelation mechanism. To begin, the mechanism collects a type report from each player. In response to any profile θ of reports, the mechanism specifies a social alternative $\bar{s}(\theta)$ that is optimal if the reports are truthful:

$$\bar{s}(\theta) \in \operatorname*{arg\,max}_{s \in S} \sum_{j \in I} U^j(s, \theta^j). \tag{1}$$

As we noted earlier, a social alternative in the present context is a strategy profile in S.

The mechanism also specifies transfers $t^i(\theta)$ for each player as a function of the profile of type reports. The transfers are chosen to render truth-telling a weakly dominant strategy. These transfers, paid from the planner to the players, are of the form

$$t^i(\theta) = \sum_{j \neq i} U^j(\bar{s}(\theta), \theta^j) - \tau^i(\theta^{-i}).$$

The first term in this expression equals the total payoffs obtained by i's opponents at the efficient alternative $\bar{s}(\theta)$. The second term is an arbitrary function of i's opponents' reports. It has no effect on player i's incentives, but allows the planner some flexibility in determining the levels of the transfers.[14]

The best known form of the VCG mechanism, the so-called *pivotal* or *Clarke mechanism*, is obtained when $\tau^i(\theta^{-i})$ equals the total payoffs obtained by i's opponents at the strategy profile that maximizes their payoffs, conditional on the reports θ^{-i} being truthful:

[14] One can also generalize our optimal price scheme by modifying prices in a way that does not affect incentives. However, one can show that the price scheme p^* defined above is the only price scheme that implements efficient behavior while never imposing a nonzero price on the outside option.

$$\tau_C^i(\theta^{-i}) = \sum_{j \neq i} U^j(\bar{s}_{-i}(\theta^{-i}), \theta^j), \quad \text{where } \bar{s}_{-i}(\theta^{-i}) \in \operatorname*{argmax}_{s \in S} \sum_{j \neq i} U^j(s, \theta^j). \tag{2}$$

In the present context, in which social alternatives are strategy profiles in an underlying game, we can also consider this interesting alternative:

$$\tau_{mC}^i(\theta^{-i}) = \sum_{j \neq i} U^j(\tilde{s}_{-i}(\theta^{-i}), \theta^j), \quad \text{where } \tilde{s}_{-i}(\theta^{-i}) \in \operatorname*{argmax}_{s \in S:\, s^i = 0} \sum_{j \neq i} U^j(s, \theta^j). \tag{3}$$

The transfer $\tau_{mC}^i(\theta^{-i})$ is similar to $\tau_C^i(\theta^{-i})$, except that it is based on the strategy profile that is optimal for j's opponents *among those in which i chooses the outside option*. In this case, the transfers $t^i(\theta)$ are computed by comparing the payoffs of j's opponents at the strategy profile that is optimal when i is present to their payoffs at the strategy profile that is optimal when i is completely absent from the economy. We call the VCG mechanism corresponding to τ_{mC} the *modified Clarke mechanism*.

The Clarke mechanism, the modified Clarke mechanism, and the optimal price scheme p^* can all be described as instruments that ensure efficiency by "making each player pay for the externalities he imposes on others". Yet the transfers specified by the two VCG mechanisms and by the optimal price scheme are generally different, even in equilibrium.

Example 5.2: Let $S^i = \{0, a, b\}$, and assume that the payoffs to actions a and b are separable and reflect positive externalities: $u_a(x) = v_a(x_a)$ and $u_b(x) = v_b(x_b)$, where v_a and v_b are strictly increasing.

Suppose that under type profile θ, the efficient strategy profile $\bar{s}(\theta)$ assigns player i to the outside option and all other players to action a. As we have seen, player i's transfer at this profile under the optimal price scheme is zero, as the outside option is never priced.

In contrast, the Clarke mechanism may require players assigned to the outside option to pay nonzero transfers, even in equilibrium. In being assigned to the outside option, player i denies his opponents the positive externalities he would bring them were he to be assigned to action a. His equilibrium payment to the planner equals the value of these externalities:

$$-t_C^i(\theta) = (N-1)(v_a(N) - v_a(N-1)) > 0.$$

Under the modified Clarke mechanism, as under the optimal price scheme, players assigned to the outside option neither pay nor receive transfers.[15] §

Example 5.3: In the positive externality setting above, suppose that under type profile θ, the efficient strategy profile assigns all players to action a: $\bar{s}^j(\theta) = a$ for all $j \in I$. Then under the optimal price scheme, each player receives a transfer of

$$-p_a^*(x(\bar{s}(\theta))) = -p_a^*(Ne^a) = (N-1)(v_a(N) - v_a(N-1)) > 0$$

at this state. This quantity is the value of the positive externality that each player creates for his opponents.

Under the two VCG mechanisms, different players obtain different equilibrium transfers, with each player's transfer depending on whether or not his announcement is pivotal. Under the Clarke mechanism, player i is *not* pivotal if the optimal strategy profile $\bar{s}(\theta)$ remains optimal when player i's welfare is ignored: in other words, if $\bar{s}_{-i}^j(\theta^{-i}) = a$ for all $j \in I$. In this case, i's transfer under the Clarke mechanism is zero. If instead the best profile for i's opponents is the one in which all players choose b (i.e., if $\bar{s}_{-i}^j(\theta^{-i}) = b$ for all $j \in I$), then player i is pivotal. In this case, player i must pay the planner

$$-t_C^i(\theta) = (N-1)(v_b(N) - v_a(N)) + \sum_{j \neq i}(\theta_b^j - \theta_a^j) > 0,$$

which is the cost that i's announcement ultimately imposes on his opponents.

Under the modified Clarke mechanism, the meaning of not being pivotal is necessarily different. In the present example, player i is not pivotal under θ if from the point of view of i's opponents, the optimal strategy profile *among those in which i plays the outside option* has *all remaining players* play a: in notation, $\tilde{s}_{-i}^j(\theta^{-i}) = a$ for all $j \neq i$ (and, by definition, $\tilde{s}_{-i}^i(\theta^{-i}) = 0$). If i is not pivotal in this sense, he receives a transfer of

$$t_{mC}^i(\theta) = (N-1)(v_a(N) - v_a(N-1)) > 0,$$

[15] This statement follows from the following general property of the modified Clarke mechanism: if $\bar{s}(\theta)$ is the unique solution to (1) and satisfies $\bar{s}^i(\theta) = 0$, then the solution $\tilde{s}_{-i}(\theta^{-i})$ to (3) is unique and equal to $\bar{s}(\theta)$. For if profile $\hat{s} \neq \bar{s}(\theta)$ (with $\hat{s}^i = 0$) solved (3), then by adding player i's payoff of θ_0^i to the objective function in (3), we would find that $\hat{s}$ does at least as well as $\bar{s}(\theta)$ in problem (1), contradicting the uniqueness of $\bar{s}(\theta)$.

This transfer is equal to the one that *all players* receive at the efficient state under the optimal price scheme. Evidently, the modified Clarke mechanism is *not* a pivotal mechanism, as non-pivotal players can obtain non-zero transfers. If we suppose instead that player i is pivotal, in the sense that $\tilde{s}^j_{-i}(\theta^{-i}) = b$ for all $j \neq i$, then i's payment under the modified Clarke mechanism is

$$-t^i_{mC}(\theta) = (N-1)(v_b(N-1) - v_a(N)) + \sum_{j \neq i}(\theta^j_b - \theta^j_a).^{16} \ \S$$

These examples show that equilibrium transfers under VCG mechanisms can replicate certain specific features of equilibrium prices under the optimal price scheme. But they also show that in general, VCG transfers and equilibrium prices are distinct. Indeed, the second example shows that under both the Clarke and modified Clarke mechanisms, pivotal players' transfers depend explicitly on type announcements. These transfers necessarily differ from those generated by the optimal price scheme, as under the price scheme type announcements are not made.

Given these differences, it may seem paradoxical that each of these mechanisms can be described as "making each player pay for the externalities he imposes on others". But the paradox vanishes once we realize that in each case, the externality for which each player pays is defined with respect to his choice variable in the mechanism at hand. Under the VCG mechanisms, the choice variables are type announcements, so the externality for which the player must pay is determined by the effect that his *announcement* has on the assignment (i.e., the action profile) the planner chooses. Under the optimal price scheme, the choice variables are actions in the underlying game, so the externality for which the player must pay is the direct effect which his choice of *action* has on other players' payoffs. As the examples above show, there is no reason to expect the magnitudes of these externalities to be equal, even in equilibrium.[17]

---

[16] Unlike under the standard Clarke mechanism, we cannot sign the pivotal player's transfer here. By the definition of $\tilde{s}_{-i}(\theta^{-i})$, we know that this transfer would be positive if we replaced $v_a(N)$ with $v_a(N-1)$. Thus, player i's payment to the planner is positive if $v_a(N) - v_a(N-1)$ is sufficiently small, but it is negative if $v_a(N) - v_a(N-1)$ is sufficiently large.

[17] It is interesting to note that like our price scheme, VCG mechanisms can be analyzed by means of potential functions. But unlike the potential functions considered above, potential functions for revelation mechanisms take social alternatives and type profiles as arguments—see Jehiel, Meyer-ter-Vehn, and Moldovanu [17].

5.3 Choosing a Mechanism

In our view, whether a VCG mechanism or a price scheme is preferable depends on the context at hand. In an environment where the planner can collect and evaluate type reports at little cost and can directly choose the allocation of players to strategies, a VCG mechanism would seem to be the appropriate choice. But in some environments, particularly those involving large numbers of agents, these criteria may be violated. When these environments also entail a repeated interaction, it may be preferable to employ a price scheme.

As an example, consider the problem of highway congestion.[18] The planner in charge of a highway network would like drivers to use the network efficiently. But because he lacks information about the drivers' preferences, he does not know what efficient use of the network entails. To address this problem using a revelation mechanism, the planner would need to ask all drivers to submit type reports, use these reports to compute the efficient assignment of drivers to routes, announce this assignment to the drivers, and take measures to ensure that the assignment is obeyed. Compared to this "command and control" approach, price schemes seem far less intrusive: the planner imposes prices directly on the routes, and allows the players to decide for themselves which routes to take.

This same reasoning applies in environments with positive externalities, including the model of consumer technology choice from Section 2, as well as macroeconomic models with positive spillovers and hidden information.[19] In the later setting, the planner would like to coordinate production on certain sectors to the exclusion of others. However, which sectors should be emphasized depends on features of technologies and preferences about which the planner is incompletely informed. Under a revelation mechanism, the planner would ask all agents in the economy to report their preferences; he would then use these reports to compute an efficient allocation of agents to activities, announce this allocation to the agents, and then monitor behavior to ensure obedience. In contrast, the price scheme simply subsidizes each activity in proportion to the externalities it generates. If each individual agent then chooses for himself which activities he will pursue, the agents' aggregate behavior will be efficient in the long run.

---

18 See Sandholm [26, 27].

19 For a survey of work on complementarities in macroeconomics, see Cooper [8].

6. Limitations

The proof of our main theorem relies on a useful property of the logit choice rule: that in potential games, the stationary distribution generated by this rule places nearly all mass on the maximizer of potential when the noise level is small. It is clearly of interest to know the extent to which our results depend on this specification. We believe that our results can be generalized to other choice rules. Recent results of Benaïm and Weibull [2] and Hofbauer and Sandholm [15] provide a framework for investigating this question in settings with large numbers of players.

An important criticism of our result concerns the waiting times needed before the predictions upon which it is based become relevant. It is well known that in evolutionary models which rely upon rare errant choices, the expected time before departure from an equilibrium which is not stochastically stable can grow exponentially in the population size.[20] For this reason, predictions of play which rely upon stochastic stability may be somewhat suspect when the population size is large. On the other hand, it is known that this waiting time critique loses force when agents are geographically dispersed and tend to interact with others who live nearby. Under local interaction, contagion dynamics can cause a stochastically stable strategy which gains a foothold in a single locale to quickly spread throughout the population.[21] It is therefore noteworthy that Blume's selection theorem continues to hold in models of local interaction (Blume [6], Young [32]). This fact gives us hope that the pricing mechanisms studied here can be adapted for use in local interaction models, where long waiting times are far less of an issue. Establishing evolutionary implementation results for such settings is an important question for future research.

References

[1] S. P. Anderson, A. de Palma, J.-F. Thisse, Discrete Choice Theory of Product Differentiation, MIT Press, Cambridge, 1992.

[2] M. Benaïm, J. Weibull, Deterministic approximation of stochastic evolution in games, Econometrica 71 (2003), 873-903.

[3] M. Beckmann, C. B. McGuire, C. B. Winsten, Studies in the Economics of Transportation, Yale University Press, New Haven, 1956.

[20] See Ellison [11], Binmore, Samuelson, and Vaughan [4], and Sandholm and Pauzner [29].

[21] See Blume [5, 6], Ellison [11, 12], Young [32], and Morris [22].

[4] K. G. Binmore, L. Samuelson, R. Vaughan, Musical chairs: modeling noisy evolution, Games Econ. Behav. 11 (1995), 1-35.

[5] L. E. Blume, The statistical mechanics of strategic interaction, Games Econ. Behav. 5 (1993), 387-424.

[6] L. E. Blume, Population games, in: W. B. Arthur, S. N. Durlauf, and D. A. Lane (Eds.), The Economy as an Evolving Complex System II, Addison-Wesley, Reading, MA, 1997, 425-460.

[7] E. Clarke, Multipart pricing of public goods, Public Choice 11 (1971), 19-33.

[8] R. J. Cooper, Coordination Games: Complementarities and Macroeconomics, Cambridge University Press, Cambridge, 1999.

[9] S. N. Durlauf, Statistical mechanics approaches to socioeconomic behavior, in: W. B. Arthur, S. N. Durlauf, and D. A. Lane (Eds.), The Economy as an Evolving Complex System II, Addison-Wesley, Reading, MA, 1997, 81-104.

[10] R. Durrett, Probability: Theory and Examples, 3rd ed., Thomson–Wadsworth–Brooks/Cole, Belmont, CA, 2005.

[11] G. Ellison, Learning, local interaction, and coordination, Econometrica 61 (1993), 1047-1071.

[12] G. Ellison, Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution, Rev. Econ. Stud. 67 (2000), 17-45.

[13] D. Foster, H. P. Young, Stochastic evolutionary game dynamics, Theor. Pop. Biol. 38 (1990), 219-232.

[14] T. Groves, Incentives in teams, Econometrica 41 (1973), 617-631.

[15] J. Hofbauer, W. H. Sandholm, Evolution in games with randomly disturbed payoffs, J. Econ. Theory, forthcoming.

[16] J. Hofbauer, K. Sigmund, Theory of Evolution and Dynamical Systems, Cambridge University Press, Cambridge, 1988.

[17] P. Jehiel, M. Meyer-ter-Vehn, B. Moldovanu, Potentials and implementation, unpublished manuscript, University College London and University of Bonn, 2004.

[18] M. Kandori, G. J. Mailath, R. Rob, Learning, mutation, and long run equilibria in games, Econometrica 61 (1993), 29-56.

[19] M. Kandori, R. Rob, Bandwagon effects and long run technology choice, Games Econ. Behav. 22 (1998), 30-60.

[20] M. L. Katz, C. Shapiro, Systems competition and network effects, J. Econ. Lit. 8 (1994), 93-115.

[21] D. Monderer, L. Shapley, Potential games, Games Econ. Behav. 14 (1996), 124-143.

[22] S. Morris, Contagion, Rev. Econ. Stud. 67 (2000), 57-78.

[23] A. C. Pigou, The Economics of Welfare, Macmillan, London, 1920.

[24] R. W. Rosenthal, A class of games possessing pure strategy Nash equilibria, Int. J. Game Theory 2 (1973), 65-67.

[25] W. H. Sandholm, Potential games with continuous player sets, J. Econ. Theory 97 (2001), 81-108.

[26] W. H. Sandholm, Evolutionary implementation and congestion pricing, Rev. Econ. Stud. 69 (2002), 667-689.

[27] W. H. Sandholm, Negative externalities and evolutionary implementation, Rev. Econ. Stud. 72 (2005), 885-915.

[28] W. H. Sandholm, Affine analysis and potential functions for population games, unpublished manuscript, University of Wisconsin, 2005.

[29] W. H. Sandholm, A. Pauzner, Evolution, population growth, and history dependence, Games Econ. Behav. 22 (1998), 84-120.

[30 W. Vickrey, Counterspeculation, auctions, and competitive sealed tenders, J. Finance 16 (1961), 8-37.

[31] H. P. Young, The evolution of conventions, Econometrica 61 (1993), 57-84.

[32] H. P. Young, Individual Strategy and Social Structure, Princeton University Press, Princeton, 1998.