

Cost Benefit Analyses versus Referenda

Martin J. Osborne and Matthew A. Turner

University of Toronto

We consider a planner who chooses between two public policies and ask whether a referendum or a cost benefit analysis leads to higher welfare. We find that a referendum leads to higher welfare than a cost benefit analysis in a “common value” environment. Cost benefit analysis is better in a “private value” environment.

I. Introduction

Both referenda and cost benefit analyses are widely used to choose public policies. Moreover, many public policy decisions resolved by cost benefit analyses could alternatively be resolved by referenda, and conversely. Given the prevalence and apparent interchangeability of cost benefit analyses and referenda, it is natural to ask when one method leads to a better decision than the other. We seek to answer this question.

We consider public policy decisions that may have both “private value” and “common value” components. The agents in a private value policy decision, like those in a private value auction, understand and are familiar with the outcome of the policy but assign different values to this outcome. Policies that principally determine the distribution of familiar goods or services fall into this class. The information structure of a common value policy decision is the same as the information structure of a common value auction. That is, some agents are unfamiliar with, or uncertain about, the consequences of the policy, but all agents agree on what constitutes a good outcome. If the policy is implemented, the

We are grateful to Maureen Cropper, Gilles Duranton, Nicholas Flores, Nat Keohane, Stephan Schott, the editors, and two referees for comments that substantially improved the paper and to the Social Sciences and Humanities Research Council of Canada for financial support.

[*Journal of Political Economy*, 2010, vol. 118, no. 1]
© 2010 by The University of Chicago. All rights reserved. 0022-3808/2010/11801-0005\$10.00

agents will know how they value it, but at the point of making a decision, they do not. We argue (in Sec. V.B) that common value policy decisions are pervasive and that much of the literature on cost benefit analysis is concerned with exactly such decisions.

We restrict attention to binary policy choices and compare the outcome of a referendum with that of a cost benefit analysis. In a cost benefit analysis, the policy maker first elicits a value for each policy from each agent and then selects a policy by summing the agents' reports and choosing the policy that generates the largest sum of benefits net of costs. While we discuss the implications of our work for other methods (see Sec. V.D), our model fits most closely the "stated preference" method of cost benefit analysis, in which the policy maker elicits agents' valuations by asking the agents to report them. In a referendum, each agent either votes for one of the policies or abstains, and the policy that receives the most votes is implemented.

Most cost benefit analyses maintain the assumption that each subject reveals her preferences truthfully. Under this assumption and the assumption that individuals vote strategically in a referendum, cost benefit analyses and referenda differ in two significant respects. First, a cost benefit analysis may be better at eliciting cardinal information about preferences than a referendum. Referenda allow voters to reveal only ordinal information, whereas cost benefit analyses, with their larger report spaces, allow the communication of cardinal information. We show that, as a consequence, policy decisions obtained from cost benefit analyses result in higher welfare than policy decisions obtained from referenda when agents know the values they assign to the policy and these values vary across agents. That is, cost benefit analyses lead to better outcomes than do referenda for private value policy decisions.

The second significant respect in which cost benefit analyses and referenda differ is that the "swing voter's curse" (Feddersen and Pesendorfer 1996) may lead uninformed agents to abstain from a referendum. Such abstention allows uninformed voters to delegate to like-minded but informed voters. Thus, the well-informed agents are oversampled by a referendum, whereas no such self-selection occurs in a cost benefit analysis. This implies that the outcome of a referendum is superior to that of a cost benefit analysis when individuals have similar preferences but different information—that is, for a common value policy decision.

In summary, we show that the outcome of a cost benefit analysis is superior when individuals have diverse preferences but similar information, whereas the outcome of a referendum is superior when individuals have similar preferences but different degrees of uncertainty. More succinctly, a cost benefit analysis is better in a private value environment, but a referendum is better in a common value environment.

This strong result hinges on agents behaving differently in the two

mechanisms. Specifically, it requires strategic voting in referenda but truthful revelation of preferences in cost benefit analyses. The assumption of truthful revelation is easy to defend for a cost benefit analysis based on revealed preference information such as housing prices or travel costs but may be questioned for a cost benefit analysis using stated preference methods. We find that if agents approach cost benefit analyses with the same sophistication that they apply to referenda, the two mechanisms are equivalent.

Even without taking a position on the extent to which agents behave strategically when participating in a cost benefit analysis, we can make a strong statement about the comparative advantages of cost benefit analyses and referenda: a cost benefit analysis is at least as good as a referendum in a private value environment, whereas the converse is true in a common value environment. For reasons we discuss in Section V.C, we are inclined to give some credence to the widely maintained assumption of truthful revelation in cost benefit analyses. In this case, cost benefit analyses are strictly preferred to referenda in private value environments, and conversely in common value environments.

After discussing the background for the problem in Section II, we describe the model in Section III and analyze it in Section IV. Section V contains detailed discussion.

II. Background: Cost Benefit Analyses and Referenda

Executive Order 12866 of the U.S. government requires that “each agency shall assess both the costs and the benefits of . . . regulation” (Clinton 1993, sec. 1.b.6). Similarly, the U.S. Office of Management and Budget is required to submit to Congress “an estimate of the costs and benefits of Federal rules and paperwork” each year (FY2001 Treasury and Government Appropriations Act, sec. 624.a.1). Several major U.S. regulations also mandate measurement of the benefits and costs of regulation. In addition, as of 1996, 10 states required an analysis of the benefits and costs of regulation (Hahn 2000).

Given these directives, it is unsurprising that cost benefit analyses are often used to evaluate public policy and that these analyses appear to influence regulators in favor of policies that generate higher estimated benefits at lower estimated costs. For example, Cropper et al. (1992) provide evidence that the U.S. Environmental Protection Agency’s decisions to regulate dangerous pesticides are influenced by estimates of the costs and benefits of the pesticide in question. Smith (2000) provides anecdotal evidence that cost benefit analysis has affected air quality regulation for the Grand Canyon. Viscusi (1996) argues that the U.S. Department of Transportation began to pursue regulations with larger

estimated net benefits after it incorporated cost benefit analysis into its decision process.

This evidence suggests that many important allocation decisions are made roughly according to a “cost benefit decision rule” that operates in two steps: measure the cost and benefit of a proposed action and then choose the action if and only if the net benefit is positive. Indeed, this stylized decision rule is broadly consistent with the injunction of Executive Order 12866 to “propose or adopt a regulation only upon a reasoned determination that the benefits of the intended regulation justify its costs” (Clinton 1993, sec. 1.b.6), with similar mandates present in many state laws (Hahn 2000) and with the exhortations of professional economists (e.g., Arrow et al. 1996).¹

Estimating the costs and benefits of public policies is not trivial. But economists have responded to the problem with considerable ingenuity, and many techniques are now available. Of these, the three most common are stated preference methods, travel cost methods, and hedonic analysis. We develop a stylized description of stated preference cost benefit analysis and discuss later (in Sec. V.D) the extent to which our intuition applies to the travel cost and hedonic methods.

The stated preference method draws a sample from the affected population and asks each respondent to reveal information about the benefit he or she would derive from a particular policy. The concept is simple, but the method often involves sophisticated survey techniques. For example, stated preference surveys often describe policies and their consequences in detail, elicit demographic information, debrief respondents to assess their understanding of the questions (Arrow et al. 1993; Hanemann 1994), and even allow for the possibility that agents’ responses violate the axioms of revealed preference (Cherry, Crocker, and Shogren 2003; Settle, Cherry, and Shogren 2003).

Referenda are another important mechanism for collective decision making. Aside from their pervasive use in choosing government officials, they are also widely used to resolve questions that might otherwise be left to regulators.² For example, in California alone, 2004 saw some 16 state-level referendum measures on topics ranging from health care to gambling to criminal law (State of California 2004). All of these decisions could have been made on the basis of cost benefit analyses.³

¹ Each of these sources also allows for the possibility that factors other than the costs and benefits of a policy, e.g., distributional implications, should influence the chosen policy.

² Lupia and Matsusaka (2004) report that in 2004 more than 70 percent of the U.S. population lived in jurisdictions in which referenda were used in this way. They also report widespread use of referenda in Europe.

³ Choosing government officials by cost benefit analysis may seem odd, but it is almost certainly practical. If we identify a candidate for office with a bundle of policies and attributes, the choice problem could then be resolved by a cost benefit analysis.

III. Model: A Simple Public Policy Decision

We consider a planner who must choose one of two policies. She faces an environment in which there are n agents ($n \geq 3$) and two states of the world. Each agent is one of four possible types. We use the following notation:

States: $S = \{0, 1\}$,

Policies: $X = \{0, 1\}$,

Agent types: $T = \{0, 1, i, u\}$.

The number of agents of each type t is n_t . Agents of types 0 and 1 are *partisans* who respectively prefer policies 0 and 1 in both states of the world. Agents of types i and u are *independents* who prefer the policy to match the state (i.e., policy 0 in state 0, policy 1 in state 1). Agents of type i are informed: they know the state before experiencing the consequences of the policy. Agents of type u are uninformed: they do not learn the state until they experience the consequences of the policy choice. All agents are expected payoff maximizers.

The planner and the uninformed independents believe that the probability of state 0 is α , with $\alpha < \frac{1}{2}$. (Because the partisans' preferences do not depend on the state, a specification of their information is unnecessary.)

We denote by $v^j(x, s)$ the payoff of agent j for policy x in state s . We assume that these payoffs take the following forms, where π_0 and π_1 are positive constants:

$$\text{0-partisans: } v^j(x, s) = \begin{cases} \pi_0 & \text{if } x = 0 \\ 0 & \text{if } x = 1 \end{cases} \quad \text{for } s = 0, 1,$$

$$\text{1-partisans: } v^j(x, s) = \begin{cases} 0 & \text{if } x = 0 \\ \pi_1 & \text{if } x = 1 \end{cases} \quad \text{for } s = 0, 1,$$

$$\text{independents: } v^j(x, s) = \begin{cases} 1 & \text{if } x = s \\ 0 & \text{if } x \neq s. \end{cases}$$

This specification gives each agent's payoff net of any cost she bears. This formulation simplifies the exposition by relieving us of separately accounting for costs.

It is sometimes illuminating to consider two simple polar cases. In the *pure private values* case, all agents are partisans ($n_i = n_u = 0$). In this case, agents disagree about the best policy because their tastes differ. In the *pure common values* case, all agents are independents ($n_0 = n_1 = 0$). In this case, agents are in perfect agreement about the best policy conditional on the state, but some agents do not know the state.

A. *Welfare Maximization*

We measure welfare by summing the payoffs of all agents. Thus welfare for policy x in state s is

$$W(x, s) = \sum_{j=1}^n v^j(x, s).$$

When the state is s , the planner would like to choose the policy x that maximizes $W(x, s)$. Given the forms we assume for the payoffs, policy 0 is welfare-maximizing in state 0 if and only if $\pi_0 n_0 + n_i + n_u \geq \pi_1 n_1$ or if and only if

$$\pi_0 n_0 - \pi_1 n_1 + n_i + n_u \geq 0. \tag{1}$$

Similarly, policy 0 is welfare-maximizing in state 1 if and only if

$$\pi_0 n_0 - \pi_1 n_1 - n_i - n_u \geq 0. \tag{2}$$

The welfare-maximizing policy in state 0, as a function of the number of agents of each type, is illustrated in figure 1 for a case in which 1-partisans feel more strongly about the policy than 0-partisans do—that is, $\pi_1 > \pi_0$.

The planner's problem is to choose a policy based on information that she obtains from the agents. We consider two means for making such a choice: a cost-benefit analysis and a referendum.

We note that for the pure common values case, delegating the choice of policy to an informed agent results in a welfare-maximizing policy. The difficulty with this method is that it is not useful in environments in which there are even a small number of partisans. Such partisans




FIG. 1.—An example of the dependence of the welfare-maximizing policy in state 0 on the number of agents of each type. In the example, $\pi_0 < \pi_1$. The welfare-maximizing policy is policy 0 below the solid line and policy 1 above this line.

have no incentive to reveal themselves, so it is impossible for a planner to reliably identify an informed agent to whom to delegate the decision.⁴

B. Cost Benefit Analysis

In a cost benefit analysis, each survey respondent reports her estimate of the difference between her valuations of the two policies—that is, her “willingness to pay” to switch from policy 0 to policy 1. The policy maker sums these reports and chooses policy 1 if the sum is positive, policy 0 if the sum is negative, and each policy with probability one-half if the sum is zero.⁵

As with many aspects of stated preference methods, the extent to which agents respond to survey questions strategically is contentious. To avoid taking a position on this issue, we conduct our analysis under two competing assumptions. In the first, each agent reports (the expectation of) her true valuation (she acts “sincerely”). In the second, each agent chooses her report strategically.

In our model of cost benefit analysis with sincere reporting, each 0-partisan reports $-\pi_0$ and each 1-partisan reports π_1 . Each informed independent reports -1 in state 0 and 1 in state 1. Each uninformed independent reports the expected difference between her valuations of the two policies, given her belief that the probability of state 0 is α . That is, each uninformed independent reports $\alpha(0 - 1) + (1 - \alpha)(1 - 0) = 1 - 2\alpha (> 0)$.

In our model of cost benefit analysis with strategic reporting, each agent is free to choose any report. We restrict the reports to be bounded and allow nonparticipation. Precisely, for some $B > 0$ we require each agent either to report a number in the interval $[-B, B]$ or not to participate; we denote nonparticipation by ϕ . That is, the set of actions of each agent is $[-B, B] \cup \{\phi\}$.⁶ By bounding the set of permissible reports, we are implicitly assuming that the survey administrator has a

⁴ This problem of partisans masquerading as informed independents is essentially the same problem as the one studied in Banerjee and Somanathan (2001).

⁵ Stated preferences surveys are generally administered to only a subset of the affected population and are therefore subject to sampling error. An analysis of these errors is tangential to our inquiry. To abstract from this problem we assume that the policy maker samples the whole population.

As a referee suggests, other formalizations of cost benefit analysis are possible. In particular, one might imagine that the planner retains some of the rent or bases her decision on the median report. While interesting, these rules appear to correspond less closely to the cost benefit analyses described in Sec. II than the one we study.

⁶ Our assumption that an agent who participates reports a real number appears to be at variance with the accepted best practice in stated preference methods, which calls for the elicitation of bounds for an individual’s valuation (Hanemann 1994). Our assumption allows cardinal information to be revealed more precisely.

prior about the range of possible values and either rejects or truncates responses that are unreasonably large or unreasonably small.⁷

Because the policy maker's decision is based on the sum of the reports, nonparticipation (the action ϕ) is equivalent to a report of 0. Thus, to simplify the notation, we restrict the set of allowable reports to $[-B, B]$ without loss of generality. These assumptions define a Bayesian game with two states. Each informed independent knows the state; each uninformed independent does not. Partisans' payoffs are independent of the state, so it does not matter whether they know the state. A strategy for an informed independent is a pair of reports, one for each state. A strategy for every other agent is simply a report. Our solution concept is a variant of the standard notion of Nash equilibrium, which we explain in Section IV.

C. *Referendum*

In a referendum, each agent chooses whether to vote and if so for which policy. The policy that receives the most votes is selected. In the event of a tie, each policy is selected with probability one-half.

Our formal model of a referendum, like our model of cost benefit analysis with strategic reporting, is a Bayesian game. The models differ only in that for a referendum each agent's set of actions is $\{-1, 0, 1\}$ rather than $[-B, B]$, where the action 1 is a vote for policy 1, the action -1 is a vote for policy 0, and the action 0 is nonparticipation. We use the same notion of equilibrium as we do for our model of a cost benefit analysis with strategic reporting and assume, as before, that the policy maker chooses policy 1 if the sum of the agents' actions is positive, policy 0 if it is negative, and each policy with probability one-half if it is zero.

IV. Results

A. *Cost Benefit Analysis with Sincere Reporting*

When agents report sincerely, the sum of the reported values of changing from policy 0 to policy 1 is

$$\begin{cases} -\pi_0 n_0 + \pi_1 n_1 - n_i + (1 - 2\alpha)n_u & \text{if } s = 0 \\ -\pi_0 n_0 + \pi_1 n_1 + n_i + (1 - 2\alpha)n_u & \text{if } s = 1. \end{cases} \quad (3)$$

We wish to understand the conditions under which it is welfare-maximizing to choose policy 1 when this sum is positive and policy 0 when it is negative.

⁷ One of the functions of the "closed-ended questions" discussed in Hanemann (1994) is to impose such bounds.

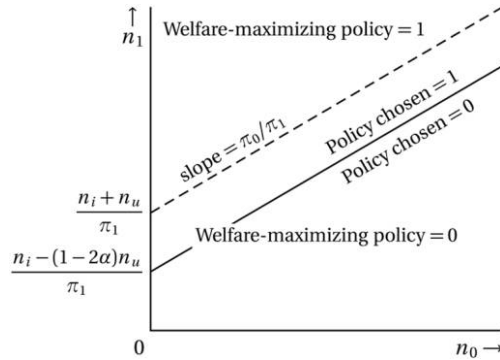


FIG. 2.—An example of the difference between the welfare-maximizing policy in state 0 and the policy chosen in state 0 under cost benefit analysis with sincere reporting. In the example, $\pi_0 < \pi_1$ and $n_i > (1 - 2\alpha)n_u$. The welfare-maximizing policy is policy 0 below the dashed line and policy 1 above it.

First consider a pure common value problem ($n_0 = n_1 = 0$). In state 1 each informed agent reports the true change in her valuation from moving from policy 0 to policy 1, namely 1, and each uninformed agent reports the expected value of this change, namely $1 - 2\alpha > 0$. Since the reports of agents of both types are qualitatively correct, the sum of the reports is also qualitatively correct and leads to policy 1, the welfare-maximizing policy choice.

In state 0, informed agents make qualitatively correct negative reports. However, uninformed agents make the same positive reports as they do in state 1, and these reports are now qualitatively incorrect. Thus if uninformed agents are sufficiently numerous, the sum of reports is positive. In this case, the cost benefit decision rule leads to policy 1, whereas policy 0 maximizes welfare. The condition for this erroneous decision is $n_i < (1 - 2\alpha)n_u$.

Now consider a pure private value problem ($n_i = n_u = 0$). In both states the sum of the reports is $-\pi_0 n_0 + \pi_1 n_1$, the actual welfare gain. In this case, the cost benefit analysis decision rule coincides exactly with the welfare-maximizing decision rule and thus always selects the welfare-maximizing policy.

The intuition suggested by these two polar cases generalizes naturally to environments in which both partisans and independents are present. The more uninformed independents there are, the wider the range of circumstances under which cost benefit analysis leads to an incorrect decision. Figure 2 illustrates the intuition for an example. It shows the policy chosen in state 0 as a function of the numbers of agents of each type. In the middle region, the policy chosen by cost benefit analysis with sincere reporting is not welfare-maximizing. The size of this region

is increasing in n_u . In this region, the reports of the uninformed independents, each of whom submits the positive report $1 - 2\alpha$, cause the policy maker to choose policy 1 even though policy 0 is welfare-maximizing.

The following result gives the condition for cost benefit analysis with sincere reporting to yield the welfare-maximizing policy in the general case and in the two polar cases of pure common values and pure private values. A proof is given in the Appendix.

PROPOSITION 1 (Cost benefit analysis with sincere reporting). Cost benefit analysis with sincere reporting selects the welfare-maximizing policy in state 0 if and only if either $\pi_1 n_1 - \pi_0 n_0 - n_i + (1 - 2\alpha)n_u < 0$ or $\pi_1 n_1 - \pi_0 n_0 - n_i - n_u \geq 0$. It selects the welfare-maximizing policy in state 1 if and only if either $\pi_1 n_1 - \pi_0 n_0 - n_i - (1 - 2\alpha)n_u < 0$ or $\pi_1 n_1 - \pi_0 n_0 - n_i - n_u \geq 0$. In particular, for the two polar cases, we have the following results.

- a. (Pure common values) If $n_0 = n_1 = 0$, then the welfare-maximizing policy is chosen in state 1 (the more likely state) and is chosen in state 0 if $n_i > (1 - 2\alpha)n_u$. The wrong policy is chosen in state 0 if $n_i < (1 - 2\alpha)n_u$.
- b. (Pure private values) If $n_i = n_u = 0$, then the welfare-maximizing policy is always chosen.

B. Cost Benefit Analysis with Strategic Reporting

Equilibrium notion.—Our cost benefit analysis game has many Nash equilibria. For example, it has a Nash equilibrium in which all agents report B and another equilibrium in which all agents report $-B$. In each case, no deviation by any agent affects the outcome (given that there are at least three agents). These equilibria are unappealing because some agents' actions are weakly dominated. That is, these agents have actions that are at least as attractive as their equilibrium actions regardless of the other agents' actions and *more* attractive for some actions of the other agents. For example, for the equilibrium in which every agent reports B , an informed independent who knows the state is 0 and switches to a report of $-B$ is better off if every other agent deviates to nonparticipation and is no worse off for any other deviation.

Other less extreme equilibria are unappealing even though no agent's action is dominated. Suppose, for example, that there are no partisans and that the number of uninformed independents exceeds the number of informed independents by three. Consider the strategy profile in which all informed independents report $-B$ in state 0 and B in state 1 and all the uninformed independents report B . This strategy profile is a Nash equilibrium, with the outcome policy 1 in both states, because

no change in any agent's action affects the outcome in either state. However, a deviation by an uninformed independent, say j , to $-B$ is attractive. The reason is that if exactly one of the other uninformed independents fails to follow her strategy and instead does not participate, then j 's deviation changes the outcome to policy 0 in state 0, whereas for any other failure by a single agent to follow her strategy, j 's deviation does not affect the outcome in either state. If *all* the other agents fail to follow their strategies and choose nonparticipation, then j 's deviation leads to a worse outcome (policy 0 in both states), but j should plausibly regard such a departure from the strategies of the other agents as much less likely than a deviation by a single agent.

To formulate the idea precisely, we focus on deviations by the other agents to nonparticipation (which are key in the arguments for the examples we have just described). Suppose that when choosing an action, each agent considers the possibility that each of the other agents may exogenously be prevented from participating (e.g., because a phone rings, a child cries, or a doorbell breaks). Specifically, suppose that each agent assumes that every other agent will, with small probability, independently be prevented from participating. Then, when choosing a strategy, each agent first limits herself to strategies that are optimal when all the other agents adhere to their strategies and then, within this set, chooses on the basis of the performance of the strategies when some agents do not participate. Because the probability of a small number of nonparticipants is much higher than the probability of a large number, each agent gives most weight in her strategy choice to situations in which the number of nonparticipants is small. But if two strategies perform equally well when the number of nonparticipants is small, she compares the strategies in the case in which the number of nonparticipants is large.

More precisely, define the ϵ -perturbation of a strategy profile σ to be the strategy profile in which each player j chooses σ_j with probability $1 - \epsilon$ and nonparticipation with probability ϵ .⁸

DEFINITION 1. A strategy profile σ is an *equilibrium* if there exists $\bar{\epsilon} > 0$ such that, for all $\epsilon < \bar{\epsilon}$, the strategy σ_j of each agent j is a best response to the ϵ -perturbation of σ .

The Nash equilibrium that we consider above in which every agent reports B is not an equilibrium in this sense. When all agents report B ,

⁸ This definition is similar in spirit to the definition of a trembling hand perfect equilibrium of the strategic game in which each type of each player in the Bayesian game is a different player. It differs in that only perturbations to nonparticipation, rather than arbitrary perturbations, are considered, and the probability ϵ is assumed to be the same for all agents. It is closely related also to the assumption of Feddersen and Pesendorfer (1996) that the number of agents is random.

a deviation by an agent affects the outcome only if the number of other participants is zero or one. In both cases, an informed independent is better off deviating to $-B$ in state 0, changing the outcome in state 0 from policy 1 either to a 50-50 mixture of policies 0 and 1 or to policy 0. Thus for any small probability that each agent involuntarily does not participate, the deviation is profitable, so that the strategy profile is not an equilibrium.

With the equilibrium notion in place, we turn to a characterization of the equilibria of the cost benefit analysis game. For this result, we restrict attention to the two polar cases of pure common and pure private values.

Equilibrium with pure common values.—It is useful to recall an idea in Feddersen and Pesendorfer (1996). Consider a referendum with only two agents, both independents, one informed and one not. For the informed agent, voting for the policy that matches the state weakly dominates her other actions (abstain and vote for the policy different from the state). Given that the informed agent votes for the policy that matches the state, a vote by the uninformed agent, which changes the outcome in one of the states, is unambiguously detrimental. Therefore, the uninformed voter is best off abstaining. Thus the outcome of the referendum is determined by the informed voter and coincides with the outcome that would occur if both agents were informed and voted sincerely.

The intuition behind equilibrium behavior in the common value cost benefit analysis game is similar. Uninformed independents want to choose reports that are not pivotal and thus choose small reports. Informed agents want to influence the collective decision so that the policy matches the state and thus choose extreme reports. In the resulting equilibria, the correct policy is selected in both states.

Equilibrium with pure private values.—In a pure private value problem, a report of $-B$ weakly dominates all other reports for a 0-partisan and a report of B weakly dominates all other reports for a 1-partisan. Hence in every equilibrium, every 0-partisan reports $-B$ and every 1-partisan reports B .⁹ It follows that the policy favored by the larger group of partisans is chosen in equilibrium. This leads to an incorrect policy choice when a minority places a high enough value on one policy and the majority places a small enough value on the other policy.

The next result, proved in the Appendix, states these results formally.

PROPOSITION 2 (Cost benefit analysis with strategic reporting). Consider cost benefit analysis with strategic participation and reporting.

⁹ Similar intuition is developed by Kalai and Kalai (2001) in a different context.

- a. (Pure common values) Suppose that $n_0 = n_1 = 0$ and that at least one agent is informed ($n_i \geq 1$).
 - i. Every strategy profile in which each informed agent reports $-B$ in state 0 and B in state 1 and the report of each uninformed agent lies in $(0, B/n_u)$ is an equilibrium.
 - ii. In every equilibrium the welfare-maximizing policy is chosen in each state.
- b. (Pure private values) Suppose that $n_i = n_u = 0$ and that there is at least one partisan of each type ($n_0 \geq 1, n_1 \geq 1$).
 - i. The game has a unique equilibrium, in which every 0-partisan reports $-B$ and every 1-partisan reports B . The policy chosen is the one favored by the majority of agents.
 - ii. If $n_0 > n_1$, then the welfare-maximizing policy is selected in the unique equilibrium if and only if $\pi_0 n_0 \geq \pi_1 n_1$. If $n_1 > n_0$, then the condition is $\pi_1 n_1 \geq \pi_0 n_0$.

Comparing cost benefit analysis with sincere and strategic reporting.—Comparing propositions 1 and 2, we see that cost benefit analysis with sincere reporting always chooses the correct policy for private value problems and sometimes chooses incorrectly for common value problems, whereas cost benefit analysis with strategic reporting always chooses the correct policy for common value problems and sometimes chooses incorrectly for private value problems.

Estimates derived from stated preference methods generally rest on the assumption that the sample of respondents is representative of the population (at least on the basis of unobservable characteristics). In fact, since a high nonresponse rate creates the possibility of sampling bias, Arrow et al. (1993) give a high nonresponse rate as a reason for discounting the conclusions of a stated preference survey. Proposition 2 suggests, to the contrary, that allowing for strategic nonresponse may improve the ability of surveys to aggregate information.¹⁰

Another interesting implication of proposition 2 is that when reporting is strategic, the individuals' reports do not necessarily convey much information about the individuals' preferences. When values are private, all individuals make extreme reports regardless of their true values for the policies. When values are common, the game has multiple equilibria, and in at least one equilibrium the reports of all informed agents are qualitatively correct but extreme, whereas the reports of all uninformed agents are close to zero.

¹⁰ This point is closely related to the point made at the end of sec. 6 of Feddersen and Pesendorfer (1996).

C. *Referendum with Strategic Reporting*

The equilibria of the referendum game are similar to those of the model of Feddersen and Pesendorfer (1996). In particular, in any equilibrium all informed independents vote for the correct policy, given the state, and every partisan votes for her favorite policy. The uninformed independents vote in such a way as to cancel out, as far as possible, the votes of the partisans.

The uninformed independents would like the policy to be chosen by the informed independents. This is possible if there are enough independents to cancel out the partisans' votes. In this case, our notion of equilibrium selects the strategy profile most robust to nonparticipation. In this strategy profile, the uninformed independents cast their votes so that the margin in favor of policy 0 in state 0 is equal to the margin in favor of policy 1 in state 1. If there are too few independents to outvote the partisans, all uninformed independents vote for the same policy, maximizing the influence of the informed independents in case some partisans do not participate. These conclusions are formalized in the following result, which is proved in the Appendix.

PROPOSITION 3 (Referendum with strategic participation). Consider a referendum with strategic participation and voting. Suppose that at least one agent is an informed independent ($n_i \geq 1$). A strategy profile is an equilibrium if and only if it satisfies the following conditions.

- a. Every informed independent votes for policy 0 in state 0 and policy 1 in state 1.
- b. Every 0-partisan votes for policy 0 and every 1-partisan votes for policy 1.
- c. If $n_0 > n_1$, then the difference between the number of uninformed independents who vote for policy 1 and the number who vote for policy 0 is $\min\{n_0 - n_1, n_u\}$. If $n_0 \leq n_1$, then the difference between the number of uninformed independents who vote for policy 0 and the number who vote for policy 1 is $\min\{n_1 - n_0, n_u\}$.

In any equilibrium, the policy selected in state 0 is 0 if $n_1 - n_0 < n_i + n_u$ and 1 if $n_1 - n_0 > n_i + n_u$, and the policy selected in state 1 is 1 if $n_0 - n_1 < n_i + n_u$ and 0 if $n_0 - n_1 > n_i + n_u$.

In an equilibrium, the policy selected in state 0 is welfare-maximizing if and only if $n_i + n_u = \pi_1 n_1 - \pi_0 n_0$, $n_i + n_u > \max\{n_1 - n_0, \pi_1 n_1 - \pi_0 n_0\}$, or $n_i + n_u < \min\{n_1 - n_0, \pi_1 n_1 - \pi_0 n_0\}$. The policy selected in state 1 is welfare-maximizing if and only if $n_i + n_u = \pi_0 n_0 - \pi_1 n_1$, $n_i + n_u > \max\{n_0 - n_1, \pi_0 n_0 - \pi_1 n_1\}$, or $n_i + n_u < \min\{n_0 - n_1, \pi_0 n_0 - \pi_1 n_1\}$. In particular, if either $\pi_0 = \pi_1 = 1$ or $n_0 = n_1 = 0$ (pure common values), the policy selected in an equilibrium is welfare-maximizing in both states.




FIG. 3.—An example of the welfare-maximizing policy in state 0 and the policy chosen in state 0 in an equilibrium for a referendum with strategic behavior, with $\pi_0 < \pi_1 < 1$. The welfare-maximizing policy is policy 0 below the dashed line and policy 1 above it.

The result shows, in particular, that a referendum always arrives at the correct policy decision in common value environments ($n_0 = n_1 = 0$) and sometimes makes incorrect decisions in private value environments ($n_i = n_u = 0$). In an environment in which both independents and partisans are present, a referendum can make the wrong decision when partisans value the policies differently from each other and/or from independents.

The difference between the policy chosen by a referendum and the welfare-maximizing policy when the state is 0 is illustrated in figure 3 for a case in which $\pi_0 < \pi_1 < 1$. Below the dashed line, policy 0 is welfare-maximizing; above the line, policy 1 is welfare-maximizing. In the two regions between the solid and dashed lines, a referendum chooses the wrong policy.

The characterization of equilibria in proposition 3 is closely related to propositions 2 and 3 of Feddersen and Pesendorfer (1996) (see also Austen-Smith and Banks 1996; Feddersen and Pesendorfer 1998, 1999). Feddersen and Pesendorfer say that the equilibria “fully aggregate information”: sincere voting by an electorate in which all independents know the state would result in the same outcome. Their result is asymptotic. Because of the different way in which we have formulated our model, our result holds for all population sizes.

D. Cost Benefit Analyses versus Referenda

We can now compare a cost benefit analysis with a referendum.¹¹ Under the assumption that the subjects of a cost benefit analysis behave sin-

¹¹ To be complete, our comparison of a cost benefit analysis and a referendum should

cerely, the result of the cost benefit decision rule is given by proposition 1: in private value environments this institution arrives at the correct decision, whereas in some common value environments it makes the incorrect decision. The outcome of a referendum is described by proposition 3: in common value environments referenda choose correctly, whereas in private value environments they may choose incorrectly. Thus if the subjects of a cost benefit analysis behave sincerely, then a cost benefit analysis is strictly better than a referendum in a private value environment, whereas a referendum is strictly better than a cost benefit analysis in a common value environment.

The conclusion is different if the subjects of a cost benefit analysis behave strategically. In this case, comparing proposition 2 (cost benefit analysis) with proposition 3 (referendum), we see that the two decision rules lead to exactly the same decisions in both pure common and pure private value environments.¹²

The intuition behind these conclusions is clear. When reporting strategically in a cost benefit analysis in a private value environment, every agent submits the largest or smallest possible report. Thus the realized reports are identical, except for their names, to those realized in a referendum. The incentive for agents to strategically misreport their private values in a cost benefit analysis prevents the analyst from learning any cardinal information about their preferences; only ordinal information, which is also revealed by a referendum, is obtained. When reporting strategically in a common value environment, informed and uninformed agents face similar incentives in referenda and cost benefit analyses. Uninformed agents generally do not want to be pivotal, whereas informed agents do. The larger strategy sets in a cost benefit analysis generate more equilibria than exist for a referendum, but all these equilibria are welfare-maximizing.

consider the cost of each method. We ignore these costs for two reasons. First, we suspect that the cost of making a decision by either method is usually small compared to the value of the policies selected by these methods. For example, Executive Order 12866 requires cost benefit analysis only for policies whose impact is expected to be greater than \$100 million. Second, a preliminary investigation suggests that the costs of the two decision-making methods are approximately the same. In particular, the average value of 32 grants awarded by the U.S. Environmental Protection Agency between 2000 and 2005 for the purpose of conducting stated preference valuation was about \$283,000. In contrast, the state budget for California allocated about \$128 million to elections for 2008. During this time, the state conducted five elections to resolve 327 choices (12 ballot initiatives and 315 public offices), for an average cost per decision of about \$391,000 (authors' calculations).

¹² Proposition 2 does not cover the case of mixed environments.

V. Discussion

A. *Empirical Evidence on Referenda and Information Aggregation*

That uninformed agents are less likely to vote than informed agents has some empirical support. Using U.S. data from the early 1970s, Wolfinger and Rosenstone (1980) find that more educated individuals are much more likely to vote. Milligan, Moretti, and Oreopoulos (2004) find a similar positive effect more recently in both the United States and the United Kingdom, though the effect is much smaller in the United Kingdom. Assuming that more educated agents are better informed, this evidence is consistent with better-informed agents' being more likely to vote.

Lassen (2005) finds more direct evidence that voters are better informed than nonvoters. He examines a natural experiment in Copenhagen in which city residents were asked to vote in a referendum to decentralize municipal service provision. He finds that residents who were better informed about the proposal, by virtue of living in arbitrarily selected pilot districts, were more likely to vote than other residents. Further, the effect is substantial: by one estimate, being informed increased the propensity to vote by 20 percentage points.

To our knowledge, strategic abstention by uninformed independents has not been observed in laboratory experiments. However, similar strategic behavior has been observed. Guarnaschelli, McKelvey, and Palfrey (2000) and Ladha, Miller, and Oppenheimer (2003) report experiments in which small groups of subjects (three and six) played a common value majority rule voting game. Both find evidence for strategic voting behavior of the sort predicted by Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1998). Wit (1999) also reports an experimental analysis of a common value election (also with small groups) and finds that elections are effective at aggregating dispersed information.

B. *Relevance to Real-World Policy Decisions*

While the concept of a private value public policy decision is conventional, that of a common value public policy decision is at least superficially novel. Since much of our analysis is devoted to common value decisions, the importance of the analysis hinges on the existence and prevalence of such decisions.

The canonical example of a common value environment is an auction of offshore drilling rights (e.g., Hendricks, Pinkse, and Porter 2003). The value of the drilling rights is imagined to be fixed and constant across firms but measured with error by each firm. That is, the value of the drilling rights is common but not precisely known by any firm.

Common value elements are present in any auction in which the item for sale is subject to resale and the resale value is not known with certainty.¹³

The information structure underlying a pure common value auction is consistent with our specification of a pure common value environment: all agents agree on the value of the object or policy in question, but some are unsure of this value.

With this said, it is not clear that the examples described in the auction literature are relevant to our investigation of cost benefit analysis. In particular, common value auctions typically involve technical uncertainty that could be resolved if better data or better instruments were available. When public policy decisions involve such technical uncertainty, it is probably better to invest in these better data and better instruments rather than to survey an uninformed population. For example, it probably does not make sense to conduct a survey to determine the costs of global warming: the respondent population probably does not have much information to aggregate.

Our analysis is more usefully applied to “experience goods.” An experience good has the property that a consumer does not know her value until she actually uses it. An apple is a simple example: it is difficult to determine how well one likes an apple before eating it. Examples of experience goods studied in the empirical literature suggest that these goods are commonplace (e.g., yogurt [Akerberg 2003] and the quality of employee-to-employer matching [Light and McGarry 1998; Pries and Rogerson 2005]). The popularity of guidebooks also suggests that experience goods are common. The purpose of guidebooks is to provide information about unfamiliar goods. The popularity of these books suggests that people can benefit by relying on the author’s assessment of an unfamiliar good. Indeed, Slovic (1995) argues that individuals construct preference rankings as they are needed rather than recalling them. If so, then it is reasonable to believe that experience goods are pervasive. For our purpose, the idea of an experience good is important because the values of such goods are uncertain to consumers, and no technical method of resolving this uncertainty is available except for actual consumption.

Our analysis applies to policy decisions that involve “public experience goods.” That is, it applies to policies that agents cannot value without experiencing them. The uncertainty involved in such policy decisions will not yield to better instrumentation but will yield to experience. The

¹³ Other specific examples include procurement auctions (Hong and Shum 2002), interbank financial markets (Linzert, Nautz, and Bindseil 2007), and eBay auctions for computers (Yin 2007).

problem that the planner faces is to elicit the values placed on the policy by agents with the relevant experience.

While the idea of a public experience good may appear novel, it is new in name only. Researchers who administer stated preference surveys have long been concerned with the way that participants' valuations respond to changes in the amount and type of information participants are given. For example, in his important paper on survey methodology, Hanemann (1994) advises that "providing adequate and accurate information" (24) and allowing respondents time to "reflect and give a considered opinion" (22) increase the accuracy with which respondents report their values. This advice echoes another important reference on survey methodology, Arrow et al. (1993).

The literature on stated preference surveys has gone beyond these early exhortations and has tried to measure the effects of information on preferences. Macmillan et al. (2002) estimate the value of a policy to compensate farmers for the depredations of migrant geese. To do so, they conducted three surveys. The first sample of respondents was asked to value the policy after a brief description of the problem that did not allow an opportunity to ask questions. The second sample of respondents was given a more extensive description of the problem and was permitted to ask questions about the policy. The third survey re-sampled the second group of respondents after they had been given a week to think about the problem and discuss it with friends and family and had been given another opportunity to ask questions of the researchers. Preferences change systematically from one survey to the next. In particular, about a third of the respondents to the third survey revised their initial responses, and most of these revisions were downward. That is, Macmillan et al. show that reported valuations change systematically in response to increases in respondents' information. This effect is consistent with the policy being a public experience good as we have defined it.¹⁴

In sum, many surveys ask respondents to formulate their preferences about a public policy with which they have little experience. The possibility that the reported preferences vary systematically with the respondents' information is regarded as a fundamental issue in the literature on stated preference surveys. Such a systematic relationship between preferences and information is precisely what our common value framework is intended to reflect.

Having argued in the abstract that cost benefit analysis often concerns itself with common value environments, we now suggest three examples.

¹⁴ Shapansky, Adamowicz, and Boxall (2008) also investigate the effect of providing different levels of information to survey respondents but find that it does not affect mean responses. Legget (2002) develops econometric methods for estimating utility functions from survey data describing a public experience good.

First, consider the problem of assessing the value of a policy change that would improve the quality of drinking water. While one can provide an extensive description of the changes that would occur to the water as a consequence of the policy, it is probably difficult for survey respondents to value the change unless they have had some of the water to drink. The policy may reduce pathogens, and hence illness, but survey respondents will find the policy difficult to value unless they have had access to similar-quality water for a fairly extended period of time. Agents who have lived in areas where the better-quality water is available correspond to our “informed” type.

Second, consider the problem of assessing the value of improving water quality at a beach. As with drinking water, one can provide an extensive description of how the policy will improve water quality, but it will probably be difficult for respondents to value the change until they see and smell the cleaned-up beach. Respondents who have experience with clean and dirty beaches will be better able than respondents who know only the dirty beach to assess these values. These experienced and inexperienced respondents correspond to our informed and uninformed agents.

Third, consider the problem of choosing between two potential designs for a neighborhood park when there is a consensus among survey respondents that the park should serve principally as a venue for winter ice skating and summer ball games and picnics. Architects may provide informative drawings and even models of the two designs, but it will be difficult for residents to get a sense of what the park will be like without spending a lot of time studying the drawings and without prior experience of architectural drawings. Agents who spend a lot of time studying the drawings or who have prior experience with such drawings will be better able to assess the values of the two designs, and they correspond to our informed agents.

Policy decisions like these appear to be modeled well as common value problems. Policy decisions that have primarily distributional consequences are, by contrast, better modeled as private value problems. A typical problem of this type is the choice for the location of a municipal landfill. Every agent dislikes being close to a landfill and prefers that such a facility be located near someone else.

C. Strategic versus Sincere Behavior in Cost Benefit Analysis

Our analysis indicates that when agents respond strategically to survey questions, referenda and cost benefit analysis always lead to the same policy choice, but that when survey respondents are sincere, cost benefit analysis dominates referenda in private value environments and con-

versely in common value environments. Which behavioral assumption is relevant?

Stated preference analyses usually simply posit that survey respondents are truthful. The exceptions to this assumption seem to prove the rule. For example, Carson, Groves, and Machina (2000) contemplate the possibility of strategic responses to stated preference questions. Their object is to determine conditions when such strategic behavior will lead to truthful revelation. While their arguments are mostly informal, they conclude that it is often possible to formulate stated preference questions so that agents truthfully reveal their preferences.

Two arguments suggest that agents tend to respond to survey questions sincerely. First, if agents respond strategically, by proposition 2 we ought to observe only extreme reports in a private value environment, whereas in a common value environment we ought to see agents' responses become more extreme as the agents become better informed. These patterns of responses do not appear to be widely observed. For example, Gregory et al. (1995) and Macmillan et al. (2002) report distributions of survey responses that are not consistent with strategic behavior under either of the information environments considered in proposition 2 but are easy to reconcile with sincere reporting.

Second, many cost benefit analyses are based on revealed preferences, using, for example, travel cost or hedonic data. Even if we suspect strategic behavior in stated preference cost benefit analyses, it is implausible to suspect such behavior in travel cost and hedonic cost benefit analyses: people do not choose their housing or recreational travel with an eye to influencing a cost benefit analysis. Therefore, if strategic responses to stated preference surveys are common, we ought to see systematic differences in the conclusions of analyses based on revealed and stated preferences. The survey of many revealed and stated preference cost benefit analyses conducted by Carson et al. (1996) does not support this conclusion. In sum, both arguments suggest that respondents behave sincerely when they participate in stated preference surveys.

Finally, we note that stated preference surveys are not generally administered with the same solemnity as referenda, and the link between survey responses and outcomes is more subject to doubt than the link between a referendum and the outcome.¹⁵ People may regard stated preference surveys as an academic exercise or a nuisance but are not likely to view a referendum in this way.¹⁶ Thus, it is probably not reasonable to expect the same degree of strategic behavior from people participating in a cost benefit analysis as from those participating in a

¹⁵ To our knowledge, nowhere is there legislation guaranteeing participation in stated preference surveys. Voting rights, on the other hand, are commonly protected by law and constitution.

¹⁶ This claim is contentious. See Carson et al. (2000).

referendum. However, if, as in our formal model, stated preference cost benefit analysis were used as a mechanistic decision rule (as referenda are), then strategic behavior might emerge.

D. Implications for Travel Cost and Hedonic Methods

Like stated preference methods, revealed preference methods (e.g., travel cost and hedonic methods) seek to estimate agents' willingness to pay for a particular policy. For common value environments it is not clear that such methods will elicit the same estimates from independents as stated preference methods. Travel cost data often (but not always) describe individuals who have visited a site repeatedly and have presumably learned at least some of its characteristics. The hedonic method is based on market prices, and the extent to which these prices allow us to recover private information is open to question. Nevertheless, the basic assumptions in propositions 1 and 3 remain plausible. In common value problems some agents are better informed about the characteristics of their destination or housing locations than others, and their travel and housing decisions may reflect these differences in information. Thus values imputed to uninformed agents by revealed preference methods, like those obtained by stated preference methods, may reflect decisions based on expected values. This condition is sufficient to generate the differences between referenda and cost benefit analyses given in propositions 1 and 3: referenda are superior to cost benefit analyses in common value environments, whereas cost benefit analyses are superior to referenda in private value environments.

E. Agents' Motivations

We assume throughout that our agents are motivated by the expectation that their actions will affect the outcome. When the number of agents is large, however, the impact on the outcome of the vote or report of any single agent is small and quite possibly insufficient to outweigh the costs of voting or reporting. A possible explanation for significant participation in such situations is that agents have expressive rather than instrumental motivations. This idea has been explored in the context of an election by Feddersen and Sandroni (2006) and Feddersen, Gailmard, and Sandroni (2008), whose agents are motivated by ethical concerns. As we have noted, though cost benefit analyses and referenda are symmetric in our model, they differ considerably in practice. In particular, ethical concerns may affect agents' participation decisions in the two mechanisms differently. As a referee suggests, ethical concerns may play a more significant role in an election, where voters have been

mobilized, than in a cost benefit analysis. The implications for a comparison of referenda and cost benefit analyses are unclear.

F. Planner's Commitment to a Mechanism

In our model, the planner is committed to the mechanism mapping the agents' reports into a policy. Morgan and Stocken (2008) study a model of polls in which the policy maker is not committed to a decision-making mechanism. They find conditions under which a poll fully aggregates the agents' information. They also compare polls with referenda and show (in their sec. 4) that in a model with a binary policy space (which allows the two mechanisms to be compared), a referendum may reveal information better than a poll.

VI. Conclusion

The clarity of the language with which legislators, regulators, and professional economists call for calculations of the economic costs and benefits of policy decisions conceals the difficulty of such calculations. In response to this difficulty, the U.S. Office of Management and Budget (1992) has issued guidelines on how to perform a cost benefit analysis, as has at least one panel of distinguished economists (Arrow et al. 1996). Both sets of instructions implicitly adopt the standard, now widespread in the profession, that a cost benefit analysis is good if it produces an accurate estimate of the costs and benefits of the policy in question.

We propose a different standard: a cost benefit analysis is good or bad according to whether it leads to a better public decision than competing decision rules.¹⁷ With this standard in mind, we compare a stylized cost benefit decision rule with a referendum, another widely used institution for making public decisions.

Our conclusions depend on the extent to which agents behave strategically in their interactions with the analyst performing the cost benefit analysis.

Under the assumption of sincere reporting, a cost benefit analysis elicits cardinal information about preferences whereas a referendum elicits only ordinal information. If this cardinal information is important, then a cost benefit analysis leads to a better decision than a referendum does. Conversely, a referendum can aggregate widely dispersed information, whereas cost benefit analysis simply recovers a common prior. If information about the state of the world is important, then a referendum leads to a better decision than a cost benefit analysis does. This

¹⁷ Diamond and Hausman (1994) call for a related calculation: a comparison of decisions made with and without cost benefit analysis.

logic leads to the conclusion that a cost benefit analysis is superior to a referendum in private value environments and inferior in common value environments.

For a cost benefit analysis conducted with a stated preference methodology, it is of interest to examine the implications of strategic responses to survey questions. When agents participate and report strategically, a cost benefit analysis and a referendum elicit qualitatively identical behavior, and the two methods always result in the same policy choice. Thus, if we believe that agents behave strategically in cost benefit analyses, there is no reason to prefer one decision rule over the other.

Without taking a stand on whether agents approach stated preference cost benefit analyses strategically, we can make a strong statement about the comparative advantages of the two decision rules: cost benefit analyses are always at least as good as referenda for private value problems, whereas referenda are always at least as good as cost benefit analyses for common value problems. If our skepticism about strategic behavior in cost benefit analyses is warranted, a stronger statement is possible: cost benefit analyses are strictly better than referenda for private value problems, whereas referenda are strictly better than cost benefit analyses for common value problems.

Our results suggest that a determination of whether or not to rely on referenda or cost benefit analyses for any given public decision depends on whether valuations are private or common. This question appears to be difficult to answer *ex ante* for any given public decision. However, the fundamental difference between private value and common value environments is that preferences in common value environments change systematically after a policy decision is made, whereas preferences in private value decisions do not. This suggests the possibility of distinguishing common value from private value decisions *ex post*. Both Macmillan et al. (2002) and Shapansky et al. (2008) conduct exercises along these lines. In principal, a collection of studies performed on different public decisions could allow regulators to determine classes of public decisions in which common values are important and in which they are not. This information would, in turn, allow a determination of when referenda should be expected to outperform cost benefit analyses and when they should not.

Appendix

Proof of Proposition 1

In state 0, cost benefit analysis with sincere reporting chooses policy 0 if $-\pi_0 n_0 + \pi_1 n_1 - n_i + (1 - 2\alpha)n_u < 0$ (see [3]). Policy 0 is welfare-maximizing in state 0 if and only if $\pi_0 n_0 - \pi_1 n_1 + n_i + n_u \geq 0$ (see [1]). The result for state 0 follows immediately. A symmetric argument applies to state 1. QED

In the remaining proofs, an equivalent version of our definition of an equilibrium (definition 1) is useful. Consider a player choosing between two strategies that generate the same payoff given the other players' strategies. When ϵ is small, for any positive integer k , the effect on the player's expected payoff of the involuntary nonparticipation of $k + 1$ or more players is negligible compared with the effect of the involuntary nonparticipation of k players. Therefore, the requirement that a player's strategy be a best response to the other players' strategies for all small ϵ means that the player's choice between the strategies is based on the case of the involuntary nonparticipation by the smallest number of players for which the expected payoffs of the strategies differ. If, for example, the player's strategies r_j and r'_j yield her the same expected payoff when all players participate and also over all the cases in which one of the other players involuntarily does not participate, but the expected payoff of r_j over all cases in which two of the other players involuntarily do not participate is higher than the expected payoff of r'_j in this case, then she chooses r_j .

Precisely, the following implication of our definition of equilibrium is useful.

A strategy profile is an equilibrium if and only if for each agent and each change in her strategy, one of the following conditions is satisfied.

- For every integer k with $0 \leq k \leq n - 1$, the change does not affect her expected payoff when each set of k nonparticipants is equally likely.
- For the smallest integer k for which the change does affect her expected payoff when each set of k nonparticipants is equally likely, it decreases this expected payoff.

Proof of Proposition 2

ai: The outcome of any such strategy profile is policy 1 in state 1 and policy 0 in state 0. This outcome is the best possible outcome for every agent, so no deviation can improve any agent's payoff when all agents participate.

Now consider a perturbation of the game in which some agents do not participate. If at least one informed agent participates, the outcome is policy 1 in state 1 and policy 0 in state 0. If only uninformed agents participate, the outcome is policy 1. In both cases, the outcome is the best possible outcome for every participant (given that α , the prior probability of state 0, is less than one-half), so that no deviation by any agent increases her payoff. Thus any such strategy profile is an equilibrium.

aii: We first argue that in every equilibrium, every informed agent's report is positive in state 1 and negative in state 0.

Suppose to the contrary that the report of some informed agent j in state 0 is nonnegative. Then if no other agent participates, the outcome is policy 1 if j 's report is positive and each policy with probability one-half if j 's report is 0. Thus a deviation by j to $-B$ improves the outcome (to policy 0) in state 0 if no other agent participates and either improves the outcome or does not affect the outcome for any other set of reports. Thus every informed agent's report in state 0 is negative in any equilibrium. Similarly, every informed agent's report in state 1 is positive in any equilibrium.

Now let r be a strategy profile and for each state s let

$$R(s) = \sum_{j \in I} r_j(s) + \sum_{j \in U} r_j$$

be the sum of the reports in state s , where I is the set of informed independents and U is the set of uninformed independents. Assume that the outcome in state 1 is not correct, so that $R(1) \leq 0$. By the previous argument, $r_j(1) > 0$ and $r_j(0) < 0$ for every $j \in I$, so $R(0) < 0$ and $r_j < 0$ for some $j \in U$. Denote by m the uninformed agent for whom r_m is smallest (or any such agent if there is more than one such agent). (That is, $r_m < 0$ and $r_m \leq r_j$ for all $j \in U$.)

We claim that there is an alternative report r'_m for agent m and a nonnegative integer k such that m 's deviating to r'_m does not affect the outcome if at most k other agents do not participate and increases m 's expected payoff over all cases in which k other agents do not participate.

First suppose that $r_m - R(1) < B$. Choose $\delta > 0$ such that $r_m - R(1) + \delta \leq B$ and $\delta < R(1) - R(0)$. Then m can deviate to the report $r'_m = r_m - R(1) + \delta$ ($\geq r_m$), and if she does so, the sum of the reports in state 1 becomes positive (it becomes δ) whereas the sum of the reports in state 0 remains negative (it becomes $R(0) - R(1) + \delta$). Thus m 's deviation induces a better outcome and increases her payoff. We conclude that if $r_m < B + R(1)$, then r is not an equilibrium.

Now suppose that $r_m - R(1) = B$. Then $R(1) < 0$ because $r_m < 0$. Thus if m deviates to $r'_m = B$, the outcome improves because the policy chosen in state 1 changes from 0 to a 50-50 mixture of 0 and 1 whereas the policy chosen in state 0 remains 0. Hence r is not an equilibrium.

Finally, suppose that $r_m - R(1) > B$. In this case, no deviation by m —even a deviation to B —affects the outcome in either state when all agents participate. Suppose that one agent does not participate. Denote by j the agent for whom r_j is smallest among the agents other than m . The sum of the reports when j does not participate is $R(1) - r_j$. Given that $r_j \geq r_m$, we have $R(1) - r_j \leq R(1) - r_m < -B < 0$ (given $r_m - R(1) > B$). That is, when j does not participate, the sum of the reports in state 1 (and therefore also the sum in state 0) remains negative.

There are two cases to consider. If $r_m - R(1) + r_j \leq B$, then by the arguments in the previous two paragraphs, if agent j does not participate, agent m has a deviation that changes the outcome to either policy 1 or a 50-50 mixture of the two policies in state 1 and retains policy 0 in state 0, thus increasing her payoff. Further, the report of every other agent is at least r_j , so when any other single agent does not participate, this deviation either increases m 's payoff or has no effect on the outcome. Thus if $r_m - R(1) + r_j \leq B$, then m has a deviation that increases her expected payoff over all cases in which one agent does not participate, so that the strategy profile is not an equilibrium.

If $r_m - R(1) + r_j > B$, then no deviation by agent m affects the outcome in either state when at most one agent does not participate. In this case, we can take the agent j' whose report is smallest among the agents other than m and j and repeat the argument of the previous paragraph for the case in which j and j' do not participate. We can continue in the same way for larger sets of nonparticipants, at each step adding the agent whose report is smallest among the reports of the remaining agents. For each number l of nonparticipants, one of the following two cases occurs.

- i. Agent m has a deviation that changes the policy from 0 to 1 (or a 50-50 mixture of 0 and 1) in state 1 and retains policy 0 in state 0 for some set (or sets) of l nonparticipants and does not affect the policy in either state for other sets of l nonparticipants. In this case r is not an equilibrium.
- ii. No deviation by agent m affects the outcome in either state for any set of l nonparticipants.

In case ii we increment l and proceed to the next step. If we reach $l = n_u - 1$, the process terminates with case i because when the set of nonparticipants consists of all the uninformed agents except m , every participant is informed other than m and hence (by the earlier argument) makes a positive report in state 1.

We conclude that in any equilibrium the outcome in state 1 is correct. A similar argument shows that in any equilibrium the outcome in state 0 is correct.

bi: First observe that if a 0-partisan whose report is greater than $-B$ deviates to $-B$, either the outcome does not change or it improves from the agent's perspective. A deviation by a 1-partisan from a report of less than B to B has a similar effect.

To prove the result, we need to argue only that for any strategy profile r in which some 0-partisan reports more than $-B$ or some 1-partisan reports less than B , there is an agent m and a subset S of the other agents such that the deviation by m from r_m to $-B$ (if m is a 0-partisan) or B (if m is a 1-partisan) strictly improves the outcome from m 's perspective when the set of nonparticipants is S . Denote by m_0 the 0-partisan whose report is largest (or any such agent in the event of a tie) and by m_1 the 1-partisan whose report is smallest. Of these two agents, choose the one whose report deviates the most from the extreme report appropriate for her type. (That is, choose the 0-partisan if $r_{m_0} > -r_{m_1}$, the 1-partisan if $r_{m_0} < -r_{m_1}$, and either in the case of equality.) Denote this agent by m , and assume, without loss of generality, that she is a 0-partisan. Let S be the set of all the other agents with the exception of a single 1-partisan, say m' . When the set of nonparticipants is S , there are exactly two participants, m and m' . By the choice of m , we have $r_m + r_{m'} \geq 0$ and $-B + r_{m'} \leq 0$, with at least one of these inequalities strict (otherwise $r_m = -B$ and $r_{m'} = B$). Thus if m deviates to $-B$ when the set of nonparticipants is S , the outcome changes either from policy 1 to policy 0, from policy 1 to a 50-50 mixture of policies 0 and 1, or from a 50-50 mixture of policies 0 and 1 to policy 0. In all cases, the deviation increases m 's payoff, so r is not an equilibrium.

bii: In the unique equilibrium, if $n_0 > n_1$, then policy 0 is selected; if $n_1 = n_0$, then each policy is selected with probability one-half; and if $n_1 > n_0$, then policy 1 is selected. The welfare-maximizing policy is policy 0 if $\pi_0 n_0 - \pi_1 n_1 > 0$, policy 1 if $\pi_0 n_0 - \pi_1 n_1 < 0$, and either policy if $\pi_0 n_0 - \pi_1 n_1 = 0$ (see [1] or [2]). The result follows immediately. QED

Proof of Proposition 3

Denote by r^* a strategy profile that satisfies conditions *a-c* of the result. We first argue that r^* is an equilibrium.

If an informed independent changes her action in state 0, then depending

on the pattern of participation by the other players, either the outcome does not change or it changes to yield policy 1 with positive probability (one-half or one). If she is the only participant, then the change is of the latter type. Similar arguments apply to the action of an informed independent in state 1 and to the actions of partisans.

Now consider an uninformed independent, say j . There are two cases to consider.

1. First suppose that the absolute value of the difference between the numbers of 0-partisans and 1-partisans is less than the number of independents (informed and uninformed). Then the outcome of r^* is policy 0 in state 0 and policy 1 in state 1; the vote margin in favor of policy 0 in state 0 is equal to the vote margin in favor of policy 1 in state 1 and is equal to n_i , the number of informed independents.

Consider a case in which $n_0 > n_1$. Suppose that j votes for policy 0. If she switches to abstention, the outcome changes only if the number of nonparticipants is at least $n_i - 1$. A set of $n_i - 1$ nonparticipants causes the change in j 's strategy to affect only the outcome in state 0 and only if all nonparticipants vote for policy 0 in r^* . In this case, j 's switch changes the outcome in state 0 to policy 1 with probability one-half, which decreases j 's expected payoff. Similarly, a switch by j to vote for policy 1 is detrimental.

Similar arguments show that if j votes for policy 1 or abstains, any change in her strategy reduces her expected payoff when the number of nonparticipants is the smallest number for which the change in j 's strategy changes the outcome.

We conclude that when $n_0 > n_1$, r^* is an equilibrium. A symmetric argument applies when $n_0 \leq n_1$.

2. Now suppose that the absolute value of the difference between the number of 0-partisans and the number of 1-partisans is at least the number of independents (informed and uninformed). Then the outcome of r^* is the same policy in both states. Suppose specifically that 0-partisans outnumber 1-partisans, so that the outcome of r^* is policy 0 in both states. Then r^* specifies that every uninformed independent, and in particular j , votes for policy 1. Denote by k_0 the vote margin in favor of policy 0 in state 1. If j switches to abstention, the outcome changes only if at least $k_0 - 1$ agents do not participate, and a set of $k_0 - 1$ nonparticipants causes the change in j 's strategy to affect only the outcome in state 1 and only if all nonparticipants vote for policy 0 in r^* . In this case, j 's voting for policy 1, as r^* prescribes, yields each policy with probability one-half, whereas abstaining yields policy 0. Thus the deviation reduces j 's expected payoff. The same argument shows that a deviation by j to vote for policy 0 also reduces her payoff. A symmetric argument applies when the number of 0-partisans is at most the number of 1-partisans. Thus in this case, as in the previous case, r^* is an equilibrium.

We now argue that every equilibrium satisfies conditions *a-c* of the result. We first show that in every equilibrium, every informed independent votes for policy 0 in state 0 and 1 in state 1.

Let r_j be a strategy of agent j , an informed independent. Suppose, contrary to the claim, that $r_j(0) \neq -1$ (i.e., j does not vote for policy 0 in state 0). Then a deviation by j to vote for policy 0 in state 0 increases her payoff if no other agent participates and either increases her payoff or does not affect it for any

other set of actions of the other players. Thus $r_j(0) = -1$ in any equilibrium. Similarly $r_j(1) = 1$ in any equilibrium.

A very similar argument establishes that in every equilibrium, every 0-partisan votes for policy 0 and every 1-partisan votes for policy 1.

Finally, consider the actions of an uninformed independent, say j . First consider a case in which $n_0 > n_1$. Suppose that the vote margin in favor of the selected policy is smaller in state 1 than it is in state 0. Denote this margin in state 1 by k_1 . Suppose that j deviates from voting for policy 0 to abstention. Any minimal set of nonparticipants for which this deviation changes the outcome consists of $k_1 - 1$ agents voting for policy 1. For such a set of nonparticipants, the change increases j 's expected payoff because it changes the outcome in state 1 from each policy with probability one-half to policy 1 for sure. A similar argument shows that a deviation by an uninformed independent from abstention to voting for policy 1 in such a case increases her expected payoff for any minimal set of nonparticipants for which the change affects the outcome. We conclude that a vote margin in favor of the selected policy that is smaller in state 1 than in state 0 and some uninformed independent not voting for policy 1 are inconsistent with equilibrium.

Given that all informed independents vote for policy 0 in state 0 and policy 1 in state 1 and all p -partisans vote for policy p in any equilibrium, we conclude that if $n_0 - n_1 > n_u$, then all uninformed independents vote for policy 1. By a similar argument, if $n_1 - n_0 > n_u$, then all uninformed independents vote for policy 0.

Finally, consider a strategy profile in which all informed independents vote for policy 0 in state 0 and policy 1 in state 1, all p -partisans vote for policy p , and the vote margin in favor of the selected policy is the same in both states. For such a strategy profile, the difference between the number of uninformed independents who vote for policy 1 and the number who vote for policy 0 is equal to $|n_0 - n_1|$, completing the characterization of equilibria.

If $n_0 > n_1$, then in any equilibrium, policy 0 wins in state 0 because the votes of the uninformed independents never more than cancel out the votes of the partisans. If $n_0 < n_1$, then the difference in state 0 between the numbers of votes for policies 0 and 1 is $n_i + n_0 - n_1 + \min\{n_1 - n_0, n_u\}$, which is equal to n_i if $n_1 - n_0 \leq n_u$ and to $n_i + n_u + n_0 - n_1$ otherwise. Thus the difference is positive, in which case policy 0 wins if and only if $n_1 - n_0 < n_i + n_u$. We conclude that policy 0 is chosen in state 0 if $n_1 - n_0 < n_i + n_u$. Similarly, policy 1 is chosen in state 0 if $n_1 - n_0 > n_i + n_u$.

A symmetric argument shows that in state 1, policy 1 is chosen if $n_0 - n_1 < n_i + n_u$ and policy 0 is chosen if $n_0 - n_1 > n_i + n_u$.

Finally, the comparison of the equilibrium outcomes with the outcome of welfare maximization for state 0 follows from the fact that in state 0, policy 0 is welfare-maximizing if and only if $\pi_0 n_0 - \pi_1 n_1 + n_i + n_u \geq 0$ (see [1]) and is chosen in equilibrium if $n_1 - n_0 < n_i + n_u$, and policy 1 is welfare-maximizing if and only if $-\pi_0 n_0 + \pi_1 n_1 - n_i - n_u \geq 0$ and is chosen in equilibrium if $n_1 - n_0 > n_i + n_u$. A similar argument applies to state 1. QED

References

- Ackerberg, Daniel A. 2003. "Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination." *Internat. Econ. Rev.* 44: 1007–40.
- Arrow, Kenneth J., Maureen L. Cropper, George C. Eads, Robert W. Hahn, Lester B. Lave, Roger G. Noll, Paul R. Portney, Milton Russell, Richard Schmalensee, V. Kerry Smith, and Robert N. Stavins. 1996. "Is There a Role for Benefit-Cost Analysis in Environmental, Health, and Safety Regulation?" *Science* 272:221–22.
- Arrow, Kenneth J., Robert M. Solow, Paul R. Portney, Edward E. Leamer, Roy Radner, and Howard Schuman. 1993. "Report of the NOAA Panel on Contingent Valuation." *Fed. Register* 58:4601–14.
- Austen-Smith, David, and Jeffrey S. Banks. 1996. "Information Aggregation, Rationality, and the Condorcet Jury Theorem." *American Polit. Sci. Rev.* 90:34–45.
- Banerjee, Abhijit, and Rohini Somanathan. 2001. "A Simple Model of Voice." *Q.J.E.* 116:189–227.
- Carson, Richard T., Nicholas E. Flores, Kerry M. Martin, and Jennifer L. Wright. 1996. "Contingent Valuation and Revealed Preference Methodologies: Comparing the Estimates for Quasi-Public Goods." *Land Econ.* 72:80–99.
- Carson, Richard T., Theodore Groves, and Mark J. Machina. 2000. "Incentive and Informational Properties of Preference Questions." Manuscript, Dept. Econ., Univ. California, San Diego.
- Cherry, Todd L., Thomas D. Crocker, and Jason F. Shogren. 2003. "Rationality Spillovers." *J. Environmental Econ. and Management* 45:63–84.
- Clinton, William J. 1993. "Presidential Documents, Executive Order 12866 of September 30, 1993: Regulatory Planning and Review." *Fed. Register* 58 (October 4).
- Cropper, Maureen L., William N. Evans, Stephen J. Berardi, Maria M. Duclasoares, and Paul R. Portney. 1992. "The Determinants of Pesticide Regulation: A Statistical Analysis of EPA Decision Making." *J.P.E.* 100:175–97.
- Diamond, Peter A., and Jerry A. Hausman. 1994. "Contingent Valuation: Is Some Number Better than No Number?" *J. Econ. Perspectives* 8:45–64.
- Feddersen, Timothy J., Sean Gailmard, and Alvaro Sandroni. 2008. "Moral Bias in Large Elections: Theory and Experimental Evidence." Manuscript, Dept. Polit. Sci., Univ. California, Berkeley.
- Feddersen, Timothy J., and Wolfgang Pesendorfer. 1996. "The Swing Voter's Curse." *A.E.R.* 86:408–24.
- . 1998. "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting." *American Polit. Sci. Rev.* 92:23–35.
- . 1999. "Elections, Information Aggregation, and Strategic Voting." *Proc. Nat. Acad. Sci. United States of America* 96:10572–74.
- Feddersen, Timothy J., and Alvaro Sandroni. 2006. "A Theory of Participation in Elections." *A.E.R.* 96:1271–82.
- Gregory, Robin, Sarah Lichtenstein, Thomas C. Brown, George L. Peterson, and Paul Slovic. 1995. "How Precise Are Monetary Representations of Environmental Improvements?" *Land Econ.* 71:462–73.
- Guarnaschelli, Serena, Richard D. McKelvey, and Thomas R. Palfrey. 2000. "An Experimental Study of Jury Decision Rules." *American Polit. Sci. Rev.* 94:407–23.
- Hahn, Robert W. 2000. "The Impact of Economics on Environmental Policy." *J. Environmental Econ. and Management* 39:375–99.

- Hanemann, W. Michael. 1994. "Valuing the Environment through Contingent Valuation." *J. Econ. Perspectives* 8:19–43.
- Hendricks, Kenneth, Joris Pinkse, and Robert H. Porter. 2003. "Empirical Implications of Equilibrium Bidding in First-Price, Symmetric, Common Value Auctions." *Rev. Econ. Studies* 70:115–45.
- Hong, Han, and Matthew Shum. 2002. "Increasing Competition and the Winner's Curse: Evidence from Procurement." *Rev. Econ. Studies* 69:871–98.
- Kalai, Adam, and Ehud Kalai. 2001. "Strategic Polarization." *J. Math. Psychology* 45:656–63.
- Ladha, Krishna, Gary J. Miller, and Joe A. Oppenheimer. 2003. "Information Aggregation by Majority Rule: Theory and Experiment." Manuscript, Olin School Bus., Washington Univ. St. Louis.
- Lassen, David D. 2005. "The Effect of Information on Voter Turnout: Evidence from a Natural Experiment." *American J. Polit. Sci.* 49:103–18.
- Legget, Christopher G. 2002. "Environmental Valuation with Imperfect Information: The Case of the Random Utility Model." *Environmental and Resource Econ.* 23:343–55.
- Light, Audrey, and Kathleen McGarry. 1998. "Job Change Patterns and the Wages of Young Men." *Rev. Econ. and Statis.* 80:276–86.
- Linzert, Tobias, Dieter Nautz, and Ulrich Bindseil. 2007. "Bidding Behavior in the Longer Term Refinancing Operations of the European Central Bank: Evidence from a Panel Sample Selection Model." *J. Banking and Finance* 31: 1521–43.
- Lupia, Arthur, and John G. Matsusaka. 2004. "Direct Democracy: New Approaches to Old Questions." *Ann. Rev. Polit. Sci.* 7:463–82.
- Macmillan, Douglas C., Lorna Philip, Nick Hanley, and Begona Alvarez-Farizo. 2002. "Valuing the Non-market Benefits of Wild Goose Conservation: A Comparison of Interview and Group-Based Approaches." *Ecological Econ.* 42:49–59.
- Milligan, Kevin, Enrico Moretti, and Philip Oreopoulos. 2004. "Does Education Improve Citizenship? Evidence from the United States and the United Kingdom." *J. Public Econ.* 88:1667–95.
- Morgan, John, and Phillip C. Stocken. 2008. "Information Aggregation in Polls." *A.E.R.* 98:864–96.
- Pries, Michael, and Richard Rogerson. 2005. "Hiring Policies, Labor Market Institutions, and Labor Market Flows." *J.P.E.* 113:811–39.
- Settle, Chad, Todd L. Cherry, and Jason F. Shogren. 2003. "Rationality Spillovers in Yellowstone." Manuscript, Dept. Econ., Univ. Wyoming.
- Shapansky, Bradford, Wiktor L. Adamowicz, and Peter C. Boxall. 2008. "Assessing Information Provision and Respondent Involvement Effects on Preferences." *Ecological Econ.* 65:626–35.
- Slovic, Paul. 1995. "The Construction of Preference." *American Psychologist* 50: 364–71.
- Smith, Vernon K. 2000. "JEE and Non-market Valuation: 1974–1998." *J. Environmental Econ. and Management* 39:351–74.
- State of California. 2004. "Official Voter Information Guide, California General Election." <http://vote2004.sos.ca.gov/voterguide/english.pdf>.
- U.S. Office of Management and Budget. 1992. "Memorandum for Heads of Executive Departments and Establishments." Circular no. A-94 Revised (October). <http://www.whitehouse.gov/omb/rewrite/circulars/a094/a094.html>.
- Viscusi, W. Kip. 1996. "Regulating the Regulators." *Univ. Chicago Law Rev.* 63: 1423–61.

- Wit, Jürgen. 1999. "Social Learning in a Common Interest Voting Game." *Games and Econ. Behavior* 26:131–56.
- Wolfinger, Raymond E., and Steven J. Rosenstone. 1980. *Who Votes?* New Haven, CT: Yale Univ. Press.
- Yin, Pai-Ling. 2007. "Empirical Tests of Information Aggregation." *Internat. J. Indus. Org.* 25:1179–89.