Bayesian Improvement of the Phantom Voters Rule: An example of Dichotomic Communication

Pierre Fleckinger

Bayesian Improvement of the Phantom Voters Rule: An Example of Dichotomic Communication*

Pierre Fleckinger†

This version: October 2006

Abstract

This paper studies communication mechanisms for two players with symmetric single-peaked preferences. The peaks are privately known and drawn from a uniform distribution before the agents take a collective decision. While for the general setting Moulin (1980) characterized all strategy-proof mechanisms, much remains to be known in the Bayesian framework. The example consists of a dichotomic mechanism, that yields a strictly higher ex-ante expected utility than the best "min-max" rule. The properties of the mechanism are analyzed, then limits and possible directions for generalization are discussed.

Keywords: Single-peaked Preferences, Communication, Incentive Compatibility, Dichotomy

JEL Classification: D82, D72

1 Introduction

In a seminal paper, Moulin (1980) characterized all strategy-proof mechanisms for one-dimensional single-peaked preferences, and those satisfying in addition Pareto-optimality.

*This paper is a revised chapter of my PhD dissertation at Ecole Polytechnique. It was first presented in 2003 at University Paris I (Eurequa). I would like to thank Willie Fuchs and Jon Levin for several discussions, as well as Claude d'Aspremont, Rida Laraki, an anonymous referee, the associate editor and especially Hervé Moulin for helpful comments.

†Laboratoire d'Econométrie de l'Ecole Polytechnique, 1, Rue Descartes, 75005 Paris and INRA-LORIA. Email address:pierre.fleckinger@shs.polytechnique.fr.

Border and Jordan (1983) extended the characterization to multidimensional single-peaked preferences. More recently, it has also been extended to probabilistic mechanisms by Ehlers et al. (2002). This paper considers Bayesian mechanisms for those settings, focusing on the simplest case: two agents with peak-symmetric utility and uniform priors. While strategy proof implementation has nice normative properties and yields robust results, communication (involving beliefs) is also appealing from a positive point of view, especially in settings with few agents.

It is well known that Bayesian mechanisms generally perform better than strategy-proof ones at least since d'Aspremont and Gérard-Varet (1979), in particular for settings without transfers, because such mechanisms can use "more" information. However, this concept has not been applied to single-peaked preferences for two or more agents[1], and nothing is known about optimal Bayesian solutions. The contribution of this paper is not a characterization result, rather the simple and intuitive mechanism presented could be seen as a first step towards more general results, or at least a first glance in that direction. Indeed, the construction is itself of interest in a context where optimal mechanisms are probably very complicated. An algorithmic description is as follows. The mechanism proceeds by sequentially bisecting the type space and asks the players to which subdivision their type belongs, iterates while the messages match, and stops otherwise, selecting the splitting point as outcome. Whether this mechanism is close to the optimal one remains an open question, that is discussed after the exposition.

The next section sets up the problem and presents the dichotomic mechanism, the third section compares the outcomes with that of the best ex-post implementable solution. Finally, the last section broadens the discussion to the relevant literature while pointing out the crucial points for extensions. Most of the technicalities and proofs are relegated to the appendix.

[1] For similar models in the principal-agent framework, with only one informed party, see in particular Melumad and Shibano (1991) for a first contribution. Martimort and Semenov (2006a) use that Bayesian setting to study lobbying, but they use dominant-strategy implementation for the case of two agents.

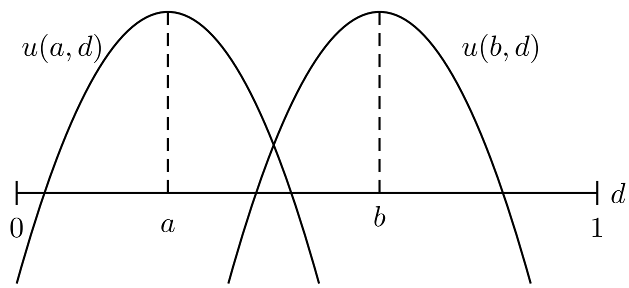

Figure 1: Setting

2 The example

2.1 setting

As already mentioned, the focus is on the simplest possible setting. Two agents must take a common decision $d \in D \equiv [0,1]$, and they have single-peaked preferences with peaks $(a,b) \in T^2 \equiv [0,1]^2$ representing their ideal decisions. They have the same utility function $u(x,d) = -(x-d)^2$ where x denotes the peak[2]. The restriction to quadratic preferences is purely for expositional clarity, and it also allows for exact calculations. The ideas presented below work for any single-peaked preferences that are symmetric and concave.

The types are privately known, and drawn i.i.d. from the uniform distribution on T. In addition, we assume comparable utilities, so that we can define a measure of ex-post surplus[3] as: $S(a,b,d) = u(a,d) + u(b,d)$. Ex-ante and interim welfare are then defined by taking the relevant expectations.
As a reference, consider the first-best case, without asymmetric information. Given the concavity[4] and the symmetry of preferences, it is straightforward to see that the surplus-maximizing rule is:

$$d^{FB}(a,b) = \frac{a+b}{2} \tag{1}$$

[2]The model could be also stated as a game, with types a, b, strategies d_a, d_b, and payoffs $U(x,d)$ if $d_a = d_b = d$, and a very bad payoff, say $-M$, with M large enough, if $d_a \neq d_b$. Then coordination is absolutely necessary, and agents "must" take the same decision. I am grateful to Jon Levin for suggesting this interpretation.

[3]This is not essential, but makes comparisons between solution concepts straightforward, and implies equal treatment of the players.

[4]In the case of linear preferences, any decision between a and b is ex-post optimal with the retained criterion.

Given its uniqueness, this rule is also optimal from an ex-ante point of view.

2.2 The dichotomic mechanism

We first state two useful definitions.

Binary Development. For any real number x between 0 and 1, write its binary development as:
$x = \sum_{k=1}^{\infty} \frac{x_k}{2^k}$ where x_k is 0 or 1, so that $(x_k) \in \{0,1\}^{\mathbb{N}^*}$.

First different digit. For two numbers a and b, define $s(a,b)$ as the smallest integer such that the corresponding binary digits of a and b differ. Formally, $s(a,b) = \inf_k \{k \mid a_k + b_k = 1\}$, and is possibly infinite[5].

We are now in position to state the definition of the dichotomic mechanism.

Definition of the direct mechanism. Let (Δ, T^2) be the following (deterministic) direct revelation mechanism, which we refer to indifferently as the "dichotomic" or "binary" mechanism.
When agents announce (a,b), let us denote by (δ_k) the binary development of $\Delta(a,b)$. The outcome is:

- $\delta_k = a_k = b_k$ for all $k < s(a,b)$
- $\delta_{s(a,b)} = 1$
- $\delta_k = 0$ for all $k > s(a,b)$

This completely defines the mechanism. Note that it would be equivalent to set $\delta_{s(a,b)} = 0$ and $\delta_k = 1$ for all $k > s(a,b)$. The outcome of the mechanism is represented in figure 2, together with the best ex-post implementable solution, anticipating a bit on the next section.

An alternative sequential definition. The mechanism is here defined in an algorithmic way, ignoring first the incentive constraints. The algorithm is initialized at $k = 1$.
Then at any stage k for which the mechanism has not stopped yet:

[5]Indeed, $s(x,x) = +\infty$, because all digits coincide. Even in such a case, all series in the sequel clearly converge.

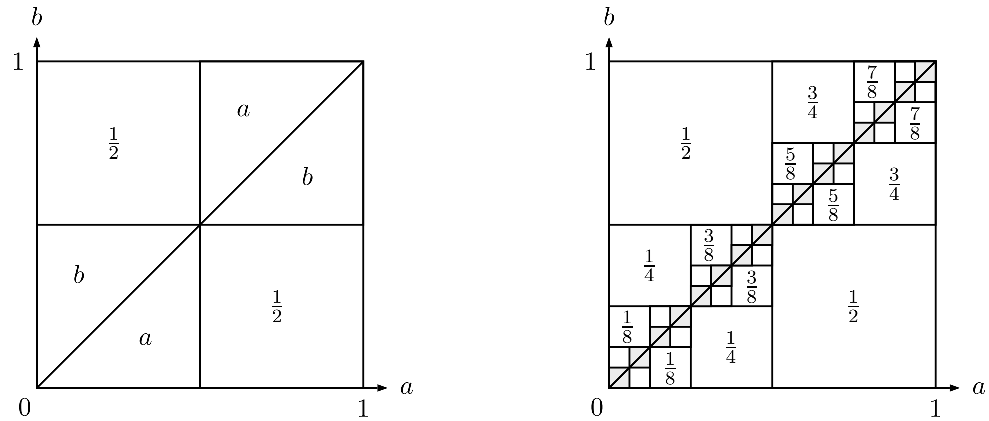

Outcome of the optimal strategy-proof mechanism (Moulin, 1980)

Outcome of the dichotomic mechanism

Figure 2: Mechanisms outcomes

- players announce simultaneously their k^{th} digits a_k and b_k;
- if $a_k = b_k$, then set $\delta_k = a_k$ and go to stage $k+1$;
- if $a_k \neq b_k$, then set $\delta_k = 1$ and $\delta_l = 0$ for all $l > k$. The algorithm stops.

This algorithm is also an extensive form representation of the communication game. In this representation, at each stage, the agents simultaneously announce a digit which means "I am in the left (0) or right (1) part of the remaining interval". If they belong to the same half, they agree on the binary digit, and continue to reveal some information in the next step, until they reach differing digits, meaning they are in different halves of the subinterval. Put differently, they exchange more and more precise information while they agree, and communication stops when they realize they have antagonist preferences, conditional on what they then know. Indeed, when the mechanism stops, this means that each player knows on which side of the own type the other one is, thus the situation then resembles a bargaining. It is easily checked that the mechanism ends in finite time when the players have different types.

2.3 Properties of the mechanism

The fundamental property of the direct dichotomic mechanism is of course incentive-compatibility.

Proposition 1 *The mechanism (Δ, T^2) has a truthful equilibrium. More precisely, truth-telling is a strict best-response to truth-telling.*

The detailed proof of this result is in the appendix. It relies mainly on the fact that a player can not influence the probability of stopping at a given stage, provided the other player is truthful.

Corollary 1 *Truth-telling by both players is a Perfect Bayesian-Nash equilibrium of the extensive form dichotomic mechanism.*

The proof of this result is virtually identical to that for the direct mechanism. The difference is that the strategies now depend on the history of messages. However, given that one player is truthful, the history of messages does not convey more information than that already contained in the fact of having reached the current stage. Therefore, arguments similar to those given in the proof of proposition 1 still apply concerning best replies against a truth-telling opponent.

On the speed of convergence. A desirable feature of the sequential mechanism is that it converges quickly. In addition, from a computational or information requirement point of view, this aspect matters for the direct mechanism as well. Therefore, we compute the expected stopping time τ, as follows. There is an unconditional probability $\frac{1}{2^k}$ that the first k digits coincide, which yields:
$\tau = \sum_{k=1}^{\infty} \frac{k}{2^k} = \sum_{k=1}^{\infty} \frac{1}{2^k} + \sum_{k=1}^{\infty} \frac{k-1}{2^k} = \sum_{k=1}^{\infty} \frac{1}{2^k} + \frac{1}{2} \sum_{k=1}^{\infty} \frac{k}{2^k} = 1 + \frac{\tau}{2}$, so that $\tau = 2$.
We can also remark that the maximal time is "small" with a high probability. Indeed, given $a \neq b$, the maximal number of consecutive matching digits starting at 1 is $\overline{K}(a,b) = \left[\frac{-\ln(|a-b|)}{\ln(2)}\right] + 1$, where $[.]$ denotes the integer part. For example, if $|a-b| \geq 0.01$, which happens more than 98 percent of the time, the maximal number of steps is only 7.

3 Comparison with ex-post implementation

3.1 Ex-post Implementation and Phantom voters

We know that there exist non-trivial ex-post implementable - or *strategy-proof* - rules. The full set of strategy-proof mechanisms has been characterized by Moulin (1980) for generic single-peaked preferences. They are the so-called "min-max" rules, that involve fixed ballots (the "phantom voters") in addition to players messages. The number of phantom voters depends on the required properties. In the present case, the preferences are completely characterized by their peak, thus *peaks-onliness* is necessarily satisfied (see Barberà (2001, section 3), Ehlers et al. (2002, section 6) and the references therein for the definitions of the properties one may want to impose on the deterministic mechanisms).

The mechanism for n voters thus works as follows: first, choose and cast $n+1$ ballots, then add the (real) votes, and take the median of all ballots as the result. If Pareto-optimality is required, then one only needs $n-1$ phantom voters. In the present Bayesian framework, we want in addition to maximize ex-ante surplus. Since Pareto-optimality is weaker than ex-post surplus maximization, itself weaker than ex-ante surplus maximization, we can here restrict to one fixed ballot. One easily obtains that with uniformly distributed types $\frac{1}{2}$ must be chosen to maximize ex-ante surplus. So, restricting to strategy-proofness, the unique optimal rule simply boils down to:

$$d^{SP}(a,b) = median(a,b,\frac{1}{2}) \tag{2}$$

Note that this mechanism is continuous, while the dichotomic one is not.

It is now possible to understand how the dichotomic mechanism improves on the min-max rule. With dominant strategy implementation, only one fixed phantom voter is used in the optimal mechanism. In turn, an infinite (but countable) set of fixed ballot is used in the binary mechanism. Intuitively, the main feature that the dichotomic mechanism adds is the possibility of *selecting the phantom voters*, by using some information on the realized preferences.

3.2 Interim comparisons

We can compare the exact solutions in terms of interim expected utility with the quadratic form. In the first-best case, one has:

$$Eu^{FB}(a) = \int_0^1 -(a - \frac{a+b}{2})^2 db = -\frac{1}{4}(a^2 - a + \frac{1}{3}) \tag{3}$$

For the phantom voters rule, we have when $a \leq \frac{1}{2}$ (and expected utility is symmetric around $\frac{1}{2}$):

$$Eu^{PH}(a) \quad = \quad a.0 + \int_a^{\frac{1}{2}} -(a-b)^2 db + \frac{1}{2}(-(a-\frac{1}{2})^2)$$

$$= -\frac{1}{2}(\frac{1}{2}-a)^2 - \frac{1}{3}(\frac{1}{2}-a)^3 \tag{4}$$

To compute expected utility in the binary mechanism, one can notice that in the truthful equilibrium, there is a probability one half of stopping at each step k, conditionally on reaching this step k. This is stated as lemma 1 in the appendix. With a slight shortcut in the notations, denote d_k the outcome if the mechanism stops at step k for a type a player. By construction, d_k does not depend on b (as already mentioned, the sentence "the mechanism stops at step k" already conveys all the necessary information on b). Thus the binary mechanism yields interim expected utility:

$$\begin{aligned}
Eu^{\Delta}(a) &= \frac{1}{2}u(a,d_1) + \frac{1}{2}\left[\frac{1}{2}u(a,d_2) + \frac{1}{2}(\frac{1}{2}u(a,d_3) + \frac{1}{2}(...)))\right] \\
&= \sum_{k=1}^{\infty} \frac{1}{2^k} u(a,d_k) \\
&= -\sum_{k=1}^{\infty}\left[\frac{1}{2^k}(\sum_{i=k}^{\infty}\frac{a_i}{2^i} - \frac{1}{2^k})^2\right] \qquad (5)
\end{aligned}$$

Despite the (very) discontinuous nature of the binary mechanism, the interim expected utility is continuous. To see this, for any $\varepsilon > 0$ define $K(\varepsilon)$ as the smallest integer such that $\frac{1}{2^{K(\varepsilon)}} < \varepsilon$. Then for all z such that $|z| < \frac{1}{2^{K(\varepsilon)}}$, $|Eu^{\Delta}(a) - Eu^{\Delta}(a+z)| < \varepsilon$. The proof of this claim is tedious although straightforward and is therefore omitted; it relies on the symmetry of the mechanism: one has to separate between two symmetric cases for each k, depending on whether the current digit match for a and $a+z$.

Surprisingly, the interim expected utility in the binary mechanism is not maximized for $\frac{1}{2}$. It is seen for example in considering the types $\frac{1}{2}$ and $\frac{3}{8}$. When $k > 3$, the terms in (5) are the same. In turn they differ for the first three stages, and we have:

$$\begin{aligned}
Eu^{\Delta}(\frac{1}{2}) - Eu^{\Delta}(\frac{3}{8}) &= \left[-0 - \frac{1}{4}(\frac{1}{4})^2 - \frac{1}{8}(\frac{1}{8})^2\right] + \left[\frac{1}{2}(\frac{1}{8})^2 + \frac{1}{4}(\frac{1}{8})^2 + 0\right] \\
&= -\frac{3}{512}
\end{aligned}$$

Figure 3 illustrates those different solutions. In particular, PV favors most the type around $\frac{1}{2}$ than Δ, while both rules give more utility to that types than the first-best. An intuitive reason for this fact is that those types have large possible deviations, while types close to the sides can not overstate their type too much (see section 4.1 for more detailed statements).

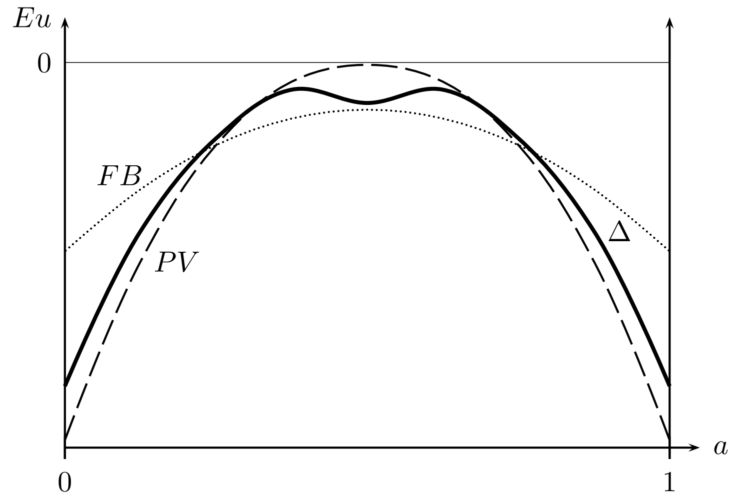

Figure 3: Interim expected utility

3.3 Ex-ante comparison

The calculations of ex-ante utility for the first-best and the phantom voters cases are straightforward. One obtains by integrating (3) and (4):

$$Eu^{FB} = -\frac{1}{24} \text{ and } Eu^{PH} = -\frac{5}{96}$$

For the binary mechanism, computations are a bit more involved. Integrating (5) yields:

$$\begin{aligned} Eu^{\Delta} &= -\int_0^1 \left(\sum_{k=1}^{\infty} \left[\frac{1}{2^k} (\sum_{i=k}^{\infty} \frac{a_i}{2^i} - \frac{1}{2^k})^2 \right] \right) da \\ &= -\sum_{k=1}^{\infty} \left[\frac{1}{2^k} \int_0^1 (\sum_{i=k}^{\infty} \frac{a_i}{2^i} - \frac{1}{2^k})^2 da \right] \end{aligned} \tag{6}$$

Note that for any $(x_j)_{1 \leq j \leq k-1}$ in $\{0,1\}^{k-1}$, we have the property:

$$\int_{\sum_{j=1}^{k-1} \frac{x_j}{2^j}}^{\sum_{j=1}^{k-1} \frac{x_j}{2^j} + \frac{1}{2^k}} (\sum_{i=k}^{\infty} \frac{a_i}{2^i} - \frac{1}{2^k})^2 da \qquad \text{does not depend on } (x_j)$$

Thus we obtain the following:

$$\int_0^1 (\sum_{i=k}^{\infty} \frac{a_i}{2^i} - \frac{1}{2^k})^2 da = 2^k \int_0^{\frac{1}{2^k}} (\sum_{i=k}^{\infty} \frac{a_i}{2^i} - \frac{1}{2^k})^2 da$$

But for the second integral, $a_i = 0$ for all $i < k$ on the integration interval, so that $\sum_{i=k}^{\infty} \frac{a_i}{2^i} = a$. Substituting in (6), one obtains:

$$Eu^{\Delta} = -\sum_{k=1}^{\infty} \int_0^{\frac{1}{2^k}} (a - \frac{1}{2^k})^2 da = -\sum_{k=1}^{\infty} \frac{1}{3} \frac{1}{8^k} = -\frac{1}{21}$$

The next table sums up those quantitative results:

Eu^{FB}	Eu^{Δ}	Eu^{PH}
$-\frac{84}{2016}$	$-\frac{96}{2016}$	$-\frac{105}{2016}$

4 Discussion and relation to the literature

4.1 On optimal Bayesian mechanisms

This subsection discusses the issue of optimal mechanisms for the simple problem considered. While for the case of one single agent Melumad and Shibano (1991) obtain a quite general characterization[6], things dramatically change for two agents. In particular, since the incentive constraint for one agent is taken in expectation with respect to the other's type, their differential approach becomes inapplicable in the present case. Even the simplest properties, such as monotonicity, can not be obtained as in their framework. Since in addition optimal mechanisms are not a priori continuously differentiable[7], and in fact not even continuous, standard tools are useless.

The first-best rule is of course not incentive-compatible, but it is still instructive to look at the deviations involved. Consider without loss of generality the corresponding direct revelation mechanism, and define the expected utility of a type a player, sending to the

[6]They show in particular that mechanisms will in general be discontinuous, even with the uniform distribution. Martimort and Semenov (2006b) give conditions under which the optimal mechanism will be continuous.

[7]Guesnerie and Laffont (1984), for the principal-agent model under adverse selection, derive conditions under which the optimal mechanism must be continuously differentiable. In such a case, one can restrict optimization to a set that we know is compact from Ascoli's Theorem, guaranteeing the existence of optimal mechanisms. In the case of single-peaked preferences, requiring differentiability may be very restrictive, as Border and Jordan (1983, p.155) point out "differentiability alone goes a long way toward eliminating nondictatorial straightforward mechanisms" (*Straightforward* is another word for *strategy-proof*, or ex-post implementable).

mechanism a message m from the set T, given that the other agent announces truthfully:

$$v(a,m) = \int_0^1 -(a - \frac{m+b}{2})^2 db \tag{7}$$

Assume $m(.)$, the mapping from types to preferred messages is almost everywhere continuously differentiable. Optimization requires that $\frac{\partial v}{\partial m}(a, m^*(a)) = 0$ at a point of differentiability, under the constraint that $0 \leq m \leq 1$, which yields:

$$m^*(a) = min(max(2a - \frac{1}{2}, 0), 1)$$

This does not match the incentive compatibility requirement, i.e. $m^*(a) = a$, except for types 0, $\frac{1}{2}$ and 1. Agents to the left of one half have incentives to understate their type, while agents to the right would overstate it. In other words there are countervailing incentives. This is the essence of the difficulty in finding the optimal mechanism (provided it exists...). In an auctions setting, Grigorieva et al. (2005)[8] construct a bisection mechanism whose central idea is nearly the same as that of the dichotomic mechanism, and they are able to show that under rather wide circumstances it is efficient and at the same time economizes on information revelation rounds[9]. Obtaining those results is possible with auctions because players never have incentives to overstate their valuations. The only deviations to prevent are understatement of the types. With single-peaked preferences, we just saw that deviations can go either way.

Another important difference is that the issue in their model, as in general with auctions, is on the interplay between participation and information revelation. Specifically, bisecting the type space allows to use mild information revelation to solve the participation problem. Once that issue is resolved, revelation is fairly easy since the procedure allocates the object at a price equal to the second-highest bid, whose revelation is costless in terms of incentives once allocation has been determined. This kind of procedure heavily relies on monotonicity with respect to type, in order to dissociate participation and revelation. On the contrary, in the present setting, revelation is the core issue while the participation issue is assumed away - which reduces the screening possibilities. Finally, the kind of mechanism presented here can not be defined independently of the priors, as is possible in the case of auctions, and the efficiency issue is thereby much more pervasive.

[8] I am grateful to the Associate Editor for having brought this reference to my knowledge.

[9] On this point, there is also a growing literature straddling mechanism design and computer sciences, see Fadel and Segal (2006) for a representative recent contribution. Complexity of the mechanism is the core concept in that perspective.

4.2 Extensions

Despite the difficulties just mentioned, the main idea behind the mechanism, namely bisecting the type space[10] in a recursive way, could probably be extended to non-uniform prior and/or non-symmetric single-peak preferences. There are a number of additional difficulties on the way. The main one for the present setting is that one would need to define at each stage a splitting point and a stopping point that may differ (while they are both the midpoint with uniform priors and symmetric peaks). Therefore it is problematic to keep sequential incentive-compatibility and Pareto-optimality at the same time. To grasp some intuition on this, assume that for a given message history, s is the splitting point and δ the stopping point, and that they are different, say $s < \delta$. If the players have been truthful up to the current stage, then for some types a, b, we have $a < s < b < \delta$. Then the stopping point is not Pareto-optimal. This is so even at the first stage as soon as the splitting and the stopping point differ. Overall, finding the general family of splitting and stopping points appears an extremely hard recursive[11] problem.

A related difficulty concerns the rather simple recursive formulae we are able to use to express expected utility. In the uniform prior case, conditional (on the type, and on the current stage) and unconditional probability of stopping in k stages from the current one never depend on types nor on current stage, which considerably simplifies the analysis. This property is lost without uniform priors, which adds a second difficulty for a generalization to non-symmetric cases.

Finally, the extension of such mechanisms to more than two agents can still be envisioned. Clearly, with bisection, there are already improvements with respect to a direct min-max mechanism. A generalization could be as follows: if all agents are in the same subinterval, iterates in that interval, otherwise, run a standard min-max mechanism calibrated on the current interval. Then, again, the recursive structure of the mechanism can be applied. However, finding the divisions in that case would be subject to the same complex recursive problem as above, in an even more constrained version.

[10] It is possible to show that divisions in more than two sub-intervals can not satisfy incentive compatibility.

[11] Considering the direct mechanism embeds the equivalent problem in the incentive constraint, but considering recursivity - therefore the sequential version - is again more illustrative.

4.3 A cheap talk perspective

Although a mechanism design approach is used here, there is a connection with long cheap-talk (Aumann and Hart, 2003; Amitaï, 1996). A bit more structure is needed to sketch, very informally, the parallel. Assume that, after some communication has taken place through long cheap talk, the decision is determined by a bargaining game, whose outcome depends only on the information that has been revealed. Thus we extend a bargaining-like situation by adding a communication phase, in the spirit of Farrell and Gibbons (1989). Assume in addition equal bargaining powers. Then in the communication phase, information can be exchanged as in the binary mechanism. Once the players discover they are in different (sub-)intervals, the bargaining process conditioned only on public information defines the same point as the dichotomic mechanism. Overall, the cheap talk allows first to identify the common interest part of the preferences, then it is common knowledge that the players are in a situation of conflict, whose outcome is defined through a Nash-like bargaining. In that case, both players would indeed prefer to exchange some information through bisection cheap talk.

A last point deserving attention is the link between the revelation dynamics described here and the jointly controlled lotteries used in Aumann and Hart's characterization. The canonical representation of the cheap-talk phase they introduce is a process alternating joint lotteries periods and revelation periods. A joint controlled lottery in that characterization is a way of generating randomness in the communication stage. On the contrary here, the messages convey information at each stage, but thanks to the uniform distribution, both messages of one player are equally likely in the eyes of the other player. In the end, continuation is thus determined jointly, and no player can alone influence, as in a joint controlled lottery. But, once again, these messages are also informative.

Crawford and Sobel (1982) pioneered the interest in cheap talk by showing that even absent mechanism (a commitment device) and monetary incentives, some information could still be revealed between parties with conflicting preferences. In their model, the amount of information revealed in equilibrium increases as the conflict of interest decreases. While the bias, as measured by the constant difference between the peaks, is perfectly known in their model, here it is not. The two players may have the same peaks, as well as completely antagonistic preferences, but they do not know by how much their preferences differ. Nonetheless, the amount of information revealed in the dichotomic mechanism (or in the suggested adaptation to long cheap talk) is also increasing as the *realized* conflict of interest decreases. Indeed, information revelation in equilibrium is more precise the closer (a, b) is to the diagonal. Therefore, the intuition that closer preferences leads to more information revelation is preserved when the (endogenous) intensity of conflict is not known ex-ante.

A Incentive Compatibility of the dichotomic mechanism

A *pure strategy* in the mechanism (Δ, T^2) is a mapping $\sigma : T \rightarrow T$, that we can equivalently define as a mapping from $\{0,1\}^{\mathbb{N}^*}$ to itself, using binary developments. We allow for mixed strategies, but seldom need them.

The *truth-telling strategy*, denoted TT, is the identity function on T.

A type $x \in T$ is said to be *finite* if: $\exists n \in \mathbb{N}^*$ such that $x_k = x_n$ for all $k > n$. Thus a finite type is such that the digits of its binary development are identical after some stage. A *n-finite type* is a finite type such that $x_{n-1} = 1 - x_n$ and $x_k = x_n$ for all $k > n$, that is, n is the first digit where the sequence consisting of only 0's or 1's begins.

A finite type has two possible binary developments, since

$$(x_1, .., x_{n-1}, x_n, .., x_n, ...) \equiv (x_1, .., 1 - x_{n-1}, 1 - x_n, .., 1 - x_n, ...)$$

In that case, we consider that both messages are truthful, as well as mixing between them, to dissipate the (minor) ambiguity in defining the strategy (anyway, these types have zero measure in T).

Since the goal is to prove that truth-telling by both players is an equilibrium, we assume in all what follows that one player tells the truth (uses strategy TT) and consider the point of view of the other player, whose type is some a. We use reduced notations whenever this is not confusing. The aim is to prove that TT (on the part of type a) is a strict best-response to TT.

We first need a proper expression of the payoffs.

Lemma 1 *When the other player uses strategy TT, the unconditional probability of stopping at a given stage k does not depend on the strategy. The payoff conditional on stopping at a* *given stage k depends only on a, k and the strategy. The expected utility of a pure strategy σ against TT is:*

$$E[u^{\Delta}(\sigma, TT)|a] = -\sum_{k=1}^{\infty}\left[\frac{1}{2^k}(a - \sum_{i=1}^{k-1}\frac{\sigma_i(a)}{2^i} - \frac{1}{2^k})^2\right]$$

Proof. Fix a, denote S a mixed strategy and (s_i) the binary development of the realization of S. Let b denote the type of the other player. Given that the other player uses TT, the probability that the i^th digits match is independent on s_i, it is $\frac{1}{2}$, if $s_i = 0$ or if $s_i = 1$, thus also $\frac{1}{2}$ overall. The probability of stopping at stage k is the probability that all the first $k-1$ digits match but not the k^{th} one. Given the uniform prior, the digits b_i are independent, thus the unconditional probability of stopping is simply $\frac{1}{2^k}$.
When the mechanism stops at a given stage k, by definition the first $k-1$ digits match, and $\delta_i(\sigma(a), b) = \sigma_i(a)$ for $i < k$, and the remaining digits do not depend on the types. The outcome conditional on stopping at stage k when σ is used is thus:

$$d^k(\sigma, TT|a) = \sum_{i=1}^{k-1}\frac{\sigma(a)_i}{2^i} + \frac{1}{2^k}$$

The expected utility can then be written in expectation over the stopping stage k, which ends the proof. ■

The next lemma inquires what kind of non-truthful strategies we have to consider.

Lemma 2 *If the other player is truthful, then the best pure strategies of a player with type a if he lies at least for one digit have the form: $(a_1, ..., a_{t-1}, 1 - a_t, a_t, .., a_t, ...) \equiv L_t(a)$ for some t.*

Proof. Fix a, denote σ a pure strategy and (σ_i) the binary development of $\sigma(a)$. A mixed strategy is defined the usual way, but as will be clear, they are needed only in very special case. Without loss of generality, assume that up to some t, possibly equal to 1, the strategy is truthful, i.e. $\sigma_i = a_i$ for all $i < t$, but that the t^{th} digit is wrong (the "first" lie). Any pure non-truthful strategy always takes this form.

We can now express the outcome of the mechanism conditional on stopping at stage k:

$$d^k(\sigma(a)) = \sum_{i=1}^{t-1}\frac{a_i}{2^i} + \frac{1 - a_t}{2^t} + \sum_{i=t+1}^{k-1}\frac{\sigma_i}{2^i} + \frac{1}{2^k} \tag{8}$$

with the usual convention that the second sum is null when $k-1 < t+1$. Because b uses TT, the expected payoff of σ, denoted $U_\sigma(a)$, is given from the first lemma by:

$$U_\sigma(a) = EU^\Delta(\sigma, TT|a) = E_k\left[u\left(a, d^k(\sigma(a))\right)\right] = -\sum_{k=1}^{\infty} \frac{1}{2^k}\left[a - d^k(\sigma(a))\right]^2$$

Given (8), $U_\sigma(a)$ can be rewritten as:

$$U_\sigma(a) = -\sum_{k=1}^{\infty}\left[\frac{1}{2^k}\left(\sum_{i=t}^{\infty}\frac{a_i}{2^i} - \frac{1-a_t}{2^t} - \sum_{i=t+1}^{k-1}\frac{\sigma_i}{2^i} - \frac{1}{2^k}\right)^2\right] \tag{9}$$

Depending on the relative position of the type to the division of the subinterval, two cases are possible: $a \geq d^t(\sigma(a))$, and $a \leq d^t(\sigma(a))$. Since they are symmetric, we consider only the first case, in which $a_t = 1$. Equation (9) writes:

$$U_\sigma(a_1, .., a_{t-1}, 1, a_{t+1}, ...) = -\sum_{k=1}^{\infty}\left[\frac{1}{2^k}\left(\sum_{i=t+1}^{\infty}\frac{a_i}{2^i} + \frac{1}{2^t} - \sum_{i=t+1}^{k-1}\frac{\sigma_i}{2^i} - \frac{1}{2^k}\right)^2\right] \tag{10}$$

Now, remark that

$$\sum_{i=t+1}^{k-1}\frac{\sigma_i}{2^i} + \frac{1}{2^k} \leq \sum_{i=t+1}^{k-1}\frac{1}{2^i} + \frac{1}{2^k} = \frac{1}{2^t} - \frac{1}{2^{k-1}} + \frac{1}{2^k} = \frac{1}{2^t} - \frac{1}{2^k}$$

From which we deduce that in (10), each of the terms in the parentheses is positive for $k > t$. Optimizing with respect to $(\sigma_i)_{i \geq t+1}$ thus requires to minimize each of these terms, which is done by setting $\sigma_i = 1$ for all $i \geq t+1$. The case $a_t = 0$ being symmetric, we have the desired conclusion. ■

It is clear that the players could use more sophisticated lying strategy by mixing on the set of L_t strategies. However, we show now that any pure strategy L_t is strictly dominated by TT, and thus no mixing can involve one.

We have to compare what a player can get in the mechanism by telling the truth or by using a strategy Ł$_t$. Fix a, t and assume $a_t = 1$ (once again, the case $a_t = 0$ is symmetric). For $k \leq t$, $u\left(a, d^k(L_t(a))\right) = u\left(a, d^k(TT(a))\right)$ by definition of L_t, and for $k > t$, we have:

$$u\left(a, d^k(L_t(a))\right) = -\left(\sum_{i=t}^{\infty}\frac{a_i}{2^i} + \frac{1}{2^k}\right)^2$$

while:

$$u\left(a, d^k(TT(a))\right) = -\left(\sum_{i=k}^{\infty} \frac{a_i}{2^i} - \frac{1}{2^k}\right)^2$$

If there exists some $z > t$ such that $a_z = 1$, then $u\left(a, d^k(TT(a))\right) > u\left(a, d^k(L_t(a))\right)$, from which we can conclude that TT strictly dominates L_t.

If $a_k = 0$ for all $k > t$, then a is a t-finite type (recall that $a_t = 1$) and then, by convention, L_t is also considered truthful, because $L_t(a) \equiv TT(a)$. Truth-telling is thus the only (strictly) preferred strategy.

Overall we have proved that truth-telling is a strict best-response to truth-telling, meaning that the mechanism is incentive-compatible.

References

Amitaï, M., 1996. Cheap-Talk with Incomplete Information on Both Sides. Discussion Paper #90, Center for Rationality, Hebrew University of Jerusalem.

Aumann, R., Hart, S., 2003. Long Cheap Talk. Econometrica 71, 1619-1660.

d'Aspremont, C., Gérard-Varet, L.-G., 1979. Incentives and Incomplete Information. Journal of Public Economics 11, 25-45.

Barberà, S., 2001. An Introduction to Strategy-proof Social Choice Functions. Social Choice and Welfare 18, 619-653.

Border, K., Jordan, J., 1983. Straightforward Elections, Unanimity and Phantom Voters. Review of Economic Studies 50, 153-170.

Crawford, V., Sobel, J., 1982. Strategic Information Transmission. Econometrica 50, 607-637.

Ehlers, L., Peters, H., Storcken, T., 2002. Strategy-Proof Probabilistic Decision Schemes For One-Dimensional Single-Peaked Preferences. Journal of Economic Theory 105, 408-434.

Fadel, R., Segal, I., 2006. The Communication Cost of Selfishness, mimeo Stanford University.

Farrell, J., Gibbons, R., 1989. Cheap Talk Can Matter in Bargaining. Journal of Economic Theory 48, 221-237.

Grigorieva, E., Herings, J.J., Müller, R., Vermeulen, D., 2006. The Private Value Single Item Bisection Auction, Economic Theory, forthcoming.

Guesnerie, R., Laffont, J.-J., 1984. A Complete Solution to a Class of Principal-Agent Problems with an Application to the Control of a Self-managed Firm. Journal of Public Economics 25(3), 329-349.

Martimort, D., Semenov, A., 2006a. Competition vs. Collusion in Legislative Politics: A Mechanism Design Approach. mimeo IDEI.

Martimort, D., Semenov, A., 2006b. Coontinuity in Mechanism Design without Transfers, Economics Letters, forthcoming.

Melumad, N., Shibano, T., 1991. Communication in Settings with no Transfers, Rand Journal of Economics 22, 173-198.

Moulin, H., 1980. On Strategy-Proofness and Single-Peakedness. Public Choice 35, 437-455.