Let Them Burn Money: Making Elections more Informative

Colin M. Campbell

Department of Economics
Rutgers University
New Brunswick, New Jersey 08901
campbell@econ.rutgers.edu

October, 2008

Abstract

A standard election in which each voter chooses a single alternative permits voters little scope to express the intensity of their preferences. Allowing more complex statements of preferences may not alleviate the problem if voters behave strategically, as only certain statements are credible. I consider the implications of allowing voters to expend money as part of the voting procedure. In an environment with two alternatives and voters with interdependent values, I find necessary and sufficient conditions for all choice functions that are minimally responsive to voter preferences to be implementable with money burning. Furthermore, I show that any choice rule that treats ex-ante identical voters symmetrically can be implemented with an arbitrarily small amount of money burnt per voter as the set of voters is replicated. Thus, for a large electorate, the informational gains of money burning can be reaped at virtually no social cost.

1. INTRODUCTION

The appropriate role for money in the realm of public choice is a matter of controversy. In many political institutions, money is viewed as a factor that may corrupt the process of collective choice and governance, and alter policy outcomes from what would obtain under some benchmark ideal. Examples are plentiful. In elections, monetary incentives or influence may affect which candidate can run a more effective campaign, or even the subset of the electorate that votes. Addressing the former, campaign finance laws constrain the amount of giving to candidates from individual sources, with the goal of making the number of donors to a given candidate of greater importance than the ability to contribute of respective donors. The possible influence of monetary effects on voter turnout has been highlighted recently by a law in the state of Georgia requiring voters to produce photographic identification, the cost of which could be sufficient to make some would-be voters abstain.

The scope for money to affect governance is also clear. In addition to straightforward quid-pro-quo corruption, citizens may implicitly purchase access to members of government by making large donations to campaigns or political committees, with the possibility that the donors could use the access to influence policy decisions in their own favor. Citizens and corporations may also employ professional lobbyists who influence the legislative process directly. The theory of regulatory capture postulates that regulators may act as advocates for the industries they oversee; a circumstance under which this might arise is when regulators are liable to be hired into the regulated industry upon departure from government service, and hence have some financial stake in maintaining the industry's good will.

In all of these examples, policy outcomes tend to skew in favor of individuals who are most willing and able to trade off personal wealth for influence. If a society believes in the principle that, say, collective outcomes are properly determined by the preferences of the numerical majority of the population, then such effects are necessarily unwelcome. However, in some cases it may be desirable to take account of varying intensities of preference over policy choices across citizens. This could be valued for its own sake; it might be allowed that a minority preference should prevail if the average preference of the minority is sufficiently stronger than that of the majority. It could also be valuable if citizen preferences are uncertain and positively related, so that one citizen's preference for a given outcome strengthens the more intensely that other citizens favor that outcome.

The theme of this paper is that money can have a beneficial role to play in environments of collective choice as an instrument for eliciting information about citizens' intensity of preference, as referred to above. To realize theoretical benefits from this possibility, the way in which individual spending by citizens translates into influence over political outcomes must be controlled very precisely. We examine this by taking a particular mechanism design approach to a simple collective choice problem with two alternatives over which citizens must choose. In many problems of mechanism design, an outcome consists of some non-monetary action (e.g., an exchange of goods) combined with transfers of money between participants. In our environment, the non-monetary action is the choice of winning alternative, but we take the approach that as a result of the transfer, some money may be "burnt," i.e., a proportion of any money expended by an individual is lost. We make this departure to capture the dissipative nature of much influence activity; we could substitute some notion of costly, nonproductive effort for money in our model and reach identical conclusions. By assuming that money expenditures are wasteful, we impose a high standard on the gross benefits of including money in the collective choice mechanism.

The nature of our results is as follows. First, we explore the degree to which allowing money burning can sharpen the responsiveness of voting outcomes to agents' preferences. Specifically, we ask under what conditions any choice mechanism that is minimally responsive to agent preferences, in a sense that is made precise, can be effected as an equilibrium of a voting mechanism with money burning. We find a necessary and sufficient condition for this possibility; roughly, it requires that an agent's private information be more important to his preferences over the alternatives than it is as a signal about other agents' information.

Our second result addresses the fact that any informational gains from money burning come at the cost of lost resources. Under a more structured assumption about the environment that allows well-defined replication, we find that as the number of agents grows, any mechanism that treats ex-ante identical agents symmetrically must have the property that the expected amount of money burnt by an agent in any equilibrium converges to zero. That is, the informational benefits of allowing money burning comes at a vanishing cost per agent as the size of the electorate grows. This holds because under symmetric treatment, the likelihood that a single agent's behavior changes the winning alternative converges to zero as the number of agents grows; thus, no agent will be willing to burn very much money to try to influence the outcome. However, as long as some money burning is possible, there is full scope for extracting information about intensity of preferences.

We emphasize that our results do not argue for any of the specific institutions or practices cited at the beginning of this introduction. Our goal is to show that a carefully designed (and indeed, potentially highly complex) system for allowing money to influence collective choice outcomes can have benefits that are purely informational. Unwanted side effects on the equality of representation may be probable, or even unavoidable. We do, however, provide an example in which a reasonable notion of equity may be preserved, in a later section of the paper.

We motivate some of the potential benefits of money burning with the following example. Suppose there is a group of n voters who will collectively choose one of two social alternatives, A or B. Every voter's nonmonetary benefit if B wins is 0. A given voter i's nonmonetary benefit if A wins is an amount x_i, observed only by voter i. Each type x_i is an independent draw from a normal distribution with mean 0 and variance 1.

A social planner must design a mechanism that elicits information from the voters and chooses a winning alternative based on their actions. One possibility is to hold a simple vote, in which each voter takes one of two actions, voting for A or voting for B.[1] An election rule would specify which of the two alternatives wins (or, more generally, with what probability each alternative wins) for every possible profile of the n votes. A majority rule, for instance, would specify that the alternative receiving more votes is the winner, along with some probabilities of each alternative winning in the event of a tie vote.

It can be seen that in any Bayes-Nash equilibrium of any election rule (majority or otherwise), a given voter i's optimal vote must be the same for all strictly positive valuations x_i for A, and her optimal vote must be the same for all strictly negative valuations x_i for A. This is because, given strategies of the other players, a player whose x_i is positive wishes to take the action maximizing the probability that A wins, while a player whose x_i is negative has the opposite incentive. Thus, an election rule can elicit information about voters' ordinal preferences over A and B, but can elicit no information about their cardinal preferences.[2]

Consider instead a different mechanism. Each voter i takes an action with two components: she casts a vote for A or for B, and she burns a nonnegative amount of money c_i (i.e., c_i is lost and cannot be redistributed). Given her action, the voter's final utility is $x_i - c_i$ if A wins the election, $-c_i$ if B wins the election. The rule for determining the winning alternative is as follows. Any voter who votes for A and burns amount of money c is given a weight $\alpha(c)$, with $\alpha(c)$ solving

$$c = \sqrt{\frac{n-1}{2\pi}}(1 - e^{-\frac{\alpha(c)^2}{2(n-1)}})$$

[1] Allowing abstention would have no effect in this particular example.

[2] This is not strictly accurate, since, if a player is indifferent about which alternative to vote for in an equilibrium, he may use his value of x_i to determine how to vote. However, in such an instance, that player's indifference implies that his vote must have no average effect on the outcome, and so the information he reveals about his type plays no role.

for $c < \sqrt{(n-1)/2\pi}$, $\alpha(c) = 0$ otherwise. Any voter who votes for B and burns c is given weight $\beta(c) \equiv -\alpha(c)$. To determine the winning alternative, the weights of all voters are summed; A wins if the sum is strictly positive, B wins if the sum is strictly negative, and the winner is chosen according to some arbitrary but fixed rule if the weights sum to 0. Thus, any voter who votes for A can increase the probability that A wins by burning a larger quantity of money, and any voter who votes for B can increase the probability that B wins by burning a larger quantity of money.

In this mechanism, it is a Bayes-Nash equilibrium for every voter i to vote for her preferred alternative, and to burn the amount of money c such that her weight $\alpha(c)$ or $\beta(c)$ is equal to her type x_i. Note that when the voters behave in this fashion, the mechanism effectively implements as the winner the alternative that maximizes the sum of the voters' payoffs, gross of money burnt, an outcome that is impossible with simple voting. To see that this behavior is an equilibrium, suppose that all voters other than i act in accordance with the proposed strategies, and that i's type is $x_i > 0$; the argument is fully analogous for $x_i < 0$. Clearly, i wishes to vote for A, irrespective of how much money she burns. If i burns an amount of money that gives her vote weight $\hat{\alpha}$, the probability that A will win is the probability that the sum of other voters' types, which is equal to the sum of the other voters' weights under the proposed behavior, is greater than or equal to $-\hat{\alpha}$. As each voter's type is independently distributed as a standard normal, the sum of $n-1$ types is distributed normally with mean 0 and variance $n-1$. Thus, the expected payoff to i of inducing weight $\hat{\alpha}$ is

$$x_i \int_{-\hat{\alpha}}^{\infty} \frac{1}{\sqrt{2\pi(n-1)}} e^{-\frac{t^2}{2(n-1)}}\, dt - \sqrt{\frac{n-1}{2\pi}}\left(1 - e^{-\frac{\hat{\alpha}^2}{2(n-1)}}\right)$$

$$= \frac{x_i}{2} + \int_{0}^{\hat{\alpha}} \frac{(x_i - t)}{\sqrt{2\pi(n-1)}} e^{-\frac{t^2}{2(n-1)}}\, dt.$$

This payoff is maximized at $\hat{\alpha} = x_i$ for all $x_i > 0$, so i best responds by also burning an amount of money inducing a weight equal to her type x_i.

Allowing money burning in this example thus permits the choice of winning alternative to respond to the intensity of voters' preferences in a way that simple voting does not. However, these informational gains come at the social cost of the money that is burnt, and whether total welfare is improved relative to a given simple voting mechanism depends on the magnitude of money burnt. To check on this, consider the amount of money burnt by a voter whose preference type is x, $\sqrt{(n-1)/2\pi}\left(1 - e^{-\frac{x^2}{2(n-1)}}\right)$. Although this expression is not globally decreasing in the total number of voters n, it does have the property that it converges to zero as n grows unboundedly, for all x. Furthermore, the ex-ante expected amount of money burnt by a voter who has not yet learned his type also vanishes as the number of voters grows. That is, the mechanism induces the voters to reveal all of their preference information, and makes use of that information, with an efficiency loss that vanishes when there are many voters.

We expand on this example to generalize its results in the next section.

2. THE MODEL

One of two alternatives, A and B, must be chosen for a society of I agents, indexed by i. Information in the society is summarized by a vector of agent types $x \equiv (x_1, x_2, \ldots, x_I)$, with each $x_i \in \Re$. Agent i can observe only her own type x_i. The benefit to agent i of A being chosen is $u_i(x)$, while the benefit of B being chosen is normalized to zero for all agents. Thus, $u_i(x) > 0$ means that i prefers alternative A ex-post, other things equal, while $u_i(x) < 0$ means that i prefers B ex-post. $u_i(\cdot)$ is assumed to be continuous and weakly increasing in its arguments, so that higher types mean that A is relatively better for all agents. At times it will be convenient to write $u_i(x)$ as $u_i(x_i, x_{-i})$, where x_{-i} is the vector of types other than i's. For one of our main results, we will assume that type vectors x are distributed with a continuous density function $f(\cdot)$ on $\Re^I$, where we write $f(x_{-i}|x_i)$ for the density of x_{-i} conditional on x_i. For the other result, we will assume a finite type space for agents, with comments on extension of the result to a continuous space. For an unbounded type space, we also impose the following condition bounding expected benefits. For any agent i and any vector of types x, let the indicator function $1_i(x)$ equal 1 if $u_i(x) > 0$, 0 else, and the indicator function $1'_i(x)$ equal 1 if $u_i(x) < 0$, 0 else. We assume that the functions of x_i $\int_{\Re^{I-1}} 1_i(x_i, x_{-i}) u_i(x_i, x_{-i}) f(x_{-i}|x_i) dx_{-i}$ and $\int_{\Re^{I-1}} 1'_i(x_i, x_{-i}) u_i(x_i, x_{-i}) f(x_{-i}|x_i) dx_{-i}$ (where we abuse notation in writing x_{-i} as a scalar) are bounded above and below for all i.[3]

The mechanism via which an alternative is chosen for the society works as follows. Simultaneously and independently, each agent i sends a message m_i and burns a nonnegative amount of money c_i. In the spirit of direct revelation mechanisms, the space of messages an agent can send is the space of possible types (the real numbers), augmented by a "null message" $m_\emptyset$. For every type x_i of agent i, there is an associated amount of money $c_i(x_i)$ that that type is expected to burn. If the agent announces a type $\hat{x}_i$ and burns money equal to $c_i(\hat{x}_i)$, then the mechanism responds as if the agent's type is $\hat{x}_i$; if the agent burns a different amount of money, then the mechanism responds as if the agent had sent the null message $m_\emptyset$. Formally, an agent's reported message m_i is translated

[3] This restriction in fact excludes the specification of the example in the introduction, as benefits are unbounded there. The conclusions of the general results nevertheless hold for that example, due to properties of the normal distribution in its tails.

into a registered message $m(m_i)$ according to the function $m(m_i) = m_i$ if $c_i = c_i(m_i)$, $m(m_i) = m_\emptyset$ if $c_i \neq c_i(m_i)$. We note in advance that in an equilibrium, each agent will announce his type x_i truthfully and burn money in the amount $c_i(x_i)$; the null message is included only to ensure that agents have the opportunity not to participate in the mechanism, and in particular, not to burn any money.

A *choice rule* $\pi : (\Re \times \{m_\emptyset\})^I \rightarrow [0,1]$ determines the winning outcome from the registered messages; $\pi(m)$ is the probability that A is selected when the set of registered messages is m, while B is implicitly selected with probability $1 - \pi(m)$. A mechanism is characterized by the functions $\{\pi(\cdot), \left(c_i(\cdot)\right)_{i=1}^{I}\}$. Only certain choice rules can be part of mechanisms that properly induce agents to reveal all information. We focus on mechanisms $(\pi(\cdot), (c_i(\cdot))_{i=1}^{I})$ that are Bayesian incentive compatible, or for which it is a Bayes-Nash equilibrium for all agents to register their types truthfully. The incentive compatibility condition for type x_i of agent i is

$$x_i \in \underset{\hat{x}_i \in \Re \cup \{m_\emptyset\}}{\text{argmax}} \int_{\Re^{I-1}} \pi(\hat{x}_i, x_{-i}) u_i(x_i, x_{-i}) f(x_{-i}|x_i) dx_{-i} - c_i(\hat{x}_i),$$

where we take $c_i(m_\emptyset) = 0$ for all i; this is without loss of generality given agent incentives, as reporting the null message and burning a positive amount of money is strictly dominated by reporting the null message and burning no money, for every type of agent. Incentive compatibility must hold for every $x_i \in \Re$ and for every i for the mechanism to implement its choice rule. We say that a choice rule $\pi(\cdot)$ is *implementable with money burning* if there is some mechanism $(\pi(\cdot), (c_i(\cdot))_{i=1}^{I})$ that is Bayesian incentive compatible.

We are interested in the implementability of a particular class of choice rules that have some desirable welfare properties. We define a choice rule $\pi(x)$ as *monotonic* if it is weakly increasing in its arguments on the domain in which no agent sends the null message. Given the assumption that each agent's benefit from alternative A, $u_i(x)$, is weakly increasing in all types, a choice rule that increases the probability of selecting A as types increase may be considered minimally responsive to social preferences.

As the example in the previous section demonstrates, it may not be possible to implement some choice rules, even if they are monotonic, when the only instrument available for the expression of preferences is a costless report. When the selection mechanism may be augmented with money expenditures, the set of choice rules that can be implemented must weakly expand. Our first main result provides the exact conditions under which any monotonic choice rule may be implemented with money.

PROPOSITION 1: *Every monotonic choice rule* $\pi(\cdot)$ *is implementable with money if and only if for every agent* i, *every profile of types of other agents* x_{-i}, *and every pair of types* x_i *and* x_i' *for* i *such that* $x_i \geq x_i'$, $u_i(x_i, x_{-i})f(x_{-i}|x_i) \geq u_i(x_i', x_{-i})f(x_{-i}|x_i')$.

PROOF: See appendix.

We comment on the significance of the necessary and sufficient condition in Proposition 1. The condition is that the function $u_i(x_i, x_{-i})f(x_{-i}|x_i)$ be nondecreasing in x_i, for every x_{-i}. This function is agent i's ex-post payoff from A winning the election under a given profile of types, weighted by the probability of that realization of types given i's own type. By assumption, $u_i(x_i, x_{-i})$ is increasing in x_i, so a corollary of Proposition 1 is that all monotonic choice rules are implementable with money burning when agents' types are statistically independent. More generally, the condition requires that the conditional density $f(x_{-i}|x_i)$ not vary more, in a particular direction, than does $u_i(x_i, x_{-i})$. Restricting the condition to two signals for i, $x_i \geq x_i'$, when $u_i(x_i, x_{-i}) \geq 0$ and $u_i(x_i', x_{-i}) \leq 0$, the condition is satisfied for any specification of $f(\cdot)$. However, if $u_i(x_i, x_{-i})$ and $u_i(x_i', x_{-i})$ are both negative, then the ratio $f(x_{-i}|x_i)/f(x_{-i}|x_i')$ must lie below the upper bound $u_i(x_i', x_{-i})/u_i(x_i, x_{-i}) \geq 1$, and if $u_i(x_i, x_{-i})$ and $u_i(x_i', x_{-i})$ are both positive, then the ratio $f(x_{-i}|x_i)/f(x_{-i}|x_i')$ must lie above the lower bound $u_i(x_i', x_{-i})/u_i(x_i, x_{-i}) \leq 1$. Roughly speaking, the condition requires that an agent's type be of greater importance to that agent's preferences over the alternatives A and B, than to that agent's assessment of what types other agents are likely to be.

In many models of agents with ordered types, it is assumed that the types are statistically affiliated. In auction models, for instance, assuming affiliation may simplify construction of equilibrium bids. Here, the standard is higher, in that we require our condition to guarantee implementability of a large class of choice rules, rather than equilibrium of a specific game. Technically, our condition guarantees satisfaction of a single-crossing property, which is that under any monotonic mechanism, the gain of reporting a higher type rather than a lower type is greater (gross of any money burnt) to a higher type of agent than to a lower type of agent. This may not hold if two types prefer the same alternative, but the type who has a weaker (i.e., closer in absolute value to 0) preference has a substantially greater potential to pivot the winning outcome via her report; affiliation does not rule out this possibility. Although distributions $f(\cdot)$ that satisfy our condition for a given set of $u_i(\cdot)$ will typically include some that do not satisfy affiliation, so that our condition is not stronger per se, affiliation is usually invoked only as a sufficient condition, and one that has the desirable property of not depending on agents' preferences.[4]

[4] We suspect that our condition would also be necessary and sufficient for incentive compatibility in an all-pay

The preceding sections highlight the potential value of money as an instrument to elicit information about agent preferences. Balancing this value is the resource cost of burnt money, which we must account in the assessment of any decision mechanism. In addition, fairness or equity concerns may make important how the costs of the mechanism are distributed across the population, and whether some agents are systematically favored by a system in which money has a role.

To explore these issues, we construct a variation of our model. First, we assume that there is a finite set X of possible private signals for agents. In addition to the private signal x_i, each agent i is assumed to have a *public type* θ_i from a set Θ. Each agent i's θ_i is assumed to be common knowledge among all agents, as well as observable to the social planner who implements a decision mechanism. When two agents have the same public type, they have symmetric benefit functions, and their private signals enter all benefit functions and the joint distribution $f(\cdot)$ symmetrically. Formally, referring to any two agents i and j, write a profile of private signals x as (x_i, x_j, x_{-ij}). Our assumption is that if $\theta_i = \theta_j$, $x_i = x'$, and $x_j = \hat{x}$, then $u_i(x', \hat{x}, x_{-ij}) = u_j(\hat{x}, x', x_{-ij})$ for all x', $\hat{x}$, and x_{-ij}; $u_k(x', \hat{x}, x_{-ij}) = u_k(\hat{x}, x', x_{-ij})$ for all $k \notin \{i, j\}$, x', $\hat{x}$, and x_{-ij}; and $f((x', \hat{x}, x_{-ij}) = f(\hat{x}, x', x_{-ij})$ for all x', $\hat{x}$, and x_{-ij}. Two agents with the same public type can have different incentives, due to different private signals, but they are ex-ante identical, and their private information affects all other agents in a symmetric fashion.

Our second main result considers money burning mechanisms for a large number of agents. To allow for some agent heterogeneity (ex ante) while performing a well-defined comparative static, we make a simplification to the environment for the purpose of exploring this result. We assume specifically that there is an underlying state of the world ω, chosen randomly from a finite set Ω. The agents and mechanism designer cannot observe the realized ω, but it is common knowledge that ω is drawn according to probabilities $p(\omega)$. Agent i's ex-post benefit from alternative A winning depends only on i's private signal, her public type θ_i, and the realized state, $\tilde{u}(x_i, \theta_i, \omega) \equiv \tilde{u}_i(x_i, \omega)$. In addition, the private signals of two agents with the same public type are assumed to be independently and identically distributed conditional on the realization of ω. Given the assumption on preferences, agent i's benefit function $u_i(x_i, x_{-i})$ can be written as the conditional expectation

$$\frac{\sum\limits_{\omega \in \Omega} \tilde{u}_i(x_i, \omega) f(x|\omega, \theta) p(\omega)}{\sum\limits_{\omega \in \Omega} f(x|\omega, \theta) p(\omega)},$$

and the joint distribution of private signals $f(x)$ is

$$\sum_{\omega} f(x|\omega, \theta) p(\omega).$$

auction environment of all selling rules in which the probability of winning is increasing in own type.

To generate a large population of agents that preserves certain size-independent characteristics, we use replication. Given an environment with n agents having public types $(\theta_1, \theta_2, \ldots, \theta_n)$, benefit function $\tilde{u}(\cdot)$, and signal distribution $f(\cdot|\omega, \theta)$, the *T-fold replication* of this environment is an environment with the same benefit and conditional signal distribution functions, and T agents of public type θ_1, T of public type θ_2, etc.

We prove our result for a particular class of mechanisms with money burning. A choice function $\pi(\cdot)$ is said to be *anonymous up to public type* if it is symmetric for any two agents who have the same public type. Formally, if $\theta_i = \theta_j$, then $\pi(\tilde{x}, x', x_{-ij}) = \pi(x', \tilde{x}, x_{-ij})$ for all i, j, $\tilde{x}$, x', and x_{-ij}. Choice functions in this class treat all agents with the same public type equally, in the sense that it responds only to the entire profile of signals reported by a given class of agent, and not to which agent within that class makes which report. Such a choice function may differentiate across agents with different public types.

Our result is a statement about the amount of money that can be burnt in an equilibrium that implements such a choice function:

PROPOSITION 2: *For any n agent environment and any $\epsilon > 0$, there exists $\overline{T}$ such that in any T-fold replication with $T > \overline{T}$, any choice function that is anonymous up to public type and implementable with money burning may be implemented such that the money burnt by every type of every agent is less than ϵ.*

This result states that any choice function that is anonymous up to public type can be implemented with vanishing burnt money in the limit. Thus, in a large society, it is possible to reap all of the potential informational gains associated with money burning with negligible inefficiency. The combination of replication and anonymity up to public type ensure that in the relevant limit, the report made by any one agent must have a vanishing effect on the outcome. Thus, no agent will find it in her interest to burn more than a very small amount of money; however, the amount of money she would be willing to burn in principle to affect the outcome is always positive, so the value of money burning as an instrument does not diminish.

Proposition 2 characterizes a limiting property of incentive compatible mechanisms, but is silent on how these mechanisms change as a class under replication. For instance, we may inquire whether the necessary and sufficient condition for implementability with money burning of all monotonic choice mechanisms found in Proposition 1 is robust to replication in the environment constructed for Proposition 2. Letting $p(\omega|x, \theta)$ be the posterior probability of state ω to an agent of private type x and public type θ, the condition of Proposition 1 is necessarily satisfied, for any replication, if the product $u(x, \theta, \omega)p(\omega|x, \theta)$ is increasing in x for all ω and θ. Like the original condition, this restricts type x to contain more information about private preferences than about the realization of other payoff-relevant variables. We note that this condition precludes a "pure common value" environment, in which $u(x, \theta, \omega)$ does not depend on x.

We obtain Proposition 2 for a finite set of types for agents, and the proof depends critically on the assumption that all types occur with strictly positive probability. In fact, the result, as stated, does not hold for an infinite type space. We conjecture that a similar result could be recovered if some conditions were put on the sequence of choice functions so that each successive choice function were forced to extend the prior choice function in some appropriate way.

3. DISCUSSION

As discussed in the introduction, there are many possible concerns about a system of public choice that explicitly or implicitly allows money a role in determining outcomes. The issue of corruption or purchased influence is outside the scope of our model, as we assume the alternatives from which voters choose to be fixed, and in particular independent of the flow of money. More relevant are issues of equal representation arising from the direct effect of spending on electoral outcomes. In general, a system in which a voter's say increases in the money she expends may be expected to result in the relative overrepresentation of those best able to pay.

What mitigates these concerns is that we are considering an environment in which the particular influence of money can be controlled precisely, as formalized by the specification of the voting mechanism. That is, while a given reduced-form mechanism that maps money burnt (or money contributed, e.g., in some models of lobbying) can have undesirable properties with respect to equity, the scope to choose the mechanism may allow one to adjust for these problems.

As an example, consider the following environment. Invoking the semi-symmetric model established in the previous section for the statement of Proposition 2, let voter i's public characteristic θ_i be a strictly positive real number, entering her benefit function $u(x_i, x_{-i}, \theta_i)$ as $(1/\theta_i)u(x_i, x_{-i})$. Her payoff in an election from burning c dollars is therefore $(1/\theta_i)u(x_i, x_{-i}) - c$ if A wins, $-c$ if B wins. By an affine transformation, these are the same preferences she would have if her payoff were $u(x_i, x_{-i}) - \theta_i c$ if A wins, $-\theta_i c$ if B wins. Thus, we may interpret θ_i as voter i's marginal utility of income. Assume also that the joint distribution of private types $(x_1, \ldots, x_n)$ is symmetric; by the symmetry assumptions on $u(\cdot)$, voters are thus distinguished ex-ante only by their marginal utilities of income.

Consider now a mechanism with money burning that treats voters symmetrically solely with respect to their private types x_i, i.e., the induced choice function $\pi(x)$ is symmetric in the types. Such a mechanism may properly be said to treat voters equally in an ex-ante sense, and in particular does not discriminate between voters on the basis of their willingness to trade off money for electoral outcomes. Let $c_i(x_i)$ be the money burning function for voter i in the mechanism, where voter i has public characteristic θ_i. By symmetry of the distribution of private types and of the choice function $\pi(\cdot)$, the mechanism is supported by specifying a money burning function for voter j with public characteristic θ_j equal to $c_j(x_j) = (\theta_i/\theta_j)c_i(x_j)$ for all x_j. That is, when voters i and j have the same private type $\hat{x}$, they burn money in the proportion θ_j/θ_i. Thus, a voter with a greater marginal utility of income is required to burn less money in order to express a given preference. In this example, when marginal utility of income can be observed by the mechanism designer, a mechanism that strives for equality in representation has the natural property that it requires those who value income highly to burn less money. This contrasts sharply with explicit or implicit institutions in which influence is always proportional to total expenditures.

Related Literature

Our work has precedent in various strands of literature. One strand is models of influence buying, which share with ours the feature that a group of agents can spend resources to affect collective decisions. This work includes, among others, Tullock (1967), Becker (1983), Bernheim and Whinston (1986), and Grossman and Helpman (1994). The main goals of these papers are to characterize equilibria and comparative statics of influence games; in Bernheim and Whinston and in Grossman and Helpman, in which resources are transferred to a decision maker rather than dissipated, there is also focus on the potential for efficient decisions to be made in equilibrium. Our work is different from the cited papers in two especially important respects. First, these papers assume no private information among agents, meaning that there is perforce no informational role for expenditures to play. Second, these papers each assume a particular specification for the way that expenditures map into final decisions, while ours is a mechanism design approach that considers a general class of such mappings.

The issue of information aggregation in simple elections is treated in Feddersen and Pesendorfer (1997). They also consider an environment in which a group of voters chooses between two alternatives. Each voter may vote for one of the two alternatives, and the winning alternative is determined according to a prespecified proportional threshold (possibly a supermajority rule). Voters have two-dimensional private information: a private characteristic that affects only that voter's payoff; and a signal about an underlying state that affects all voters' payoffs. Feddersen and Pesendorfer wish to characterize when voters will successfully aggregate their information about the state, in the following particular sense: the alternative that wins the election is the alternative that would have won had the state been common knowledge, and the voters voted simply according to their private characteristics. They find sufficient conditions such that the probability of this coincidence goes to one as the number of voters grows.

A recent literature explores possibilities for eliciting information about intensity of preferences when the alternatives over which voters choose are multidimensional. For instance, votes on several referenda may be held simultaneously, so that an outcome is a description of the winner in each individual contest. Similarly, if there are multiple elections over time, then voters have preferences over a sequence of winners. Several papers, including Cassella (2005), Cassella and Gelman (2008), and Jackson and Sonnenschein (2007), observe that if voters are given a "budget" of votes to distribute over the individual contests, rather than being limited to a single vote on each contest, then information can be gained about voters' intensities of preference, as a voter has an incentive to trade off votes in a contest in which her preference is weak for one in which her preference is strong. The papers show, to varying degrees, that the informational gains from vote budgets help result in more efficient outcomes. Jackson and Sonnenschein show specifically that their budgeting rule allows for sufficient information revelation for implementation of the efficient outcome in the limit. One way of summarizing the contrast between our limiting result for large electorates and those of Jackson and Sonnenschein and of Feddersen and Pesendorfer is that they show that costless voting mechanisms can achieve approximate efficiency in the limit, while our result shows that a costly mechanism can achieve perfect efficiency, gross of those costs, at a cost that vanishes in the limit.

Our result for large electorates has a close precedent in work by McLean and Postlewaite (2002, 2004) on mechanism design with informationally small agents. They show that if the impact of every agent's private information on aggregate beliefs about a payoff-relevant underlying state goes to zero, then in a mechanism that implements the efficient outcome, the cost of inducing agents to reveal that information, as measured by informational rents earned by the agents, vanishes also. Our limiting result, which is for all mechanisms that treat ex-ante identical agents symmetrically, includes but does not restrict attention to efficient decision rules (which maximize the sum of agent utilities gross of money burnt). Furthermore, our specification in which a voter's type serves as both a private preference parameter and information about an underlying state is not covered by their results.

The limiting result of Feddersen and Pesendorfer in particular may cast some doubt on the usefulness of money burning, since it seems that simple, costless voting mechanisms may effectively aggregate information in the limit anyway. We thus provide an example not covered by Feddersen and Pesendorfer's environment to demonstrate the potential power of money burning. Taking the replication environment described previously, assume that there are two underlying states, ω_1 and ω_2; that voters are ex-ante symmetric (i.e., they have identical public characteristics); and that voters have purely private values ($u_i(x_i, \omega) = x_i$ for $\omega \in \{\omega_1, \omega_2\}$). Thus, the only effect of ω is to alter the distribution of voter types, $F(x_i|\omega)$. Assume further that in both states, the ex-ante probability that a voter prefers A to B is identical: $F(0|\omega_1) = F(0|\omega_2)$. In a simple election in which voters only vote for A or vote for B, the unique equilibrium is for all voters to vote for their preferred alternative; this behavior does not reveal anything about the realized state, as the probabilistic distribution of votes is the same in either state. Thus, any fixed election rule yields identical expected results across states, for any number of voters.[5]

In contrast, if $x_i p(\omega_j|x_i)$ is weakly increasing in x_i for all ω_j, then any monotonic choice rule can be implemented with money burning, and at vanishing cost in the limit if voters are treated symmetrically. In particular, since voters can be induced to reveal their profile of types exactly, it would be possible to implement a mechanism such that in the limit, one alternative wins almost surely in state ω_1, and the other alternative wins almost surely in state ω_2.

4. CONCLUSION

We have shown a theoretically beneficial role for money in elections, and that for a large electorate, these benefits are robust even when the money that is expended is dissipated. A natural question is how the theoretical benefits might be realized in practice. As with many problems of mechanism design, the mechanisms that achieve certain desiderata are liable to be complex, and to depend very particularly on details of the specifications of preferences and underlying uncertainty. The controversy that would likely surround any attempt formally to institutionalize expenditures into the workings of public choice would only be exacerbated by a lack of transparency, particularly if individual citizens were treated differently.

Against this, it may be argued that systematic differences in participation already exist, even in institutions seemingly untouched by the influence of money. In elections in the United States,

[5] Because there is correlation of types, it is theoretically possible that there is a more revealing equilibrium of a mechanism in which voters do not burn money, but send richer messages than simply voting for A or voting for B. However, the scope for such mechanisms to reveal information in the private values case is limited, more so the closer the types are to being independently distributed.

in which voting is voluntary, there are patterns in participation suggesting that the act of voting itself is costly, and that those for whom the cost is lowest tend to vote more often, ceteris paribus. An example in support of this is the high rate of turnout among retirees. Those for whom getting to the polls is particularly burdensome, such as single parents and people holding multiple jobs, may be expected to abstain more on average. Hence, the default alternative to an electoral system that allows money burning is not one in which participation is costless (or equally costly) to all participants, and some form of expenditure mechanism or transfers may be needed simply to redress such inequities.

APPENDIX

Proof of Proposition 1. We first prove that the stated condition on the $u_i(\cdot)$ and $f(\cdot)$ (called "condition *" hereafter) is necessary for all monotonic choice rules to be implementable with money burning. Suppose that condition * is violated, i.e., for some i, some x_{-i}, and some $\tilde{x}_i > x_i'$, $u_i(\tilde{x}_i, x_{-i}) f(x_{-i}|\tilde{x}_i) < u_i(x_i', x_{-i}) f(x_{-i}|x_i')$. By continuity of the $u_i(\cdot)$ and $f(\cdot)$ functions, there necessarily exists an open rectangle X_{-i}, i.e., the set of all type profiles for other agents in which each x_j lies in some interval $(\underline{x}_j, \overline{x}_j)$, such that $u_i(\tilde{x}_i, x_{-i}) f(x_{-i}|\tilde{x}_i) < u_i(x_i', x_{-i}) f(x_{-i}|x_i')$ for all $x_{-i} \in X_{-i}$. Consider the monotonic choice rule $\pi^*(\cdot)$ with the following features on the domain in which no agent makes the null report. If $x_j \leq \underline{x}_j$ for any $j \neq i$, then $\pi^*(x_i, x_{-i}) = 0$ for all x_i. If $x_j > \underline{x}_j$ for all j and $x_j \geq \overline{x}_j$ for any j, then $\pi^*(x_i, x_{-i}) = 1$ for all x_i. If $x_{-i} \in X_{-i}$ and $x_i \leq x_i'$, $\pi^*(x_i, x_{-i}) = 0$; if $x_{-i} \in X_{-i}$ and $x_i > x_i'$, $\pi^*(x_i, x_{-i}) = 1$. Let $c_i(x_i)$ be the money burning function for i in an arbitrary mechanism with $\pi^*(\cdot)$ as the choice rule. Necessary conditions for the mechanism to be Bayesian incentive compatible are

$$\int\limits_{\Re^{I-1}} \pi(\tilde{x}_i, x_{-i}) u_i(\tilde{x}_i, x_{-i}) f(x_{-i}|\tilde{x}_i) dx_{-i} - c_i(\tilde{x}_i) \geq \int\limits_{\Re^{I-1}} \pi(x_i', x_{-i}) u_i(\tilde{x}_i, x_{-i}) f(x_{-i}|\tilde{x}_i) dx_{-i} - c_i(x_i')$$

and

$$\int\limits_{\Re^{I-1}} \pi(x_i', x_{-i}) u_i(x_i', x_{-i}) f(x_{-i}|x_i') dx_{-i} - c_i(x_i') \geq \int\limits_{\Re^{I-1}} \pi(\tilde{x}_i, x_{-i}) u_i(x_i', x_{-i}) f(x_{-i}|x_i') dx_{-i} - c_i(\tilde{x}_i).$$

A necessary condition for satisfaction of these inequalities is

$$\int\limits_{\Re^{I-1}} (\pi(\tilde{x}_i, x_{-i}) - \pi(x_i', x_{-i}))(u_i(\tilde{x}_i, x_{-i}) f(x_{-i}|\tilde{x}_i) dx_{-i} - u_i(x_i', x_{-i}) f(x_{-i}|x_i') dx_{-i}) \geq 0.$$

However, this condition is violated for the constructed choice rule: the first term in the integrand is 0 except for on $x_{-i} \in X_{-i}$, where it is strictly positive, but the second term in the integrand is strictly negative on $x_{-i} \in X_{-i}$ by assumption.

We show that condition * is sufficient for implementability of all monotonic choice rules with money burning by construction. Let $\pi(\cdot)$ be an arbitrary monotonic choice rule, and fix an agent i. By the monotonicity of $\pi(x_i, x_{-i})$ in x_i and the assumed continuity of $u_i(x_i, x_{-i})$ and $f(x_i, x_{-i})$, the function $\pi(x_i, x_{-i})\frac{\partial}{\partial x_i}\{u_i(x_i, x_{-i})f(x_{-i}|x_i)\}$ is Riemann-Stieltjes integrable in x_i for any x_{-i}. Furthermore, boundedness of the expectations $\int_{\Re^{I-1}} 1_i(x_i, x_{-i})u_i(x_i, x_{-i})f(x_{-i}|x_i)dx_{-i}$ and $\int_{\Re^{I-1}} 1_i'(x_i, x_{-i})u_i(x_i, x_{-i})f(x_{-i}|x_i)dx_{-i}$ guarantees that the expectation over x_{-i},

$$\int_{\Re^{I-1}} \pi(x_i, x_{-i})\frac{\partial}{\partial x_i}\{u_i(x_i, x_{-i})f(x_{-i}|x_i)\}dx_{-i},$$

is also Riemann-Stieltjes integrable. Choose an arbitrary type x_i^0, and define

$$\begin{aligned}\hat{c}_i(x_i) \equiv &\int\limits_{\Re^{I-1}} \pi(x_i, x_{-i})u_i(x_i, x_{-i})f(x_{-i}|x_i)dx_{-i} - \int\limits_{\Re^{I-1}} \pi(x_i^0, x_{-i})u_i(x_i^0, x_{-i})f(x_{-i}|x_i^0)dx_{-i} \\ &- \int\limits_{x_i^0}^{x_i}\int\limits_{\Re^{I-1}} \pi(t, x_{-i})\frac{\partial}{\partial t}\{u_i(t, x_{-i})f(x_{-i}|t)\}dx_{-i}dt.\end{aligned}$$

By the monotonicity of $\pi(\cdot)$ and the hypothesis that $u(x_i, x_{-i})f(x_{-i}|x_i)$ is increasing in x_i, we have

$$\begin{aligned}&\int\limits_{\Re^{I-1}} \pi(x_i, x_{-i})(u_i(x_i, x_{-i})f(x_{-i}|x_i) - u_i(x_i^0, x_{-i})f(x_{-i}|x_i^0))dx_{-i} \\ =&\int\limits_{x_i^0}^{x_i}\int\limits_{\Re^{I-1}} \pi(x_i, x_{-i})\frac{\partial}{\partial t}\{u_i(t, x_{-i})f(x_{-i}|t)\}dx_{-i}dt \\ \geq&\int\limits_{x_i^0}^{x_i}\int\limits_{\Re^{I-1}} \pi(t, x_{-i})\frac{\partial}{\partial t}\{u_i(t, x_{-i})f(x_{-i}|t)\}dx_{-i}dt \\ \geq&\int\limits_{x_i^0}^{x_i}\int\limits_{\Re^{I-1}} \pi(x_i^0, x_{-i})\frac{\partial}{\partial t}\{u_i(t, x_{-i})f(x_{-i}|t)\}dx_{-i}dt \\ =&\int\limits_{\Re^{I-1}} \pi(x_i^0, x_{-i})(u_i(x_i, x_{-i})f(x_{-i}|x_i) - u_i(x_i^0, x_{-i})f(x_{-i}|x_i^0))dx_{-i}.\end{aligned}$$

The first and last expressions in the last series of inequalities are bounded, and the middle expression is the final term in $\hat{c}_i(x_i)$. Thus, all the terms in $\hat{c}_i(x_i)$ are bounded, and $\hat{c}(x_i)$ is a bounded function. Therefore, it has an infimum $\underline{c}_i$. Let the money burning function in the constructed mechanism $c_i(x_i)$ satisfy $c_i(x_i) = \hat{c}_i(x_i) - \underline{c}_i$. This $c_i(x_i)$ necessarily takes on only nonnegative values. Furthermore, for any types x_i and x_i', we have

$$\begin{aligned}
&\int\limits_{\Re^{I-1}} \pi(x_i, x_{-i})u_i(x_i, x_{-i})f(x_{-i}|x_i)dx_{-i} - c_i(x_i) \\
&= \int\limits_{\Re^{I-1}} \pi(x_i^0, x_{-i})u_i(x_i^0, x_{-i})f(x_{-i}|x_i^0)dx_{-i} + \int\limits_{x_i^0}^{x_i} \int\limits_{\Re^{I-1}} \pi(t, x_{-i})\frac{\partial}{\partial t}\{u_i(t, x_{-i})f(x_{-i}|t)\}dx_{-i}dt + \underline{c}_i \\
&\geq \int\limits_{\Re^{I-1}} \pi(x_i^0, x_{-i})u_i(x_i^0, x_{-i})f(x_{-i}|x_i^0)dx_{-i} + \int\limits_{x_i^0}^{x_i'} \int\limits_{\Re^{I-1}} \pi(t, x_{-i})\frac{\partial}{\partial t}\{u_i(t, x_{-i})f(x_{-i}|t)\}dx_{-i}dt \\
&\quad + \int\limits_{x_i'}^{x_i} \int\limits_{\Re^{I-1}} \pi(x_i', x_{-i})\frac{\partial}{\partial t}\{u_i(t, x_{-i})f(x_{-i}|t)\}dx_{-i}dt + \underline{c}_i \\
&= \int\limits_{\Re^{I-1}} \pi(x_i', x_{-i})u_i(x_i, x_{-i})f(x_{-i}|x_i)dx_{-i} - c_i(x_i').
\end{aligned}$$

Thus, every type x_i prefers to announce her true type and burn $c_i(x_i)$ than to misregister as a different type.

The only remaining feature to construct is how the mechanism responds to null reports. Because equilibria in which no agent makes a null report are sought, the specification for $\pi(\cdot)$ when there are multiple null reports is irrelevant. However, $\pi(\cdot)$ must have the property that no type of agent wishes to make a null report unilaterally. Denote the message profile in which agent i makes the null report and other agents report vector of types x_{-i} by $(m_\emptyset, x_{-i})$. Consider arbitrary agent i and the constructed function $c_i(\cdot)$. It is known that the infimum of $c_i(\cdot)$ is 0. If there is a type $\tilde{x}_i$ such that $c_i(\tilde{x}_i) = 0$, then assume that $\pi(m_\emptyset, x_{-i}) = \pi(\tilde{x}_i, x_{-i})$ for all x_{-i}; then by incentive compatibility, no type of agent will have an incentive to make the null report, as no type of agent except $\tilde{x}_i$ has an incentive to report $\tilde{x}_i$ and burn no money. However, it is possible that there is no type x_i such that $c_i(x_i) = 0$. In this case, we construct $\pi(m_\emptyset, x_{-i})$ as follows. Because the infimum of $c_i(\cdot)$ is 0, there necessarily exists a sequence of types of agent i, $(x_i^t)_{t=1}^\infty$, such that $\lim_{t\to\infty} c_i(x_i^t) = 0$, and such that $(x_i^t)_{t=1}^\infty$ is either increasing or decreasing. By the monotonicity of $\pi(\cdot)$, as well as the fact that $\pi(\cdot)$ is bounded between 0 and 1, the limit $\lim_{t\to\infty} \pi(x_i^t, x_{-i})$ therefore exists for every x_{-i}; let $\pi(m_\emptyset, x_{-i}) = \lim_{t\to\infty} \pi(x_i^t, x_{-i})$. Choose any type x_i of agent i; by incentive compatibility of $\pi(\cdot)$ for type reports, we have

$$\int\limits_{\Re^{I-1}} \pi(x_i, x_{-i})u_i(x_i, x_{-i})f(x_{-i}|x_i)dx_{-i} - c_i(x_i) \geq \int\limits_{\Re^{I-1}} \pi(x_i^t, x_{-i})u_i(x_i, x_{-i})f(x_{-i}|x_i)dx_{-i} - c_i(x_i^t),$$

for all t. Therefore, we have

$$\begin{aligned}
&\int_{\Re^{I-1}} \pi(x_i, x_{-i}) u_i(x_i, x_{-i}) f(x_{-i}|x_i) dx_{-i} - c_i(x_i) \\
\geq &\int_{\Re^{I-1}} \lim_{t\to\infty} \pi(x_i^t, x_{-i}) u_i(x_i, x_{-i}) f(x_{-i}|x_i) dx_{-i} - \lim_{t\to\infty} c_i(x_i^t) \\
= &\int_{\Re^{I-1}} \pi(m_\emptyset, x_{-i}) u_i(x_i, x_{-i}) f(x_{-i}|x_i) dx_{-i}.
\end{aligned}$$

Thus, no type has an incentive to report $m_\emptyset$ and burn no money. ∎

Proof of Proposition 2. Fix a pre-replication environment, and any sequence of choice rules $\pi^T(\cdot)$ and money burning functions $\hat{c}_i^T(\cdot)$ such that, for any replication T', $\pi^{T'}(\cdot)$ is anonymous up to public type and implementable with money burning function $\hat{c}_i^{T'}(\cdot)$ for the T'-fold replication. Given $\hat{c}_i^{T'}(\cdot)$, denote the minimum of the ex-ante support of money burnt by agent i as $\underline{c}_i^{T'}$. If $\underline{c}_i^{T'} > 0$, modify i's money burning functions to $c_i^{T'}(x_i) = \hat{c}_i^{T'}(x_i) - \underline{c}_i^{T'}$. This makes the minimum of the ex-ante support of money burnt for every agent 0, without altering the performance of the mechanism. With this feature, for every agent i there exists at least one type $\underline{x}_i^{T'}$ such that $c_i^{T'}(\underline{x}_i) = 0$.

Fix any signal type x_i of any agent i. By incentive compatibility of the given mechanism with money burning, and the specification of $\underline{x}_i^{T'}$, we have

$$\begin{aligned}
c_i^{T'}(x_i) &\leq \sum_{\omega\in\Omega} p(\omega|x_i) \sum_{x_{-i}} (\pi^{T'}(x_i, x_{-i}) - \pi^{T'}(\underline{x}_i^{T'}, x_{-i})) \tilde{u}_i(x_i, \omega) f^{T'}(x_{-i}|\omega) \\
&= \sum_{\omega\in\Omega} p(\omega|x_i) \tilde{u}_i(x_i, \omega) \sum_{x_{-i}} (\pi^{T'}(x_i, x_{-i}) - \pi^{T'}(\underline{x}_i^{T'}, x_{-i})) f^{T'}(x_{-i}|\omega).
\end{aligned} \tag{A}$$

By symmetry, the function $\pi^{T'}(\cdot)$ may only depend on the number of agents of each public type who have a given private signal. If we assume without loss of generality that the support of private types for agents with different public types are disjoint (i.e., two agents with different public types necessarily have different private types), then we may substitute for $\pi^{T'}(\cdot)$ the function $\tilde{\pi}^{T'}((n(\hat{x}))_{\hat{x}\in\cup_i X_i})$, where $n(\hat{x})$ is the number of agents with private type $\hat{x}$. In what follows it will be convenient to write a profile of private types for the T'-fold replication as $(n(x_i), \tilde{n}^{T'})$, where $\tilde{n}^{T'}$ is the profile of the number of agents who have each type other than x_i and $\underline{x}_i^{T'}$. Because the total number of agents is fixed at $T'N$, the number of agents with type $\underline{x}_i^{T'}$ is determined by $(n(x_i), \tilde{n}^{T'})$, and equal to $T'N - n(x_i) - \sum_{\hat{x}\neq x_i, \underline{x}_i^{T'}} n(\hat{x})$. To complete some simplifying notation, let $Y(\tilde{n}^{T'}) \equiv T'n - \sum_{\hat{x}\neq x_i, \underline{x}_i^{T'}} n(\hat{x})$, i.e., the number of agents having type x_i or $\underline{x}_i^{T'}$, and let $f(\tilde{n}^{T'}|\omega)$

be the marginal probability of realizing a profile with distribution $\tilde{n}^{T'}$ of types other than x_i or $\underline{x}_i^{T'}$, conditional on ω.

Fix a state ω, and consider the term $\sum_{x_{-i}} (\pi^{T'}(x_i, x_{-i}) - \pi^{T'}(\underline{x}_i^{T'}, x_{-i}))f^{T'}(x_{-i}|\omega)$. Using the notation defined above, this may be rewritten

$$\sum_{\tilde{n}^{T'}} f(\tilde{n}^{T'}|\omega) \sum_{j=0}^{Y(\tilde{n}^{T'})-1} [\hat{\pi}^{T'}(j+1, \tilde{n}^{t'}) - \hat{\pi}^{T'}(j, \tilde{n}^{t'})] \frac{(Y(\tilde{n}^{T'})-1)!}{j!(Y(\tilde{n}^{T'})-1-j)!} f(x_i|\omega)^j f(\underline{x}_i^{T'}|\omega)^{Y(\tilde{n}^{T'})-1-j}.$$

Suppose that $x_i > \underline{x}_i^{T'}$ (a parallel argument holds if the opposite is true). Then, since we are considering only monotonic choice rules, for each $\tilde{n}^{T'}$ and each $j \in \{0, 1, \ldots, Y(\tilde{n}^{T'}) - 1\}$, we have $\hat{\pi}^{T'}(j+1, \tilde{n}^{t'}) - \hat{\pi}^{T'}(j, \tilde{n}^{t'}) \geq 0$. Furthermore, $\sum_{j=0}^{Y(\tilde{n}^{T'})-1} [\hat{\pi}^{T'}(j+1, \tilde{n}^{t'}) - \hat{\pi}^{T'}(j, \tilde{n}^{t'})] = \hat{\pi}^{T'}(Y(\tilde{n}^{T'})-1, \tilde{n}^{t'}) - \hat{\pi}^{T'}(0, \tilde{n}^{t'}) \leq 1$. This implies that

$$\begin{aligned}
&\sum_{j=0}^{Y(\tilde{n}^{T'})-1} [\hat{\pi}^{T'}(j+1, \tilde{n}^{t'}) - \hat{\pi}^{T'}(j, \tilde{n}^{t'})] \frac{(Y(\tilde{n}^{T'})-1)!}{j!(Y(\tilde{n}^{T'})-1-j)!} f(x_i|\omega)^j f(\underline{x}_i^{T'}|\omega)^{Y(\tilde{n}^{T'})-1-j} \\
&\leq \max_{j \in \{0,1\ldots,Y(\tilde{n}^{T'})-1\}} \frac{(Y(\tilde{n}^{T'})-1)!}{j!(Y(\tilde{n}^{T'})-1-j)!} f(x_i|\omega)^j f(\underline{x}_i^{T'}|\omega)^{Y(\tilde{n}^{T'})-1-j}.
\end{aligned}$$

Choose any number $M > \max\{0, \tilde{u}_i(x_i, \omega)\}$ for all ω. Note that

$$\max_{j \in \{0,1\ldots,Y(\tilde{n}^{T'})-1\}} \frac{(Y(\tilde{n}^{T'})-1)!}{j!(Y(\tilde{n}^{T'})-1-j)!} f(x_i|\omega)^j f(\underline{x}_i^{T'}|\omega)^{Y(\tilde{n}^{T'})-1-j}$$

is the maximum probability on any realized draw of a binary distribution (times a conditioning scalar). As $Y(\tilde{n}^{T'})$ approaches infinity, this maximum probability approaches 0; in particular, it is possible to identify a number $\overline{Y}$ such that the maximum is less than $1-\sqrt{1-\frac{\epsilon}{M}}$ for all $Y(\tilde{n}^{T'}) > \overline{Y}$. Furthermore, because replication acts on every public type, the probability that the proportion of agents of public type θ_i who have realized private type x_i is very different from $f(x_i|\omega, \theta_i)$ vanishes; in particular, there exists $\overline{T}$ such that for $\hat{T} > \overline{T}$, the probability that fewer than $\overline{Y}$ agents have realized private type x_i or $\underline{x}_i^{T'}$ is less than $1 - \sqrt{1-\frac{\epsilon}{M}}$. For $\hat{T} > \overline{T}$, this leads to the following chain of inequalities, beginning with (A) from above:

$$\begin{aligned}
c_i^{T'}(x_i) &\leq \sum_{\omega \in \Omega} p(\omega|x_i)\tilde{u}_i(x_i, \omega) \sum_{x_{-i}} (\pi^{T'}(x_i, x_{-i}) - \pi^{T'}(\underline{x}_i^{T'}, x_{-i}))f^{T'}(x_{-i}|\omega). \\
&\leq M \sum_{\tilde{n}^{T'}} f(\tilde{n}^{T'}|\omega) \max_{j \in \{0,1\ldots,Y(\tilde{n}^{T'})-1\}} \frac{(Y(\tilde{n}^{T'})-1)!}{j!(Y(\tilde{n}^{T'})-1-j)!} f(x_i|\omega)^j f(\underline{x}_i^{T'}|\omega)^{Y(\tilde{n}^{T'})-1-j} \\
&\leq M[(1-\sqrt{1-\frac{\epsilon}{m}})(1) + (\sqrt{1-\frac{\epsilon}{m}})(1-\sqrt{1-\frac{\epsilon}{m}})] \\
&= \epsilon.
\end{aligned}$$

REFERENCES

Becker, Gary S, A Theory of Competition among Pressure Groups for Political Influence, *Quarterly Journal of Economics* **98** (1983), pp. 371-400.

Bernheim, B. Douglas and Whinston, Michael D, Menu Auctions, Resource Allocation and Economic Influence, *Quarterly Journal of Economics* **101** (1986), pp. 1-31.

Casella, Alessandra, Storable Votes, *Games and Economic Behavior* **51** (2005), pp. 391-419.

Casella, Alessandra and Gelman, Andrew, A Simple Scheme to Improve the Efficiency of Referenda, *Journal of Public Economics* **92** (2008), pp. 2240-2261.

Feddersen, Timothy and Pesendorfer, Wolfgang, Voting Behavior and Information Aggregation in Elections with Private Information, *Econometrica* **65** (1997), pp. 1029-1058.

Grossman, Gene M. and Helpman, Elhanan, Protection for Sale, *American Economic Review* **84** (1994), pp. 833-850.

Jackson, Matthew and Sonnenschein, Hugo, Overcoming Incentive Constraints by Linking Decisions, *Econometrica* **75** (2007), pp. 241-257.

McLean, Richard P. and Postlewaite, Andrew, Informational Size and Incentive Compatibility, *Econometrica* **70** (2002), pp. 2421-2454.

McLean, Richard P. and Postlewaite, Andrew, Informational Size and Efficient Auctions, *Review of Economic Studies* **71** (2004), pp. 809-827.

Tullock, Gordon , The Welfare Costs of Tariffs, Monopolies and Theft , *Western Economic Journal* **5** (1967), pp. 224-232.