Essays on Strategic Voting and Mechanism Design

Vinayak Tripathi

A Dissertation
Presented to the Faculty
of Princeton University
in Candidacy for the Degree
of Doctor of Philosophy

Recommended for Acceptance
by the Department of
Economics
Adviser: Dilip Abreu

June 2008

© Copyright by Vinayak Tripathi, 2008.

All Rights Reserved

Abstract

In the first chapter, we consider the problem of first-best implementation in dynamic, interdependent-value environments where agents have quasi-linear preferences and private types are correlated. Our first candidate solution (I.M.) relies on methods typically seen in games with imperfect monitoring. We derive conditions under which this approach gives the planner the flexibility to implement efficient mechanisms that are ex-post budget balanced and collusion proof. If agents are required to make one final payment after they exit the system, the I.M. mechanism can also satisfy individual rationality and adapt to finite horizon settings. When such payments cannot be enforced, but a cyclic monotonicity condition is satisfied, the planner may adopt our second mechanism (P.A.). This alternative approach employs a technique originally developed for surplus extraction in static auctions. It allows the planner to align every agent's preferences with his own and implement socially optimal outcomes while satisfying ex-ante budget balance and individual rationality. Both mechanisms permit the planner to implement any sharing rule when redistributing surplus.

In the second chapter (co-authored with Sambuddha Ghosh), we examine a paradox of strategic voting. Consider an electorate whose individual rankings of alternative policies may change between the time they vote for a candidate and the date a policy is implemented. Rankings may change following common or idiosyncratic shocks. Voters choose, via a simple majority election, between a candidate who is committed to one alternative, and an unbiased candidate who implements the ex-post optimum. We show that, even when idiosyncratic shocks lead to slight changes in rankings, the unique symmetric pure strategy Nash equilibrium often entails strategic voters choosing the committed candidate. The resulting welfare loss increases with the likelihood of a common shock.

In the third chapter (co-authored with Konstantinos Rokas), we examine an electoral framework in which agents can strategically choose when to vote and can observe the votes of those who precede them. This is in contrast to the literature on Condorcet's Jury Theorem which has limited its attention to simultaneous and sequential voting games. We show that in a common value setting, such an election aggregates information efficiently regardless of the voting rule or the size of the electorate. The proposed election format ensures that the Jury Theorem obtains for an electorate with differentially informed agents. Under more restrictive conditions, this framework can also produce the efficient outcome in environments where voters have conflicting interests or the size of the electorate is uncertain. Finally, we examine the performance of the election when private signals are multidimensional.

Acknowledgements

I am deeply indebted to my adviser, Prof. Dilip Abreu, for his guidance and patience these last few years. I also wish to express my sincerest gratitude to Profs. Markus Brunnermeier, Stephen Morris, Wolfgang Pesendorfer, Hyun Song Shin, and Wei Xiong for their invaluable advice. Many thanks to the Fellowship of Woodrow Wilson Scholars and Prof. John Darley for having supported my work intellectually and financially. I am honored to have had the extraordinary privilege of being a graduate student at Princeton University.

I must also take this rare opportunity to thank the people who have been instrumental in bringing me to this juncture - My grandmother, Sanjeevi Varanasi, who supervised my education at the very beginning; My mother, Suryakanthi, and my father, Amitava, for their unshakable support and indomitable faith in me; my teacher, Dr. Siddiqui, for awakening my intellectual curiosity; my friend, Dr. Diego Hidalgo whose generosity made everything possible; my first employer and guide, Dr. Roger Peterson, who, in singular fashion, steered my career back on track on two distinct occasions; my uncle, Rajaram Varanasi, for his advice during a particularly challenging period; my friend and classmate, Srihari Adireddy, whose sound judgment and balanced perspective have always helped me find my bearings; my friend and colleague, Eugene Visotsky, who taught me how to think and write critically; my friend and co-author, Sambuddha Ghosh, for his admirable persistence on our joint work and the conversations we shared on every conceivable subject; my friend and co-author, Konstantinos Rokas whose ingenuity brought our joint work to life; my friends and classmates, Kunal Dasgupta, Prashanth Hande, Ashok Mantravadi, and Konstantin Milbradt whose acquaintance has greatly enriched my life; and finally, my fiancee, Ty Wilde, with whom I shared my harshest and happiest moments at Princeton.

To my parents.

Contents

[1] Co-authored with Sambuddha Ghosh.

[2] Co-authored with Konstantinos Rokas.

Chapter 1

Correlation in Dynamic Mechanism Design

1.1 Introduction

Multi-agent mechanism design has long been a subject of interest for economists, but the study of dynamic environments has only recently come to the fore. Dynamics can be critical in a variety of settings. Governments repeatedly elicit the time-varying preferences of interested parties to optimally adjust the provision of a public good. Firms regularly consult with their financiers when making sequential investment decisions. Project managers periodically conduct peer reviews to collect information necessary to incentivize effort in their teams. Legislative committees are interminably lobbied by special interest groups from whom they try to glean all information critical for efficient policy-making. Market participants benefit from sharing trade secrets with their long-term strategic partners, and countries often find themselves communicating private information when negotiating multi-lateral agreements fraught with externalities. The primary motivation for this work is therefore to examine a problem that appears in a multitude of important economic settings.

Recent papers by Athey and Segal (henceforth, A&S) [4], Bergemann and Valimaki (henceforth, B&V) [10] and Cavallo, Parkes and Singh [13] examine the mechanism design problem when player types evolve independently and agents have private values. Under these conditions, B&V build on the celebrated Vickrey-Clark-Groves mechanism to construct an efficient direct mechanism that satisfies dynamic individual rationality. A&S focus instead on ex-post budget-balanced implementation, extending the results of d'Aspremont and Gerard-Varet [17] to this setting. In contrast, we construct mechanisms that implement the first-best allocation in Markovian, interdependent value environments where private types are correlated serially and across agents. Table 1. below summarizes how our mechanisms differ from those proposed previously.

In private value settings with independently evolving types, the planner can accurately forecast the future types of agent j without relying on agent i's reports. In other words, agent i and the planner always share common beliefs about agent j. In previous work, this property is instrumental in designing the transfer payments that incentivize truth-telling. The mechanisms of A&S and B&V cannot implement the socially optimal allocation when interested parties possess private information about each other's types that is not available to the planner. This is an important limitation as one may easily imagine settings where special interest groups understand the needs of their rivals better than the planner.

	Values	Types	I.C.*	Budget Balance*	I.R.
A&S	Private	Independent Serially Corr.	Interim	Ex-Post	Only for large δ
B&V	Private	Independent Serially Corr.	Ex-Post	None	Ex-Post
I.M.	Any	Cross-Corr. Markov	Ex-Post	Ex-Post†	Ex-Ante
P.A‡.	Any	Cross-Corr. Markov	Ex-Post	Ex-Ante	Ex-Ante

Table 1. I.M.- Imperfect Monitoring, P.A.- Preference Aligning, *= within-period

†=under strict identifiability, ‡=requires cyclic monotonicity

Our first approach, the imperfect monitoring mechanism (henceforth I.M.), exploits the fact that the future report of agent j serves as an imperfect signal of whether agent i's current report was truthful. This feature along with the planner's ability to commit intertemporally to state-contingent continuation values permits him to incentivize truth-telling in agent i via suitably defined lotteries over future promised values. If an identifiability condition reminiscent of those developed in the literature on repeated games with imperfect monitoring is satisfied, this mechanism can implement efficient outcomes in correlated, interdependent-value environments.

The I.M. approach gives the planner several degrees of freedom. Under a stricter version of the identifiability condition, this flexibility allows I.M. to satisfy a strong version of budget balance (see Table 1). We can also derive conditions under which the I.M. approach can implement a collusion-proof mechanism; a quality that is important in many contexts as coalitions and cartels often distort outcomes.

The I.M. mechanism has two limitations. First, it fails to satisfy dynamic individual rationality. Specifically, it does not rule out the possibility of an agent profitably sending a false report and exitting the mechanism immediately thereafter to avoid the consequences. Second, due to the nature of its construction, the I.M. does not adapt naturally to finite horizon settings. Since incentives for truth-telling come from the threat of low continuation values, the I.M. approach fails to incentivize agents in the final period of a finite mechanism. Each of these problems can be dealt with if the planner requires agents to make one final report and payment after they exit the mechanism. A full rank condition that is strictly more restrictive than the identifiability condition is also necessary for this approach to be feasible.

Alternatively, one may consider adopting the Preference Aligning (or P.A.) approach which permits the planner to implement first-best outcomes in environments that satisfy cyclic monotonicity. The P.A. mechanism satisfies dynamic ex-ante individual rationality and adapts easily to finite horizon settings; it does so even in the absence of payments and reports from agents who have exitted the system. This approach employs a technique originally developed by Cremer and McLean (C&M) [16] to extract surplus in static auctions; it allows the planner to align an agent's preferences over future evolution with his own (hence, the name). Using C&M's construction, one can exploit the correlation in private signals to transform every agent's expected continuation value into an affine transformation of the social welfare function. When combined with transfers that exploit cyclic monotonicity to neutralize conflicting interests in the stage game, the C&M construction yields a mechanism that induces truth-telling in the dynamic setting.

Both our mechanisms permit the planner to redistribute excess surplus according to any sharing rule. This is in particular contrast to the work of B&V who specify conditions under which all efficient direct mechanisms in independent, private-value environments distribute surplus according to the marginal contributions of the agents. Indeed, our mechanisms can also be interpreted as optimal contracts for a self-interested principal seeking to extract all the surplus for himself while satisfying individual rationality.

One issue worth acknowledging is that the P.A. approach, like the auction proposed by Cremer and Mclean, is vulnerable to a recent critique detailed in the work of Neeman [37]. The P.A. mechanism relies critically on the assumption that an agent's beliefs uniquely pin down her preferences. While some consider this an irretrievable failing, the assumption is well suited for a number of economic settings and is commonly seen in applied work; for instance, the use of affiliated private values is wide-spread in auction theory. Furthermore, as discussed above, the P.A. approach offers unique insights about this problem that are worth documenting. The I.M. based mechanism does not fall prey to Neeman's critique.

Turning to the issue of incentivization in dynamic environments, we should emphasize that serial correlation in player types severely complicates the planner's problem. It is no longer sufficient for the planner to provide incentives at time t for each of player i's period-t types, presuming all along that the player's past reports have been truthful. Doing so may eliminate profitable one-shot deviations, but leaves the door open for the player to construct profitable multi-period deviations that involve distorting the planner's beliefs about the true state over a period of time. We address this problem within the framework of Markov Decision Processes. With persistent types, the incentive problem also suffers from a "curse of dimensionality" since the planner must take every possible private history into consideration when designing transfer schemes. The Markov framework, which admits many economically interesting environments, allows us to avoid this difficulty without skirting the issue of multi-period deviations. With a finite state space and Markovian evolution, the planner can implement truth-telling by incentivizing all private histories. Agents cannot then manipulate a planner's beliefs via elaborate "double" deviations. The mechanisms we construct may be extended, within limits, to *generalized* Markov processes where conditional on a history of fixed and finite length, the distribution over future states is independent of the past.

The remainder of this paper is organized as follows. We begin with a short review of other relevant literature before moving on to a description of the model and the Sequential-Groves mechanism. Next, we formally define our notions of budget balance and individual rationality. We then construct the I.M. based mechanism and derive conditions under which it satisfies individual rationality, ex-post budget balance and collusion proofness. Finally, we examine the P.A. mechanism and compare the two approaches before concluding.

1.2 Related Literature

The problem of eliminating "double" deviations has been studied fairly extensively in the context of dynamic contracting models with hidden states and unobserved actions (see Kocherlakota [33]); typically the adverse selection takes the form of hidden savings or unobservable cash flows. The possibility of profitable multi-period deviations arises because the planner and agent do not share common beliefs about the future state. A number of techniques, including for instance the first order method(see Williams [44], [45]), have been developed to deal with persistent private types. The imperfect monitoring based approach that we examine is similar in spirit to a technique first proposed (to the best of this author's knowledge) by Fernandes and Phelan and later explored in Doepke & Townsend [19]. They too address the incentive problem by employing a vector of utility promises for every possible private history.

Another literature worth pointing out in the context of this paper is the work on common agency first developed by Bernheim and Whinston [11]. In a static setting where all players are symmetrically informed, they show the existence of an efficient "truthful" equilibrium where players bid over a menu of allocations and the planner simply implements that which maximizes his revenue. This literature is at the far end of the spectrum from the work by A&S and B&V who assume that players know nothing about each other's private values. Our work considers an intermediate problem where agents possess correlated private information. Bergemann and Valimaki [9] examine a dynamic version of the common agency setup and conclude that a Markov equilibrium in truthful strategies is unique if and only if every principal receives her marginal contribution to the social value as a payoff. In contrast, any division of the surplus is feasible with our environment.

The work of Aoyagi [2] on collusion in repeated auctions is also indirectly related to this paper. Aoyagi considers environments in which allocations do not influence future evolution, but private types may be correlated within a period. He derives conditions under which bidders may communicate with one another and sustain an "efficient" coalition that can extract the entire surplus from the seller. This is done by coordinating on which bidder enters the auction and pays the reserve price. Coalitions are sustained not with side contracts, but via the threat of punishment. In contrast, we consider environments where side contracts are permitted and derive conditions under which there exist direct mechanisms that allow a seller to allocate efficiently, redistribute surplus among agents and remain immune to collusive behavior while doing so.

1.3 The Model

Consider a discrete-time infinite horizon framework with $N \geq 2$ players and a planner. At time t, each agent i receives a private signal $\theta_{i,t}$ that can take one of the $k_i \geq 2$ values in the set Θ_i. We refer to the signal vector $\theta_t = [\theta_{1,t}, \theta_{2,t}, \ldots, \theta_{N,t}]'$ as the state at time t. We use the notation x^t to denote the history of the variable x - i.e. $x^t = \{x_0, x_1, \ldots x_t\}$. Each period, the planner takes an action $a_t \in A$; this action may influence the evolution of the state which follows a law of motion denoted by $\Pi(\theta_{t+1}|\theta^t, a^t)$ [1]. The planner also specifies a transfer $p_{i,t}$ for each player. The action and transfer determine each agent's total flow utility which we will assume takes the quasi-linear form

$$u_i = g_i\left(a^t, \theta^t\right) + p_{i,t}$$

The direct flow utility functions, $g_i(a^t, \theta^t)$, are bounded; i.e.$g_i(a^t, \theta^t) \in [-\frac{\Delta}{2}, \frac{\Delta}{2}]$. Quasi linearity is a restrictive assumption, but one that is quite standard in mechanism design and public economics. In the absence of quasi-linearity, as in the case where agents are risk averse, first best implementation is usually no longer possible.

The timing of this mechanism is as follows - every period, agent i observes her private signal $\theta_{i,t}$ and makes an announcement $\widehat{\theta}_{i,t}$. The planner, who we assume can credibly commit to an action rule and transfer scheme, determines a_t and the vector of transfers $p_t = [p_{1,t}, p_{2,t}, \ldots, p_{N,t}]'$ on the basis of these announcements. The public history can then be denoted $H_t = \{\widehat{\theta}^t, a^t, p^t\}$. An agent's private history also includes the sequence of her private signals; we denote agent i's information set $H_{i,t} = \{\theta_i^t, \widehat{\theta}^t, a^t, p^t\}$.

Finally, we will use the notation $[x, y]$ when appending column vectors x and y to form a matrix and $x[i]$ to denote the component of vector x corresponding to argument i.

1.3.1 The Planner's Problem

Suppose that the planner commits to a mechanism $M = \{a_t(H_{t-1}, \widehat{\theta}_t), p_t(H_{t-1}, \widehat{\theta}_t)\}_{t=0}^{\infty}$ and the agents respond in equilibrium by following reporting strategies denoted by $\sigma_i : H_{i,t-1} \times \theta_{i,t} \to \Theta_i$. We provide a formal definition of equilibrium once all notation is in place. We adopt the standard utilitarian notion of social welfare which under

[1] Any action plan $\{a_t\}$ uniquely specifies a probability measure over $\Theta = \prod \Theta_i$. The Tulcea product theorem gives us the existence and uniqueness of this measure. We refer the reader to A&S for a more detailed discussion.

quasi-linear utilities is given by the expected sum of discounted utilities over all agents in the mechanism. That is, the social welfare at the beginning of period t can be written as

$$W_t(H_{t-1}; M, \{\sigma\}_i) = E\left[\sum_{s=t}^{\infty} \delta^{s-t} \sum_{i=1}^{N} g_i\left(a^s, \widehat{\theta}^s\right) \middle| H_{t-1}; M, \{\sigma\}_i\right]$$

As the definition suggests, $W_t(H_{t-1}; M, \{\sigma\}_i)$ specifies welfare before the state at time t has been reported. Agent i's value from this mechanism at the beginning of period t (prior to her receiving a private signal) is denoted $V_{i,t}(H_{i,t-1}; M, [\sigma_i, \sigma_{-i}])$. For notational simplicity, we will drop the dependence on M, and σ in the remainder of this paper.

It is the planner's objective to implement the *first-best* (or surplus maximizing) action plan. In a complete information setting, the first-best action plan $\{a_t^*\}_{t=0}^{\infty}$ is defined in terms of the true evolution of the state as shown below

$$\{a_t^*\}_{t=0}^{\infty} = \arg\max_{\{a\}_{t=0}^{\infty}} \sum_{i=1}^{N} g_i\left(a^t, \theta^t\right) + E\left[\sum_{s=t+1}^{\infty} \delta^{s-t} \sum_{i=1}^{N} g_i\left(a^s, \theta^s\right) \middle| \theta^t\right]$$

$W_t^*(H_{t-1})$ and $W_t^*(H_{t-1}, \theta_t)$ denote the ex-ante and ex-post period-t social welfare under the first-best policy, respectively. We will restrict our attention to truth-telling mechanisms. The planner's problem is to incentivize truthful revelation by committing to a suitably defined transfer scheme while simultaneously implementing $\{a_t^*\}_{t=0}^{\infty}$. We will henceforth call such mechanisms *efficient*.

Note that, consistent with the literature, we have not included transfer payments in the definition of social welfare. In other words, our definition of efficiency does not aspire to any distributive goal. Indeed, as it stands, an efficient mechanism M^* may result in the planner extracting the entire surplus for himself. Typically, such mechanisms will not however satisfy a budget balancing condition. In keeping with the literature, we will continue to maintain a distinction between efficiency and budget balance. A formal discussion of our notion of budget balance appears below.

1.3.2 Equilibrium

We adopt the notion of equilibrium defined in B&V [10]. Consider a mechanism M and suppose that all agents in the set $-i$ reveal their signals truthfully. There are no profitable one-shot deviations from truth-telling for agent i if in every period t,

$$g_i\left(a^t,\theta^t\right)+p_{i,t}\left(\theta^t\right)+\delta V_{i,t+1}(H_{i,t})\geq g_i\left(a^t(\widetilde{\theta}^t),\theta^t\right)+p_{i,t}\left(\widetilde{\theta}^t\right)+\delta V_{i,t+1}(\widetilde{H}_{i,t})$$

for all $\widetilde{\theta}_i$ such that $\widetilde{\theta}^t=[\theta^{t-1},\theta_{-it},\widetilde{\theta}_i]$. We are using a fairly strong notion of incentive compatibility. Agent i prefers to tell the truth in every period even after she observes the announcements of the other players. This is referred to as *within-period ex-post incentive compatibility.*

In environments where there is serial correlation in the private signals of individual players, the absence of profitable one shot deviations does not immediately rule out the possibility of more complex profitable deviations. An agent may be able to construct profitable multi-period deviations that involve manipulating the beliefs of the planner through a sequence of false reports. We will have to be wary about this complication when we construct efficient mechanisms later in the paper.

One final remark about ex-post implementation is in order. Under this form of incentive compatibility, the model can be augmented to also include a moral hazard component. That is, suppose executing the first best action plan required each player to take an unobservable action, $a^*_{i,t}$, upon the request of the planner. Suppose further, that for every action a planner may recommend, there exists a report that an agent can make which will induce that action; in other words the first best policy is a surjective mapping from agent i's reports to allocations. With ex-post incentive compatibility, it follows immediately then that agents are also obedient; they see no profit in deviating from the planner's recommended action. Under a weaker form of incentive constraints, say interim, this would not be true since the planner's recommended action would convey information to the player about the private signals of others and obedience would not be guaranteed.

1.3.3 The Sequential Groves Mechanism

In keeping with the work of Cavallo, Parkes and Singh [13], we refer to the simplest efficient mechanism as the Sequential Groves Mechanism (SGM)[2]. This mechanism implements the first-best action plan under the very general conditions described above. According to this mechanism, agent i's transfer at time t is given by

$$p_{i,t}\left(\widehat{\theta}^t\right) = \sum_{j\neq i} g_j\left(a^*\left(\widehat{\theta}^t\right), \widehat{\theta}^t\right)$$

Under this mechanism, if agents have private values (i.e. $g_i\left(a^t,\theta^t\right) = g_i\left(a^t,\theta_i^t\right)$) and all players in $-i$ are reporting truthfully, it is player i's best response to do the same. Intuitively, this is because the transfer scheme eliminates any conflict of interest between player i and the planner. By reporting the truth, agent i induces the planner to take the very action that the she would take if she were in control. We omit a formal proof of this argument as it has appeared previously in the literature.

With general interdependent values, inducing truthful revelation takes a little more work. The planner must implement the transfer by asking each agent to report her direct flow utility from the action and then paying each agent a transfer equal to the sum of everyone else's utilities. Implementing transfers in this manner was first suggested by Mezzetti (see [35], [36]), in his work on static mechanisms. It is worth noting however that his equilibrium concept has been criticized for assuming that when a player is indifferent, she breaks ties in favor of reporting truthfully.

[2] A&S refer to this as the Team Mechanism.

Multi-period deviations do not exist in this mechanism as the transfers do not involve forming expectations about the future. We will see below that implementing efficient dynamic mechanisms will typically require the planner to use beliefs about the distribution over future states to compute continuation values for the players.

The SGM has two important shortcomings. First, it is not budget balanced. In order to implement this efficient mechanism, the planner has to infuse $(n-1)\sum_j g_j\left(a^*\left(\widehat{\theta}^t\right),\widehat{\theta}^t\right)$ units of wealth into the mechanism every period. Second, it does not satisfy participation constraints. If an agent can guarantee herself a higher value outside the mechanism at any time t, she might prefer to leave. Our focus in this paper is to construct efficient mechanisms that satisfy these properties which are discussed at length in the next section.

1.3.4 Budget Balance and Individual Rationality

Before we provide formal definitions of these criteria, a brief discussion is in order. The literature on static mechanisms typically defines budget balance in its strictest form as being satisfied when transfers sum to zero in every state, $\sum p_i(\widehat{\theta}) = 0$ for all $\widehat{\theta}$. This definition has been the standard in this literature going back to the work of d'Aspremont and Gerard-Varet [17]. Consider for a moment a weaker definition of budget balance; one which only requires that the planner not infuse any of his own wealth into the mechanism; i.e. $\sum p_i(\widehat{\theta}) < 0$. If direct flow utilities are bounded above, a planner can trivially meet this weaker requirement by simply charging every agent a sufficiently large membership fee at time $t = 0$ to finance the operation. This weaker form of budget balance must therefore be coupled with an individual rationality constraint to be meaningful. On the flip side, a mechanism can trivially satisfy individual rationality if the planner simply pays a large, flat participation transfer to every agent to ensure that reservation levels of utility are met, but this would violate budget balance. It is in this sense that budget balance and individual rationality are inextricably linked. When the budget balance constraint takes the form of an equality, however, one can consider it in isolation.

The strictest form of dynamic budget balance, *within period ex-post budget balance,* is satisfied by a mechanism if for all histories and in all states, $\sum p_{i,t}\left(H_{t-1}, \widehat{\theta}_t\right) = 0$. We will also examine a weaker definition, similar to one that appears in Athey and Miller (A&M) [3]; their definition requires that in every period t, $E\left[\sum p_{i,t}\right] = 0$. While this is appropriate for environments where player types are drawn from independent and identical distributions every period, we will employ the following natural modification in our Markovian setting,

Definition 1 *A mechanism satisfies* ***within period ex-ante budget balance*** *if in every period t,*

$$E\left[\sum_{i=1}^{N} p_{i,t+1} \middle| H_t\right] = 0$$

That is, we assume the presence of a banker who is willing to provide unbounded insurance as long as the per-period balance of payments is zero in expectation. We can interpret this condition a participation constraint for the banker; on average, he need not infuse any wealth into the mechanism, but must absorb all imbalances. Indeed, mechanisms that satisfy this dynamic constraint do not require the planner to commit at time $t = 0$ to continue providing insurance to projects that might subsequently prove to be financially draining.

We now move to individual rationality. Suppose every agent i could guarantee herself a constant value of r_i at any time t by leaving the mechanism. This is unlike B&V [10] who define individual rationality in a framework where agents leaving the mechanism continue to consume the externalities that it generates. Our assumption is standard in dynamic contract theory. B&V's unconventional definition of individual rationality is influenced in part by their focus on finding conditions under which all efficient mechanisms are marginal contribution mechanisms. Which definition is more reasonable is context dependent. The definition adopted by B&V would be suitable for instance in settings where a government uses reports to determine the level of a non-excludable public good it should provide or auction off licences that might influence a rival's competitiveness in the market.

We formally define our notion of dynamic individual rationality below.

Definition 2 *A mechanism is* ***within-period ex-ante individually rational*** *if for all* t*, and every* i*,*

$$E\left[V_{i,t}\left(H_{i,t-1}\right)\right] \geq r_i.$$

Under the ex-ante constraint, players may only leave the mechanism in period t before their private signals are realized. Drawing on a poker analogy, players must ante up (or commit to participating) before they can be dealt a hand.

Finally, we introduce our notion of social rationality.

Definition 3 *The first-best policy is* ***socially rational*** *if for all* t*,*

$$W_t^*(H_{t-1}, \theta_t) \geq \sum_{i=1}^{N} r_i.$$

That is, the expected social welfare of an efficient mechanism is bounded below by the value that can be generated by allowing each agent to pursue her outside option. As we shall see, social rationality will be essential for our mechanisms.

1.4 Markov Decision Processes

Let us continue with our mission to construct efficient mechanisms that satisfy budget balance and individual rationality in the senses defined above. To do so, we first limit our attention to a environments that can be classified as Markov Decision Processes (MDP). As in A&S, this involves imposing additional restrictions:

(i) The flow utility that each agent receives from the planner's action $g_i(a^t, \theta^t)$ depends only on contemporaneous variables and can be expressed as $g_i(a_t, \theta_t)$.

(ii) The law of motion follows a Markov Process:

$$\Pi\left(\theta_{t+1}|\theta^t, a^t\right) = \Pi\left(\theta_{t+1}|\theta_t, a_t\right)$$

Assumption (ii) does not impose any restrictions on the correlation between agent i's current private signal, $\theta_{i,t}$, and the current or future private signals of agent j. As we noted in the introduction, this is an important departure from the independence assumption made in previous work. Agents may well have access to information about the private types of the other strategic players that is not available to the planner.

Assumptions (i) and (ii) in a finite state space model ensure that for some $0 < \overline{\delta} < 1$, if $\delta \in [\overline{\delta}, 1]$, the social welfare maximizing policy is a Markov policy and can be expressed as $a_t^* = a^*(\theta_t)$. An optimum Markov policy is called a Blackwell policy; it induces a Markov chain on the states with transition probabilities given by $\Pi\left(\theta_{t+1}|\theta_t, a_t^*\right)$. In such environments, one can construct efficient mechanisms that, not surprisingly, involve transfer schemes which are also time-invariant. For more information on MDPs, we refer the reader to [40].

We will also assume for simplicity that off-equilibrium reports are not possible in this framework. There is no loss of generality in this assumption. One can always augment our mechanisms with a "shoot them all" feature that will deter off-equilibrium reports when the first-best policy does not induce full support over the space of messages. It is important to note that heavy monetary penalties associated with a "shoot them all" mechanism may not be implementable if agent's may leave after the announcement stage but before the payments are made. This is not an issue when the mechanism planner only requires ex-ante individual rationality as we do.

1.4.1 An Imperfect Monitoring Approach

In dynamic environments, conflicts of interest between the planner and an agent arise on two fronts. In addition to the misalignment in flow payoffs, the planner must also account for disagreement among agents regarding the future evolution of the system. Our first mechanism addresses this two-fold conflict using incentivizing techniques typically employed in environments with imperfect monitoring. When private signals are correlated serially in time and across players, future reports can serve as noisy signals of whether current reports were truthful. The planner can therefore induce truth-telling in agent i at time t by offering her a set of lotteries over future continuation value to choose from via her current report.

A little more notation is regrettably unavoidable. Suppose agents in $-i$ report truthfully at time t so that $\widehat{\theta}_{-i,t} = \theta_{-i,t}$. For each $\widehat{\theta}_{-i,t}$, and for every signal-announcement pair $\{\theta_{i,t} = x, \widehat{\theta}_{i,t} = y\}$, consider the $K = \prod_{i=1}^{N} k_i$- column vector of beliefs,

$$\Pi_{xy}\left(\widehat{\theta}_{-i,t}\right) = \Pi\left(\theta_{t+1}\middle|[\theta_{i,t} = x, \theta_{-i,t}], a^*\left([\widehat{\theta}_{i,t} = y, \widehat{\theta}_{-i,t}]\right)\right)$$

This vector represents agent i's beliefs at time t about the state at time $t+1$ when agents $-i$ report truthfully, agent i with private signal x reports y and the planner implements the first best policy. Stacking the subset of such vectors that share the same $\widehat{\theta}_{-i,t}$ and report $\widehat{\theta}_{i,t} = y$, we can construct $k_i \times K$ matrices of the form $\Pi\left(\widehat{\theta}_{-i,t}, a^*_y(\widehat{\theta}_{-i,t})\right) = [\Pi_{1y}, \Pi_{2y}, \ldots, \Pi_{ky}]'$ where $a^*_y(\widehat{\theta}_{-i,t}) = a^*\left([\widehat{\theta}_{i,t} = y, \widehat{\theta}_{-i,t}]\right)$. The Markov property of the state transitions implies that these matrices do not depend on t. Our full rank condition is stated below -

Definition 4 (**Identifiability**) *An MDP is identifiable if for each i, $\theta_{-i,t}$, and $a \in A$, the set of vectors $S(\theta_{-i,t}, a) = \{\Pi\left(\theta_{t+1}\left|[\theta_{i,t}, \theta_{-i,t}]\right., a\right), \forall \theta_{i,t}\}$ is linearly independent.*

We use the term "identifiability" since it is this property that allows the planner to statistically distinguish false reports. Under this condition, all matrices of the form $\Pi\left(\widehat{\theta}_{-i,t}, a_y^*(\widehat{\theta}_{-i,t})\right)$ are full row rank. The reader will recognize that identifiability is analogous to the individual full rank condition defined by Fudenberg, Levine and Maskin [28].

To jump start this mechanism at $t = 1$, we will require that prior to the first period all agents and the planner share common beliefs about the distribution over states in the first period, $\Pi(\theta_1)$, and that the planner can compute $W_1^*(H_0)$ correctly. We can now construct a set of lotteries for each agent i to choose from at time t. Lotteries are indexed by reports at time t; the payoffs of each lottery are contingent upon messages received at time $t+1$ and represent the planner's promised values to the *truthful* agent. For each $\widehat{\theta}_{-i,t}$, we denote the k_i possible lotteries for agent i by length K column vectors, $w_{i,t+1}\left([\widehat{\theta}_{i,t} = y, \widehat{\theta}_{-i,t}]\right)$, $y = 1, \ldots, k_i$.

To further simplify notation, we use the short hand $g_{i,xy}(\theta_{-i,t})$ to represent agent i's period-t direct flow utility when her private signal is x and she reports that it is y, i.e. $g_{i,xy}(\theta_{-i,t}) = g_i\left(a^*\left([\widehat{\theta}_{i,t} = y, \theta_{-i,t}]\right), [\theta_{i,t} = x, \theta_{-i,t}]\right)$. If every agent were to report truthfully at time t, the following relations would hold for each i and $\theta_{i,t} = x$,

$$w_{i,t}\left(\widehat{\theta}_{t-1}\right)[x, \theta_{-i,t}] = g_{i,xx}\left(\theta_{-i,t}\right) + p_{i,t}\left(\widehat{\theta}_{t-1}, [x, \theta_{-i,t}]\right) + \delta V_{i,t+1}(H_{i,t}) \qquad (1.1)$$

where $w_{i,t}\left(\widehat{\theta}_{t-1}\right)[x, \theta_{-i,t}]$ represents the element of $w_{i,t}\left(\widehat{\theta}_{t-1}\right)$ corresponding to a period-t announcement of $[x, \theta_{-i,t}]$. These are typically referred to as the *promise-keeping* constraints of the planner.

We solve for the K vector, $w_{i,t+1}\left(y, \widehat{\theta}_{-i,t}\right)$ using the following system of k_i equations and K variables -

$$\begin{aligned}
\Pi'_{yy}\left(\widehat{\theta}_{-i,t}\right) w_{i,t+1}\left(y,\widehat{\theta}_{-i,t}\right) &= \alpha_{i,t+1}\left(y,\widehat{\theta}_{-i,t}\right) \\
\delta\Pi'_{xy}\left(\widehat{\theta}_{-i,t}\right) w_{i,t+1}\left(y,\widehat{\theta}_{-i,t}\right) &= w_{i,t}\left(\widehat{\theta}_{t-1}\right)[x,\theta_{-i,t}] - g_{i,xy}\left(\theta_{-i,t}\right)\ldots \\
&\quad \ldots - p_{i,t}\left(\widehat{\theta}_{t-1},[y,\theta_{-i,t}]\right) - \epsilon
\end{aligned} \tag{1.2}$$

where $\epsilon > 0$ and $\alpha_{i,t+1}\left(y,\widehat{\theta}_{-i,t}\right)$ defined explicitly below represents the expected continuation value $V_{i,t+1}(H_{i,t})$ for agent i when $\theta_t = \left[y,\widehat{\theta}_{-i,t}\right]$. Notice that when identifiability holds, $\Pi_{yy}\left(\widehat{\theta}_{-i,t}\right)$ and $\Pi_{xy}\left(\widehat{\theta}_{-i,t}\right)$ are linearly independent and an agent can expect a continuation value equal to $\alpha_{i,t+1}\left(y,\widehat{\theta}_{-i,t}\right)$ only when she is reporting truthfully. For all i, $\alpha_{i,t+1}\left(\widehat{\theta}_t\right) \geq r_i$ and satisfies

$$\sum_{i=1}^{N} \alpha_{i,t+1}\left(\widehat{\theta}_t\right) = W^*_{t+1}(H_{t-1},\widehat{\theta}_t) \tag{1.3}$$

That is, the planner's promises sum up, in expectation, to the value generated by the mechanism. It is this feature that will ensure within period ex-ante budget balance is satisfied. One possible choice is to set

$$\alpha_{i,t+1}\left(\widehat{\theta}_t\right) = r_i + \gamma_i\left(W^*_{t+1}(H_t) - \sum_{j=1}^{N} r_j\right) \tag{1.4}$$

where $0 < \gamma_i < 1$ and $\sum_{i=1}^{N}\gamma_i = 1$ and we define $\gamma_m = \min\{\gamma_i,\ \forall i\}$. Indeed, any sharing rule, including those that vary proportions stochastically in time, or allocate less than r_i to agent i, can be sustained under this mechanism. As pointed out earlier, this is in contrast to the marginal contribution mechanisms constructed by B&V for environments where private types evolve independently. We exclude sharing rules for which $\sum_{i=1}^{N}\gamma_i < 1$ only because these will necessarily violate budget balance.

We can show that with the lotteries specified above it will be ex-post incentive compatible for agent i to report the truth. Promises made in period $t-1$ are kept in period t by adjusting the transfer $p_{i,t}\left(\widehat{\theta}_{t-1},[x,\theta_{-i,t}]\right)$ in (1.1) to make up for any remaining difference once the direct flow utility and expected continuation value have been accounted for. This brings us to the first main result.

Proposition 1 *If an MDP is identifiable, and the first best policy is socially rational, there exists an efficient mechanism that satisfies within-period ex-ante budget balance.*

Proof. First we note that the construction of the lotteries ensures that there are no profitable one-shot deviations. This is because for every private signal $\theta_{i,t}=x$, and false report $\widehat{\theta}_{i,t}=y$, agent i gets a utility of

$$g_{i,xy}\left(\theta_{-i,t}\right)+p_{i,t}\left(\widehat{\theta}_{t-1},[y,\theta_{-i,t}]\right)+\delta\Pi_{xy}^{T}\left(\widehat{\theta}_{-i,t}\right)w_{i,t+1}\left(y,\widehat{\theta}_{-i,t}\right)=w_{i,t}\left(\widehat{\theta}_{t-1}\right)[x,\theta_{-i,t}]-\epsilon$$
$$<w_{i,t}\left(\widehat{\theta}_{t-1}\right)[x,\theta_{-i,t}]=g_{i,xx}\left(\theta_{-i,t}\right)+p_{i,t}\left(\widehat{\theta}_{t-1},[x,\theta_{-i,t}]\right)+\delta V_{i,t+1}(H_{i,t})$$

The incentive compatibility holds in an ex-post sense within period t; the agent prefers to report the truth even after observing the reports of the remaining players. The Markov property ensures that the first best policy is only a function of contemporaneous reports. Furthermore, the planner does not rely on past reports to compute lotteries and is therefore immune to any attempts by an agent to distort his beliefs. Under these condtions, the standard argument applies - if there are no profitable one-shot deviations, then no profitable multiperiod deviations can exist.

Finally, we check whether within period ex-ante budget balance is satisfied. Applying the promise keeping constraint,

$$\begin{aligned}E\left[\sum_{i=1}^{N}p_{i,t+1}\left(\widehat{\theta}_{t},\widehat{\theta}_{t+1}\right)\middle|H_t\right] &= E\left[\sum_{i=1}^{N}w_{i,t+1}\left(\widehat{\theta}_{t}\right)[\widehat{\theta}_{t+1}]-g_{i,t+1}\left(a_{t+1}^{*},\widehat{\theta}_{t+1}\right)\right.\\ \left.\ldots-\delta V_{i,t+2}\left(\widehat{\theta}_{t+1}\right)\middle|H_t\right] &= \sum_{i=1}^{N}\alpha_{i,t+1}\left(\widehat{\theta}_{t}\right)-W_{t+1}^{*}(H_t)=0\end{aligned}$$

where the third equality follows from (1.3). This completes the proof. ■

A few remarks are in order. First, a planner can redistribute surplus in any manner he wishes. This is an important observation in light of existing literature. The work of Neeman demonstrates that full surplus extraction in a static setting with correlated types is generically impossible. The I.M. approach, however, does not fall prey to Neeman's critique and permits full surplus extraction in dynamic, correlated environments. It also provides an interesting counterpoint to previous work by B&V, whose mechanisms distribute surplus according to the marginal contributions of the participants.

For any agent i, the ex-ante continuation value is $\alpha_{i,t+1}\left(\widehat{\theta}_t\right)$, a quantity that may be chosen to be no less than her outside option as in (1.4). However, this will not in general be enough for the I.M. mechanism to satisfy dynamic individual rationality. In particular, the mechanism we have constructed fails to rule out the possibility of an agent lying in period t and exitting the mechanism before period $t+1$, thereby avoiding the consequences of her false report. The mechanism as defined, therefore, only works when participation is mandatory.

One simple technique to deter "lying and leaving" is to require each agent to make a deposit with the planner at time 0 which she forfeits if she chooses to exit prematurely. Since the gains from a one-shot lie are bounded, this deposit can be made large enough to deter such behavior. This approach is not typically seen in the literature and we will not employ it. One potential argument against it is that is requires the planner to keep a portion of the surplus on behalf of the agent's until the mechanism ends. Although, this does not have any direct impact on the utility of the unconstrained, risk neutral agents, one may still consider this undesirable.

The I.M. mechanism also does not adapt readily to finite horizon settings. Since incentives are provided via lotteries over future promised values, the mechanism is hard-pressed to induce truth-telling in the final period.

Under a more restrictive version of the identifiability condition, each of the limitations above can be overcome if agents make one final report and payment after they leave the mechanism. A "lie and leave" strategy can be deterred for instance, by requiring the exitting agent to make one final payment that is contingent on the next period's reports of the remaining agents. The exitting agent is also required to report her type in the period after she leaves; since this report does not affect her payoff outside the mechanism, we will assume that she reports truthfully. We denote agent i's exit payment, to be paid in period $t+1$ if she leaves after t, by $e_{i,t+1}(\widehat{\theta}_t)[\widehat{\theta}_{-i,t+1}]$. For each $[y, \widehat{\theta}_{-i,t}]$, we solve for the $K_{-i} = \prod_{j \neq i} k_j$ dimensional vector $e_{i,t+1}(y, \widehat{\theta}_{-i,t})$ using the k_i equations,

$$\begin{aligned} w_{i,t}\left(\widehat{\theta}_{t-1}\right)[x, \theta_{-i,t}] \quad \geq \quad & g_{i,xy}\left(\theta_{-i,t}\right) + p_{i,t}\left(\widehat{\theta}_{t-1}, [y, \widehat{\theta}_{-i,t}]\right) \dots \\ & \dots + \delta r_i + \delta \Pi'\left(\theta_{-i,t+1} | [x, \widehat{\theta}_{-i,t}], a^*(y, \widehat{\theta}_{-i,t})\right) e_{i,t+1}(y, \widehat{\theta}_{-i,t}), \ \forall x \end{aligned} \tag{1.5}$$

where it is necessary for $k_i \leq K_{-i}$. Typically, this condition will require $N > 2$, unless $k_i = k$, $\forall i$. For a solution to exist, we will require a full rank condition over distributions of the form $\Pi\left(\theta_{-i,t+1} | [x, \widehat{\theta}_{-i,t}], a_t^*(y, \widehat{\theta}_{-i,t})\right)$. Formally, this is given by,

Definition 5 (**FR**) *An MDP is full rank if for each i, $\theta_{-i,t}$,, and $a \in A$, the set of vectors $S(\theta_{-i,t}, a) = \{\Pi\left(\theta_{-i,t+1} \left| [\theta_{i,t}, \theta_{-i,t}]\right., a\right), \forall \theta_{i,t}\}$ is linearly independent.*

It is easy to see that an MDP which is full rank is also identifiable. We can now state our result formally.

Proposition 2 *If an MDP is full rank, and the first best policy is socially rational, there exists an efficient mechanism that satisfies within-period ex-ante budget balance and within-period ex-ante individual rationality.*

Proof. Since the FR condition implies identifiability, efficiency and ex-ante budget balance follow from an earlier proposition. An agent's continuation value after reporting truthfully is

$$\Pi'_{yy}\left(\widehat{\theta}_{-i,t}\right) w_{i,t+1}\left(y,\widehat{\theta}_{-i,t}\right) = \alpha_{i,t+1}\left(y,\widehat{\theta}_{-i,t}\right) \geq r_i.$$

Under full rank, the system of equations in (1.5) has a solution and the exit payments are such that a "lie and leave" strategy is not profitable. Since agent i's report in the period after she leaves is truthful, it can be used in the construction of exit payments for other agents. ■

Returning to the finite horizon problem, let T denote the final period and suppose that agents must make reports and payments in period $T+1$ as well. Suppose further that agent i's transfer in period $T+1$ depends on her period T report and the reports of $-i$ in period $T+1$. As above, let us assume that indifference is always resolved in favor of truthful revelation. Lotteries over future promises constructed in period T then take distinct values over K_{-i} states. In keeping with our notation, we denote these lotteries by K_{-i} vectors of the form $w_{i,T+1}(y,\widehat{\theta}_{-i,T})$; for each $\widehat{\theta}_T$,

$$\Pi'\left(\theta_{-i,T+1}|[y,\widehat{\theta}_{-i,T}],a^*(y,\widehat{\theta}_{-i,T})\right) w_{i,T+1}\left(y,\widehat{\theta}_{-i,T}\right) = 0$$

$$\delta\Pi'\left(\theta_{-i,T+1}|[x,\widehat{\theta}_{-i,T}],a^*(y,\widehat{\theta}_{-i,T})\right) w_{i,T+1}\left(y,\widehat{\theta}_{-i,T}\right) = w_{i,T}\left(\widehat{\theta}_{T-1}\right)[x,\theta_{-i,T}]\ldots$$
$$\ldots - g_{i,xy}\left(\theta_{-i,T}\right) - p_{i,T}\left(\widehat{\theta}_{T-1},[y,\theta_{-i,T}]\right) - \epsilon$$

Transfers in period T, $p_{i,T}\left(\widehat{\theta}_{T-1},[y,\theta_{-i,T}]\right)$, still satisfy (1.1), although continuation values are now zero in expectation. The lotteries satisfy ex-ante budget balance in period $T+1$ since the sum of promises equals zero in expectation.

The mechanism we have proposed relies critically on every player's type evolving stochastically for all time. If any one player's type enters an absorbing state after which it remains constant or evolves deterministically, the reports of that player can no longer serve to incentivize the rest. Furthermore, if the MDP enters an absorbing state, the planner need not elicit any information in the future as he can infer it himself. The identifiability condition rules out the possibility of agents' types entering such absorbing states.

We request the reader to consider for a moment an important class of Markov decision processes for which evolution is only influenced by the choice of allocation; i.e. $\Pi(\theta_{t+1}|\theta_t, a_t) = \Pi(\theta_{t+1}|a_t)$. For this class of processes, the identifiability condition is necessarily violated and the imperfect monitoring approach is ineffective. An example of this scenario is the dynamic optimal insurance model examined by Doepke and Townsend [19] where each period agents receive income shocks which are correlated with past effort choices, but are independent of past income.

Finally, we note that our first proposition above is proved under the assumption that state transitions follow a Markov process. Our mechanism can, in fact, be extended, with a suitably redefined full rank condition, to any finite memory process whose evolution is of the form

$$\Pi(\theta_{t+1}|\theta_t, \ldots, \theta_{t-h}, a_t, \ldots, a_{t-h})$$

The planner must, in this case, take into account all private histories of length-h when constructing lotteries. As a consequence, the system of equations (1.2) will expand to k_i^h equations and K variables. The mechanism cannot be adapted to memories larger than $h = \log(K/\max(k_i))$. Within-period budget balance however will typically not be achievable.

1.4.2 Ex-Post Budget Balance and Collusion Proofness

In environments where a stronger version of the identifiability condition (one that is distinct from FR) holds, the I.M. mechanism can be improved considerably. The mechanism constructed above leaves the planner with several unused degrees of freedom. These can be employed by the planner to implement within-period ex-post budget balance, a condition that eliminates the need for a banker who is willing to finance temporary imbalances.

Furthermore, if the number of private types is sufficiently "large", the planner can also design a mechanism that is collusion proof; a quality that is often important in such environments. Inefficiency can arise, for instance, in repeated auction settings where agents form cartels that can profitably distort outcomes.

As a first step towards these objectives, we define the following stronger notion of identifiability.

Definition 6 (**Strict Identifiability**) *An MDP satisfies strict identifiability if for every $a \in A$, the set of vectors $S(a) = \{\Pi(\theta_{t+1} | \theta_t, a), \forall \theta_t\}$ is linearly independent.*

Each distribution $\Pi(\theta_{t+1} | \theta_t, a)$ is a vector of K elements and there are K vectors in each set. This condition does not therefore impose any restrictions on N or k_i. Strict identifiability is unlike its weaker counterpart which only requires states that are single-player deviations apart to be statistically distinguishable under action a.

Ex-post budget balance can now be achieved, if in addition to (1.2), the vectors $\{w_{i,t+1}\left(\widehat{\theta}_t\right), i = 1, \ldots, N\}$ are also required to satisfy the following K equations,

$$\sum_{i=1}^{N} w_{i,t+1}\left(\widehat{\theta}_t\right)[\widehat{\theta}_{t+1}] = \sum_{i=1}^{N}\left(g_{i,t+1}\left(\widehat{\theta}_{t+1}, a^*_{t+1}\right)\right) + \delta W^*_{t+2}(\widehat{\theta}_{t+1}), \ \forall \widehat{\theta}_{t+1} \tag{1.6}$$

for each $\widehat{\theta}_t$. The combined system, (1.2) and (1.6), is composed of NK variables and $\sum_{i=1}^{N} k_i + K$ equations for each $\widehat{\theta}_t$. The equations in (1.6) are linearly independent of each other and if strict identifiablility is satisfied the system has a well-defined solution whenever $k_i \geq 2, \forall i$ and $N \geq 2$. We state our next result below.

Proposition 3 *If an MDP is strictly identifiable and the first best policy is socially rational, there exists an efficient mechanism that satisfies within-period ex-post budget balance.*

Proof. For every $\widehat{\theta}_t$ and $a^*\left(\widehat{\theta}_t\right)$, strict identifiability ensures that the following $\left(\sum_{i=1}^{N} k_i + K\right) \times NK$ matrix

$$\begin{bmatrix} \Pi\left(\widehat{\theta}_{-1,t}, a^*\left(\widehat{\theta}_t\right)\right) & 0 & \dots & 0 \\ 0 & \Pi\left(\widehat{\theta}_{-2,t}, a^*\left(\widehat{\theta}_t\right)\right) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \Pi\left(\widehat{\theta}_{-N,t}, a^*\left(\widehat{\theta}_t\right)\right) \\ I_K & I_K & I_K & I_K \end{bmatrix}$$

,where matrices of the form $\Pi\left(\widehat{\theta}_{-i,t}, a^*\left(\widehat{\theta}_t\right)\right)$ are $k_i \times K$ and I_K is the K-dimensional identity matrix, is full row rank. The combined system of (1.2) and (1.6) can then be solved for $\{w_{i,t+1}\left(\widehat{\theta}_t\right), i = 1, \dots, N\}$. Finally, ex-post budget balance follows because for every $\widehat{\theta}_{t+1}$,

$$\begin{aligned} \sum_{i=1}^{N} p_{i,t+1}\left(\widehat{\theta}_t, \widehat{\theta}_{t+1}\right) &= \sum_{i=1}^{N}\left(w_{i,t+1}\left(\widehat{\theta}_t\right)[\widehat{\theta}_{t+1}] - g_{i,t+1}\left(\widehat{\theta}_{t+1}, a^*_{t+1}\right) - \delta V_{i,t+2}(H_{t+1})\right) \\ &= \sum_{i=1}^{N} w_{i,t+1}\left(\widehat{\theta}_t\right)[\widehat{\theta}_{t+1}] - \sum_{i=1}^{N} g_{i,t+1}\left(\widehat{\theta}_{t+1}, a^*_{t+1}\right) - \delta W^*_{t+2}(\widehat{\theta}_{t+1}) = 0 \end{aligned}$$

where the first equality follows from (1.1) and the second from (1.6). This completes the proof. ■

Of course, this mechanism can also be extended to finite horizons and augmented to satisfy individual rationality when the FR condition is also met.

We now turn to collusion-proofness. As discussed earlier, this may be a very desirable property in a number of economic settings. A relatively recent literature on repeated auctions has examined this problem with a particular focus on understanding whether limiting the extent of communication between bidders or prohibiting side contracts can help thwart collusive strategies and allow the planner (or in this case, seller) to implement efficient outcomes. As we shall see, in the environment we consider, the I.M. mechanism renders all coalitions powerless without having to outlaw any behavior among bidders. Our notion of collusion proofness is defined formally below.

Definition 7 *A mechanism is* ***collusion-proof*** *if there is no subset J of agents such that (a) all agents in $-J$ follow truth-telling strategies and (b) agents in J can deviate jointly and secure a higher value for the coalition.*

We do not explicitly model the side contracts that sustain the coalition, but will presume that such contracts will share the spoils among the members of the coalition so that each of them weakly prefers to collude than tell the truth. Our definition of collusion-proofness is a natural adaptation of Aumann's *strong equilibrium* to environments with transferable utility. Aumann's refinement admits only those Nash equilibria for which there is no coalition that can deviate jointly to give every member a weakly preferred outcome.

The system in (1.2) specifies a mechanism which is immune to deviations by individual agents. Let us now consider a subset J composed of $m = |J|$ players where $1 < m \leq N$. Notice that we include collusion of the grand coalition consisting of all N agents. In order to simplify the exposition, we will limit our analysis to symmetric environments in which $k_i = k$ for all i. We focus as before on first eliminating profitable one-shot joint deviations. Since colluding players may leave the coalition after their private signals have been realized, it is natural to adopt ex-post incentive compatibility as our equilibrium concept.

In any state θ_t, the players in J can misreport in one of $(k^m - 1)$ ways; none of these is preferred to reporting the truth as long as

$$\begin{aligned}\sum_{j\in J} w_{j,t}\left(\theta_{t-1}\right)\left[\theta_t\right] \geq \sum_{j\in J}\left\{g_j\left(\left[\theta_t, a_t^*(\widehat{\theta}_{J,t}, \theta_{-J,t})\right]\right) + p_{j,t}(H_{t-1}, \widehat{\theta}_{J,t}, \theta_{-J,t})\right\}\ldots \\ \ldots + \delta\Pi^T_{\theta_{J,t}\widehat{\theta}_{J,t}}\left(\widehat{\theta}_{-J,t}\right)\sum_{j\in J} w_{j,t+1}\left(\widehat{\theta}_{J,t}, \theta_{-J,t}\right)\end{aligned} \tag{1.7}$$

for all $\widehat{\theta}_{J,t}$ and where $\Pi_{\theta_{J,t}\widehat{\theta}_{J,t}}\left(\widehat{\theta}_{-J,t}\right)$ is the natural generalization of $\Pi_{xy}\left(\widehat{\theta}_{-i,t}\right)$ to a set J. These ex-post incentive constraints are similar to the second equation in (1.2). When combined with (1.2) and (1.6), we have system of Nk^N unknowns and $Nk + k^N + \sum_{m=2}^{N} \binom{N}{m}(k^m - 1)$ equations for each θ_t. A necessary condition for this system to have a solution is then

$$(N-1)k^N \geq N + (k+1)^N - 2^N$$

For any $N > 2$, there is a $k \geq 2$, large enough such that solutions to the system can exist. We can now state our result.

Proposition 4 *For every $N > 2$, there is an $M(N)$ such that if $k > M(N)$, the MDP is strictly identifiable and the first best policy is socially rational, there exists an efficient, collusion-proof mechanism that satisfies ex-post budget balance..*

Proof. We begin by restricting our attention to mechanisms that satisfy (1.2) and (1.6). Collusion proofness introduces $\sum_{m=2}^{N} \binom{N}{m}(k^m - 1)$ additional constraints. We claim that under strong identifiability, the augmented system is linearly independent and therefore has a solution. To see this, fix any θ_t. If the augmented system is not linearly independent, there must exist a set of deviations (possibly joint) $\{\theta_{J_1,t}, \theta_{J_2,t}, \ldots, \theta_{J_m,t}\}$ from θ_t, such that for some non-zero constants $d_1, d_2, \ldots, d_m$,

$$\sum_i d_i \Pi\left(\theta_{t+1} \left| \theta_t, a_t^*(\theta_{J_i,t})\right.\right) = 0$$

This contradicts strict identifiability. With strict identifiability, any set of lotteries that satisfy (1.2) and (1.7) eliminate profitable one-shot deviations. No multi-period deviations can be profitable since the planner incentivizes every state in each period to report truthfully. It follows that a collusion proof mechanism satisfying ex-post budget balance exists. ■

The result essentially states that future reports of the truth-telling players can serve as disciplining devices that deter collusion as long as the number of private signals k increases with the number of players. In fact for any $N > 2$, if $k \geq N + 1$, collusion proofness is feasible.

We have used the flexibility of the I.M. approach to implement mechanisms that satisfy ex-post budget balance and collusion-proofness. Depending on the context, a planner may seek mechanisms that satisfy other criteria; private signals may be multi-dimensional or players may require per-period limited liability. Implementing these might would boil down to introducing additional linear constraints that could pin down the unused degrees of freedom.

1.4.3 A Preference Aligning Approach

Our exposition thus far has identified two limitations of the imperfect monitoring approach - its failure to eliminate profitable "lie and leave" strategies and its inability to incentivize agents in finite horizon environments. We could address these in environments where agents make one final report and payment after they leave the system. In practice, a planner may have difficulty enforcing such payments. Our solution to these problems also relied on a full rank condition being satisfied for agent types after they exitted. This may well be unreasonable in finite horizon settings where the system ceases to evolve after the final period. Finally, these constructions rely heavily on indifference being resolved in favor of truth-telling; an assumption which makes them vulnerable to the same criticism as Mezetti's work.

The Preference Aligning approach allows us to develop efficient, individually rational mechanisms which can be adapted to finite horizon settings without relying on the assumptions noted above. The P.A. mechanism does however, require cyclic monotonicity.

A generalization of the better known single-crossing condition, cyclic monotonicity allows the planner to construct transfers that induce full separation within a stage game; that is, such transfers alone would make a myopic agent report truthfully. One can then further augment the transfer to agent i with payments that only depend on the reports of the remaining agents $-i$ without perturbing the incentives for truth-telling in the stage game. This addition would allow the planner to exploit the correlation structure of private signals to align an agent's preferences over future evolution with his own. The two step process ensures that the conflict of interests arising from misalignment of flow payoffs and disagreement over future evolution are both neutralized. Interestingly, the FR condition is sufficient for this approach.

We will employ a technique first developed by Cremer and Mclean (C&M) [16] who exploit the correlation of private signals to construct an auction which allows a seller to extract the full surplus from every bidder. A critical feature of their auction is that the interim expected utility of every bidder can be set to zero. We exploit the same principle in our mechanism to first set each agent's expected payoff equal to zero in every period. Having done this, one can transform an agent's continuation value into a scaled version of the planner's welfare function.

If the excess surplus generated by the mechanism is large enough, no player finds it profitable to deviate once and surrender future value by exitting immediately thereafter. In contrast with the I.M. mechanism, the P.A. mechanism does not rely on "punishing" deviations, in a statistical sense, with lower promised values in the future. This is why the P.A. mechanism is able to sustain dynamic individual rationality.

Without further delay, let us construct the mechanism. Suppose agents $-i$ report the truth at every t, so $\widehat{\theta}_{-i,t} = \theta_{-i,t}$. For each of the k_i possible values of $\theta_{i,t}$, there is a corresponding efficient allocation $a^*(\theta_{i,t}, \theta_{-i,t})$. Suppose further that the planner can construct action-transfer pairs, $\{a^*(\widehat{\theta}_t), m_i(\widehat{\theta}_t)\}$ such that in the stage game, there is a fully-separating ex-post equilibrium in which each agent i truthfully reveals her type. That is, for every $\theta_{i,t}$ and $\theta'_{i,t}$,

$$g_i(a^*(\theta_{i,t},\widehat{\theta}_{-i,t}),[\theta_{i,t},\widehat{\theta}_{-i,t}])+m_i(\theta_{i,t},\widehat{\theta}_{-i,t})\geq g_i(a^*(\theta'_{i,t},\widehat{\theta}_{-i,t}),[\theta_{i,t},\widehat{\theta}_{-i,t}])+m_i(\theta'_{i,t},\widehat{\theta}_{-i,t}) \tag{1.8}$$

The action-transfer pairs allow the planner to implement $a^*_t(.)$ in the stage game. We address the question of when such $\{m_i(\widehat{\theta}_{i,t},\widehat{\theta}_{-i,t})\}$ exist later in the development. Given these action-transfer pairs, we propose a transfer rule for the dynamic mechanism of the form

$$p_{i,t}\left(\widehat{\theta}_t;H_{t-1}\right)=(1-\gamma_i)\left\{m_i\left(\widehat{\theta}_t\right)-R_i\left(\widehat{\theta}_{-i,t-1},a_{t-1}\right)[\widehat{\theta}_{-i,t}]\right\}+s_i(\theta_t,\widehat{\theta}_t) \tag{1.9}$$

where $0<\gamma_i<1$, $\sum_{i=1}^N\gamma_i=1$, and the second term, which we will construct explicitly below, does not depend on agent i's period-t report. We will henceforth refer to the second term as the future-neutralizing transfer. The third term is chosen to satisfy

$$s_i\left(\theta_t,\widehat{\theta}_t\right)=(1-\delta)r_i+\gamma_i\left(\sum_{j\neq i}g_j\left(a^*(\widehat{\theta}_t),\theta_t\right)-(1-\delta)\sum_{j=1}^N r_j\right) \tag{1.10}$$

The reader will no doubt recognize that $s_i\left(\theta_t,\widehat{\theta}_t\right)$ is a sharing rule that will help redistribute excess surplus among the agents. More importantly, $s_i\left(.,.\right)$ is an affine transformation of the total surplus generated in period t. Note that the first and third terms of the transfer in (1.9) are independent of past history.

The future neutralizing transfer for agent i at time t is constructed as follows. Suppose the reports of agents $-i$ at time $t-1$ are given by $\widehat{\theta}_{-i,t-1}(=\theta_{-i,t-1}$, under truth-telling) and a_{t-1} was the allocation chosen in that period. Conditional upon her private signal at time $t-1$, agent i's beliefs about the period-t signals of agents $-i$

is given by a $K_{-i} = \prod_{j \neq i} k_j$ dimensional vector $\Pi\left(\theta_{-i,t} | \left[\theta_{i,t-1}, \widehat{\theta}_{-i,t-1}\right], a_{t-1}\right)$.These beliefs can be inferred from the Markov decision process. There are k_i such distributions, one for each private signal $\{\theta_{i,t-1} = x\}$. Stacking these vectors together gives a $k_i \times K_{-i}$ matrix

$$\Pi\left(\widehat{\theta}_{-i,t-1}, a_{t-1}\right) = [\ \Pi_1\left(\widehat{\theta}_{-i,t-1}, a_{t-1}\right)\ \ \Pi_2\left(\widehat{\theta}_{-i,t-1}, a_{t-1}\right)\ \ \ldots\ \ \Pi_{k_i}\left(\widehat{\theta}_{-i,t-1}, a_{t-1}\right)\ ]'$$

where $\Pi_x\left(\widehat{\theta}_{-i,t-1}, a_{t-1}\right) \equiv \Pi\left(\theta_{-i,t} | [\theta_{i,t-1} = x, \widehat{\theta}_{-i,t-1}], a_{t-1}\right)$, $x = 1, \ldots, k_i$. We note that agent i's report at time $t-1$ affects the matrix $\Pi\left(\widehat{\theta}_{-i,t-1}, a_{t-1}\right)$ only through its influence on the action a_{t-1}. The construction of the future neutralizing transfer will not rely on this report being truthful. We note that under the FR condition developed earlier, matrices of the form $\Pi\left(\widehat{\theta}_{-i,t-1}, a_{t-1}\right)$ are full row rank. We recall that this condition is satisfied only if $k_i \leq K_{-i}$ for all i. Typically, since $k_i \neq k_j$, it will require $N > 2$.

We now define the future-neutralizing transfer in terms of the following system of k_i equations and K_{-i} unknowns-

$$\Pi\left(\theta_{-i,t-1}, a_{t-1}\right)' R_i\left(\widehat{\theta}_{-i,t-1}, a_{t-1}\right) = \begin{bmatrix} E\left[g_i(a^*(\widehat{\theta}_t), \theta_t) + m_i(\widehat{\theta}_t) \middle| [1, \widehat{\theta}_{-i,t-1}], a_{t-1}\right] \\ \vdots \\ E\left[g_i(a^*(\widehat{\theta}_t), \theta_t) + m_i(\widehat{\theta}_t) \middle| [k_i, \widehat{\theta}_{-i,t-1}], a_{t-1}\right] \end{bmatrix} \tag{1.11}$$

where $R_i\left(\widehat{\theta}_{-i,t-1}, a_{t-1}\right)$ is a vector whose K_{-i} elements are denoted by

$$\{R_i\left(\widehat{\theta}_{-i,t-1}, a^*_{t-1}\right)[\widehat{\theta}_{-i,t}]\}$$

When FR holds, the system in (1.11) has a solution. By construction, the second term in the period-t transfer does not depend on $\widehat{\theta}_{i,t}$ and is influenced by $\widehat{\theta}_{i,t-1}$ only through a_{t-1}; the fact that an agent's transfers can be computed without relying on her past reports will play a critical role in eliminating profitable double deviations.

Let us now return to the question of whether action-transfer pairs that implement truthful revelation in the stage game always exist. A necessary and sufficient condition for implementability with single-dimensional types is derived by Rochet [41]. This condition, which has since been termed "cyclic monotonicity", is a generalization of the more widely used Spence-Mirrlees (or single-crossing) condition that is necessary and sufficient for the implementability of monotonic functions in quasi-linear settings. For multi-dimensional types, we refer the reader to the work of Bikhchandani *et. al* [?]. We state the cyclic monotonicity condition formally below.

Definition 8 ***(Cyclic Monotonicity)*** *A flow utility function* $g_i(a, [\theta_i, \theta_{-i}])$ *and first best policy* $a^*(\cdot)$ *satisfy cyclic monotonicity if for each* θ_{-i} *and for all finite cycles* $\theta_i(0), \theta_i(1), \ldots, \theta_i(m) = \theta_i(0)$*, in* Θ_i

$$\sum_{k=0}^{m} [g_i(a^*(\theta_i(k), \theta_{-i}), [\theta_i(k+1), \theta_{-i}]) - g_i(a^*(\theta_i(k), \theta_{-i}), [\theta_i(k), \theta_{-i}])] \leq 0$$

If this condition is satisfied, we can construct action-transfer pairs $\{a^*(\widehat{\theta}_t), m_i(\widehat{\theta}_t)\}$ that satisfy (1.8). We should point out that, in finite horizon environments, the first best policy will typically not be a stationary one, and cyclic monotonicity will be required in every period. We are now ready to state the final result of this paper.

Proposition 5 *If an MDP is full rank and satisfies cyclic monotonicity, and the first-best policy is socially rational, there exists an efficient mechanism that satisfies within-period ex-ante budget balance. If* $W^*_{t+1}(H_t) \geq \sum_{j=1}^{N} r_j + \frac{(N-1)\Delta}{\delta}$*, the efficient mechanism also satisfies within-period ex-ante individual rationality.*

Proof. To begin, we show that there are no profitable one-shot deviations. Consider an agent i with private signal $\theta_{i,t} = x$ and suppose that the remaining agents report their signals truthfully, $\widehat{\theta}_{-i,t} = \theta_{-i,t}$. The expected flow utility that agent i receives at time $t+1$ conditional on a report (potentially false) at time t, $\widehat{\theta}_{i,t} = x'$ is given by

$$E\left[g_i(a^*(\theta_{t+1}), \theta_{t+1}) + p_{i,t+1}(\theta_{t+1}) \middle| \theta_t, a^*([x', \theta_{-i,t}])\right]$$

$$= \Pi_x\left(\theta_{-i,t}, a^*([x', \theta_{-i,t}])\right)' \left(g_i(a^*(\theta_{t+1}), \theta_{t+1}) + p_{i,t+1}(\theta_{t+1})\right)$$

$$\begin{aligned} = \quad & E\left[g_i(a^*(\widehat{\theta}_{t+1}), \theta_{t+1}) + (1-\gamma_i) m_i(\widehat{\theta}_{t+1}) \middle| \theta_t, a^*([x', \theta_{-i,t}]\right] ... \\ & ... - (1-\gamma_i)\Pi_x\left(\theta_{-i,t}, a^*\right)' R_i\left(\widehat{\theta}_{-i,t}, a^*([x', \theta_{-i,t}])\right) + E\left[s_i\left(\theta_{t+1}, \widehat{\theta}_{t+1}\right) \middle| \theta_t, a^*([x', \theta_{-i,t}]\right] \end{aligned}$$

$$= E\left[s_i\left(\theta_{t+1}, \widehat{\theta}_{t+1}\right) + \gamma_i g_i(a^*(\widehat{\theta}_{t+1}), \theta_{t+1}) \middle| \theta_t, a^*([x', \theta_{-i,t}]\right]$$

$$= \gamma_i E\left[\sum_{j=1}^{N} g_j\left(a^*(\widehat{\theta}_{t+1}), \theta_{t+1}\right) \middle| \theta_t, a^*([x', \theta_{-i,t}]\right] + (1-\delta)\left(r_i - \gamma_i \sum_{j=1}^{N} r_j\right)$$

Thanks to the future-neutralizing transfer, the expected flow utility in period $t+1$, from agent i's period t perspective, is linear in the total flow utility generated in that period. The expected continuation value of agent i if he were to revert to telling the truth after period t, is then given by

$$\begin{aligned} V_{i,t+1}([x, \theta_{-i,t}], a^*([x', \theta_{-i,t}])) \quad = \quad & \gamma_i E\left[\sum_{s=t+1}^{\infty} \delta^{s-(t+1)} \sum_{j=1}^{N} g_j\left(a^*(\widehat{\theta}_s), \theta_s\right) \middle| \theta_t, a^*([x', \theta_{-i,t}]\right] \\ & ... + \left(r_i - \gamma_i \sum_{j=1}^{N} r_j\right) \\ = \quad & r_i + \gamma_i\left(W_{t+1}(\theta_t, a^*([x', \theta_{-i,t}]) - \sum_{j=1}^{N} r_j\right) \end{aligned}$$

The agent's total utility at time t from reporting x' is given by

$$\begin{aligned} & (1-\gamma_i)\left\{g_i(a^*(x', \widehat{\theta}_{-it}), \left[x, \widehat{\theta}_{-it}\right]) + m_i\left(x', \widehat{\theta}_{-i,t}\right) - R_i\left(\widehat{\theta}_{-i,t-1}, a_{t-1}\right)[\widehat{\theta}_{-i,t}]\right\} + \ldots \\ & \ldots + \gamma_i g_i(a^*(x', \widehat{\theta}_{-it}), \left[x, \widehat{\theta}_{-it}\right] + \; s_i\left(\theta_t, \left[x', \widehat{\theta}_{-it}\right]\right) + \delta V_{i,t+1}(x, \theta_{-i,t}, a^*([x', \theta_{-i,t}]) \end{aligned}$$

If cyclic monotonicity is satisfied, truth-telling maximizes the first two terms of this expression. The third term cannot be influenced by the agent's report and therefore does not disturb incentives. When combined, the fourth, fifth and sixth terms constitute an affine transformation of the planner's objective function which is maximized when the first-best allocation is chosen. Since truth-telling induces the first-best allocation, it follows that no such one-shot deviation is profitable.

Next we rule out the possibility of an agent employing a profitable "lie and then leave" strategy. It suffices if for any $\theta_{i,t} = x$,

$$s_i\left(\theta_t, \left[x', \widehat{\theta}_{-it}\right]\right) + \delta r_i \leq s_i\left(\theta_t, \left[x, \widehat{\theta}_{-it}\right]\right) + \delta V_{i,t+1}(\theta_t, a_t^*(\theta_t))$$

$$\Rightarrow \sum_{j\neq i}\left(g_j\left(a^*(\theta_t'), \theta_t\right) - g_j\left(a^*(\theta_t), \theta_t\right)\right) \leq \delta\left(W_{t+1}^*(\theta_t, a^*(\theta_t)) - \sum_{j=1}^{N} r_j\right)$$

This condition is satisfied as long as $W_{t+1}^*(\theta_t, a^*(\theta_t)) \geq \sum_{j=1}^{N} r_j + \frac{(N-1)\Delta}{\delta}$.

Multi-period deviations cannot be profitable in this mechanism since the planner incentivizes every type of each agent to report truthfully. The first-best policy choice is function of contemporaneous reports alone and the computation of transfers for agent i does not rely on her past reports. The planner is therefore immune to any attempts by the agent to distort his beliefs.

Under truthful revelation, agent i's expected continuation value at time t is given by

$$V_{i,t+1}(\theta_t, a_t^*(\theta_t)) = r_i + \gamma_i\left(W_{t+1}^*(\theta_t, a^*(\theta_t)) - \sum_{j=1}^{N} r_j\right) \geq r_i$$

where the inequality follows from social rationality. This proves that ex-ante individual rationality constraints are met.

We can now turn to budget balance. The present value of imbalances is given by

$$
\begin{aligned}
E\left[\sum_{i=1}^{N} p_{i,t}\middle| H_{t-1}\right] &= E\left[\sum_{i=1}^{N}(1-\gamma_i)\left\{m_i\left(\widehat{\theta}_t\right) - R_i\left(\widehat{\theta}_{-i,t-1}, a^*_{s-1}\right)[\widehat{\theta}_{-i,s}]\right\}\dots\right. \\
&\quad \left.\dots + s_i(\theta_t, \widehat{\theta}_t)\middle| H_{t-1}\right] \\
&= E\left[\sum_{i=1}^{N} -(1-\gamma_i) g_i(a^*(\widehat{\theta}_t), \theta_t)\dots\right. \\
&\quad \left.\dots + \sum_{i=1}^{N}\gamma_i \sum_{j\neq i} g_{j,t}\left(a^*(\widehat{\theta}_t), \theta_t\right)\middle| H_{t-1}\right] = 0
\end{aligned}
$$

where the first equality follows from (1.9) and the second equality from (1.10). This completes the proof. ■

A number of observations are worth highlighting. First, the P.A. approach allows the planner to implement almost any sharing rule to redistribute excess surplus. The only requirement however, is that the proportion of excess surplus allocated to an agent remain constant for all periods.

Second, if $W^*_{t+1}(H_t) \geq \sum_{j=1}^{N} r_j + \frac{(N-1)\Delta}{\delta}$, the planner can also implement within-period ex-ante individual rationality. Our proof shows that under social rationality, an agent's continuation value can always be chosen to be greater than her outside option and that the additional constraint on $W^*_{t+1}(H_t)$ rules out the possibility of profitably exitting the mechanism after submitting a false report. The additional condition essentially states that the expected surplus a mechanism can produce starting in any state is so large that it is not worth surrendering future participation for the bounded gain from a one time false report. Finally, we note that this mechanism does not require the planner to extract a final payment from an exitting agent.

The fact that this approach lends itself to finite horizon settings is also easy to see. In the final period, denoted $t = T$, agent i's flow utility will simply be given by

$$(1-\gamma_i)\left\{g_{i,T}(a^*(x', \widehat{\theta}_{-iT}), \left[x, \widehat{\theta}_{-iT}\right]) + m_i\left(x', \widehat{\theta}_{-i,T}\right) - R_i\left(\widehat{\theta}_{-i,T-1}, a_{T-1}\right)[\widehat{\theta}_{-i,T}]\right\}\ldots$$
$$\ldots + \gamma_i g_{i,T}(a^*(x', \widehat{\theta}_{-iT}), \left[x, \widehat{\theta}_{-iT}\right]) + s_i\left(\theta_T, \widehat{\theta}_T\right)$$

Since truth-telling maximizes the first two terms, the third term cannot be influenced by i and the final two terms constitute an affine transformation of the total flow utility in period $t = T$, truth-telling is a best response for agent i. Unlike I.M., the mechanism does not rely on reports from the players made in period $T + 1$.

As compared to the I.M. approach, the preference aligning method gives the planner fewer degrees of freedom. As a consequence, we see that this mechanism typically requires $N \geq 3$ and only satisfies ex-ante budget balance. The P.A. approach requires that the future-neutralizing transfer for agent i be independent of her report. The lottery $R_i\left(\widehat{\theta}_{-i,t}, a_t\right)$ is therefore over the K_{-i} possible reports of agents $-i$ at time $t+1$. In contrast, future promised values $w_{i,t+1}\left(\widehat{\theta}_t\right)[\widehat{\theta}_{t+1}]$ constructed under the I.M. approach are fundamentally influenced by agent i's report at time $t+1$ and therefore lotteries may have distinct payoffs for every one of the K possible reports that can be received in that period. It is this underlying distinction that makes the I.M. approach more flexible.

1.5 Future Work

A number of questions remain unanswered. What does weakening the equilibrium to a Bayesian one have to offer ? Are environments where types are neither independent nor identifiable doomed ? What additional insights do our mechanisms offer when an agent exitting the system continues to consume its externalities ? How can first-best allocations be implemented in non-Markovian environments? Can we subscribe to the Wilson doctrine and implement efficient allocations using detail-free mechanisms? To what extent is it possible to move to decentralized, self-enforcing mechanisms that do not rely on the presence of a banker or planner ?

One may also investigate optimal mechanism design when players are risk averse. In such settings, giving incentives for truth-telling will necessarily mean that first-best implementation is no longer possible. Related work on models of dynamic optimal taxation and/or insurance which analyzes economies where players try to share risk and exert efficient levels of effort have recently generated a great deal of interest. Finding common ground with that work may prove fruitful. These are just a few of the many questions that can be pursued in the area of multi-player dynamic contracting.

Chapter 2

Ideologues beat Idealists: A Paradox of Strategic Voting [1]

We'd all like to vote for the best man but he's never a candidate.

— Frank Hubbard

It is often the case that voters expect to have more information after the election and before a policy choice is made, as the following hypothetical but plausible scenario shows. A presidential election is in the offing, and the result hinges on one central issue — How best to respond to a country that could pose a threat. Voters could be divided on this issue — some support a direct confrontation (policy-0), while the less hawkish prefer a diplomatic response(policy 1). Voters are aware that their ranking of policies could change after elections are held, but before a policy choice is made; these we call "shocks". These shocks could be weak or inconclusive; voters then react to this "idiosyncratic" shock depending on their personalities. However there is a chance that some very conclusive evidence will come to the fore, either for or against the said country being a threat, and cause all voters to agree on what the right policy is ; this we call a "common" shock. If it is proved that the enemy was close

[1] Co-authored with Sambuddha Ghosh.

to developing nuclear weapons, then even the pacifist prefer confrontation; similarly, everybody prefers a diplomatic response when there is no doubt that the adversary is technologically incapable of mounting a serious military threat. Two candidates contest on the following platforms. The first, B, is known to prefer policy 0 come what may; the other, K, is an unbiased candidate who credibly promises to wait until final rankings are formed after the shocks are received, and to then take the action that the majority prefer. Who stands a better chance of being voted to power in a (simple) majority election? One might expect the unbiased candidate to be the natural choice of the population, especially when there is a significant probability of a common shock that makes the committed policy bad for all voters: In the case of a common shock, the unbiased candidate always implements what all voters prefer. We find, perhaps counter-intuitively, that voters may prefer the first candidate, thereby committing to a policy rather than waiting to learn their true rankings. This problem is greatly aggravated when voters are rational and their strategies constitute a Nash equilibrium or, equivalently, voters base their decision on the scenario where they are pivotal. This voting equilibrium survives a high chance that new evidence exposes the committed alternative as being undesirable for *all* voters, i.e. a high probability of a common shock away from the policy committed to. Political satirist Frank Hubbard's quip, quoted at the start of the paper, could be turned on its head — (Often) we wouldn't like to vote for the best man even if he were a candidate!

We emphasize that the electorate is faced with a choice between two *candidates*, one of whom always offers a fixed policy from which he derives private benefits, while the other is one who offers a state-contingent plan. Had the electorate's choice had been between two *policies* in the above environment with shocks to rankings, then there would have been a positive probability of each policy being the socially optimal one. What makes our inefficiency quite stark is that the unbiased candidate must, in *any state of the world*, implement the policy that maximizes social welfare, but he still loses. When voters choose between two policies each voter chooses the one that he expects to prefer at the next date. As we shall see, this could give rise to inefficient choices as documented by earlier work discussed later in greater depth. A somewhat finer reasoning enters into the decision of the voter when he is strategic and must choose between the two candidates B and K. To evaluate K's plan the voter must now take into account not merely his own ranking of policies but others' rankings and their equilibrium voting strategies; this is because he needs to keep track of where the ex-post social optimum will be. This added element leads the pivotal voter to a new form of inefficiency when he conditions his decision on being pivotal, which fuels his fear of being in a precarious majority even if that is unconditionally very unlikely. This point is elaborated in our examples and proofs.

What are the implications of the common shock for our result ? A common shock captures a situation in which a large number of voters change their rankings in a correlated fashion, as opposed to independently. [2] This feature enriches the insights to be gained from our model: Presence of common uncertainty both adds realism and serves as a test of the robustness of our results. In the event of an common shock, the unbiased candidate provides perfect insurance to all voters in the electorate and therefore a high chance of such an event pushes voters away from the suboptimal or committed candidate. We show that the inefficiency is robust to a high probability of a common shock. When the committed candidate comes to power, the common shock exacerbates the degree of social inefficiency as it creates the possibility of the entire electorate suffering a poor policy choice. On a more reassuring note, in the absence of common shocks, the probability that an elected committed candidate implements the socially optimum policy approaches unity for large populations.

The choice between the committed and the unbiased candidate can also be thought

[2] This can be rationalized as follows - each individual's utility function is comprised of a private value component and a common value component. The common shock might significantly alter the common values component of the utility thereby precipitating a correlated shift in rankings.

of as the choice between picking a policy immediately and waiting until further information makes a more informed choice possible. We show that even with a substantial chance of a common shock, the electorate might choose to act in haste and do the former. Examples of public referenda fit this formulation more naturally. In October 1992, the Swedish Nuclear Fuel and Waste Management Company (SKB), an organization charged with the responsibility of safely disposing nuclear waste, proposed to conduct a study to determine the feasibility of locating a repository. One of the towns that seemed worthy of further investigation was Storuman, in northern Sweden. The proposal polarized the community into those who opposed bringing nuclear waste to Storuman and those who believed that likely economic benefits made the investigation worthwhile. The findings of the SKB would not be binding on the city and if Storuman were deemed feasible it would still be up to the city council to decide, presumably in keeping with public opinion and the interest of the city, whether to allow SKB to actually go ahead with building a nuclear waste dump. A 1995 referendum asked "whether SKB should be allowed to continue the search for a final repository location in Storuman". The outcome was an overwhelming 'no' (70.5%): the public opted to reject it outright rather than allow more information to be disclosed by a scientific study.

In an article published in the Op-Ed section of the L.A. Times, Bruce Schulman argues that changing sides has been costly in American politics of late. Candidates spend resources trying to explain away changes in their stand on key issues, from affirmative action to foreign policy. Even fairly incontrovertible evidence of having changed does not dissuade them for arguing otherwise. Schulman suggests that political candidates do not wish to come across as opportunists who pander to the electorate for political gain. This line of reasoning has also been explored by Callander [12] and later by Kartik and McAfee [32], who examine the nature of electoral competition when candidates communicate their types through their choice of platform;

they assume that the electorate has explicit preferences over unobservable traits such as candidate character. The results in our paper offer a different explanation for why political candidates might prefer to commit to an ideology rather than update their stands as new information becomes available. We argue that in an environment with changing preferences, the best conceivable flip-flopper, one who adjusts his position to what is best for society at large, cannot expect to win against an ideologue. Office seeking candidates might therefore prefer to be perceived as having ideological biases even when the electorate does not intrinsically value this quality.

The results in this paper also offer insights into the type of candidates who *enter* an election. If entering an election is costly and candidates have to choose a plan on which to run, we show that a unbiased candidate may choose not to enter the race. In this case, the electorate will not be presented with the option of voting for the best candidate, as in the tongue-in-cheek quote at the start of the paper.

Our work is closely related to three strands of literature— the first, on status-quo bias against reforms that are put to vote; the second, on pivotal voting; and the third, on spatial competition and the median-voter result. The status-quo bias in reform is well documented — changes that are known to benefit a majority ex-post are not passed ex-ante because some of the would-be winners under the reform vote against the reform; see for example Samuelson and Zeckhauser[42]. Fernandez and Rodrik[3] [26](FR) provide a rigorous explanation in the context of trade reforms, to our knowledge the first that does not appeal to risk-aversion. Their explanation is based on the identity of the winners within the majority group being unknown at the time of voting. All voters of a group are *ex-ante identical* and hence maximize the expected value of the *group*, behaving in effect like the representative voter of the group. This brings us to the second strand of related literature, which uses the concept of pivotal voting, first introduced in the '*Theory of Voting*' by Farquharson[20]. More recently

[3] Roland Benabou first drew our attention to this paper, and pointed out a very natural link between our work and theirs.

this difference has been exploited by Austen-Smith and Banks [5] and Feddersen and Pesendorfer [24] to analyze information aggregation in elections. The following comment is intended to avoid legitimate confusion about the role of pivotal voting in our model. Usually the difference between sincere and pivotal voting arises from the fact that the private signal of each voter i affects the valuation of other voters $j \neq i$: When j is pivotal, he can infer the distribution of signals of other voters $i(\neq j)$ and thus updates his rankings conditional on being pivotal. We consider a model with *private values*, i.e. any voter's ranking of alternatives does not depend on others' signals or rankings. But, as we just argued, there is no additional information about voter j's ranking of policies contained in his being pivotal; so why should his vote depend on it at all? The paradox is resolved as follows: When voter j is pivotal, he can infer the probabilities with which the alternative policies will be implemented at the next date by the unbiased candidate K ; this determines voter j 's preferences over the *candidates* and thereby influences the outcome of the voting. Lastly, our work may also be linked to models of spatial voting, notably the pioneering work of Downs and Hotelling. In some senses we provide a framework that argues why the reverse, in a very loose sense to be described later, of the median voter result might hold; in our model a candidate with an extreme position beats an unbiased candidate.

The remainder of our paper is organized as follows: Section 1 puts our contribution in perspective with illustrative examples; Section 2 sets up the model, presents the decision problems of the sincere and pivotal voters, and characterizes the resulting equilibria; Section 3 looks at situations where the voting rule at the initial date differs from that used at the subsequent date; Section 4 provides a discussion of the candidate entry problem; and Section 5 concludes.

2.1 Illustrative Examples

Example 1

We begin with an example that captures the logic of FR. To fix ideas, consider two sectors in an economy — X and Y, with 76 workers in sector in X, and 25 in sector Y. Workers of sectors X and Y are referred to as type-0 and type-1 voters respectively. A simple majority election allows voters to choose whether to stick to the status-quo policy (0) or to implement a reform policy (1). A voter who prefers the policy implemented gets a utility of 1, while others get 0. The shocks are as follows: With probability 0.2 there is a common shock that moves all voters to 1; with probability 0.8 there is an idiosyncratic shock that makes each worker of type-0 become a type-1 with probability p independently of the others. The original model of FR does not have a common shock, but we add it for realism and to facilitate comparison of their model with ours.

Case 1: p is not too small. If $p = .25$, the expected number of people who prefer policy-1 at the next date will be $0.2 * 101 + 0.8(25 + 0.25 * 76) \approx 56$ which is more than half; so the efficient outcome requires the reform to be passed. But it is not known who the beneficiaries of reform in sector X will be when there is an idiosyncratic shock because all individuals in X are ex-ante identical. Therefore the expected gain to a current type-0 from the reform is $0.2 * 1 + 0.8 * 0.25 * 1 = 0.4$, which is less than the utility from policy-0, $0.8 * 0.75 = 0.6$. The type-0 voter prefers policy 0, and the reform is defeated even though it is welfare improving.

Case 2: p very small, i.e. the idiosyncratic shock is small and does not change the balance of power. Say, $p = .1$. The expected number of people who prefer policy-1 at the next date is $0.2 * 101 + 0.8 * (25 + 0.1 * 76) \approx 46$. Then policy-0 maximizes social welfare and all type-0 voters support it because the gain from the reform is $0.2 * 1 + 0.8 * .1 * 1 \approx .28$; there is no inefficiency.

Essentially what each voter does is guard the interests of the group to which he is most likely to belong at the next date. In the above example each type-0 voter is more likely to remain a type-0. Starting with a case where type-0's are in a majority, this is inefficient if and only if each type-0 is more likely to stay than to switch *and* the type-1's are expected to be in a majority at the next date. Although the model of FR is not directly comparable to ours, their work essentially points out the above source of inefficiency.

Note that there is no way for one voter to be pivotal in the framework above. What if we were to introduce a means by which each voter could be pivotal in the example above? Suppose that instead of there being exactly 76 and 25 workers in the two sectors, we have a model in which a worker was randomly chosen by nature to be of type-0 with probability ≈ 0.75. If type-0's and type-1's vote for different candidates in equilibrium, this assigns a positive probability to each voter being pivotal. Would this change how a typical voter behaves? Note that the voter in FR knows the probability that he will prefer each alternative; he votes for that which generates the highest expected utility. This is a weakly dominant strategy, even when every voter assigns positive probability to being pivotal. Thus the logic outlined in the above example works even if each voter can be pivotal.

As we shall see, our framework allows us to distinguish between pivotal and sincere voting. We show that a combination of two features— voting over candidates instead of policies, and strategic voting — makes the inefficiency more pervasive, although either one alone would not do so. Thus inefficiency may result even when the probability of switching in response to an idiosyncratic shock p is very small, as the following example shows.

Example 2:

Let us now turn to our framework. There are 101 voters, and two alternative policies, 0 and 1. At date-0, the initial date, nature chooses each voter's type, which is the policy he ranks higher. Types are drawn independently — type 0 with probability .75, and type 1 with probability .25; each voter learns his own type. However there is uncertainty in rankings — voters understand that when it is time to implement a policy at the next date (date-1), their ranking of policies could be different. Each voter gets a utility of 1 if his preferred alternative is implemented, and 0 otherwise. There are two candidates — B, the first one, is known to have private benefits from implementing policy-0; the second, K, is the unbiased candidate who behaves like the social-planner and implements the policy that the majority prefers at date-1.

Next period one of two shocks is possible: With probability 0.2, a common shock causes all voters will switch together towards 1 and with probability 0.8 there is an idiosyncratic shock and each voter switches independently towards 1 with probability $p = .1$ as in case 2 above. We should emphasize here that candidate K is the efficient choice regardless of the shock.

Case a (Sincere Voting): The utility of a type-0 voter from voting B is $1 - 0.2 - 0.8 * 0.1 = 0.72$, since he gets a utility of 1 as long as he stays a type-0. Unconditionally, type-0's are expected to be in the majority at both dates with probabilities close to 1[4]. So if a type-0 voter is sincere and does not condition on being pivotal, his probability of agreeing with K is very close to 1, and he votes for him. There is no inefficiency as in case 2 above.

Case b (Strategic/Pivotal Voting): But the rational voter recognizes that his vote matters only when he is pivotal. The utility from voting B is still 0.72, because B chooses a fixed policy. The utility from voting K is now $1 - 0.8 * 0.9 * (1 - (0.9)^{50}) \approx$

[4]It can be checked that if the type-0's are in a majority at date-0, they can become a minority with probability of the order of 10^{-23}.

0.28. Conditional on the pivotal scenario, the majority will become the minority with probability close to 1 if there is an idiosyncratic shock. The pivotal type-0 voter then ends up in the minority if he is not among those who switch. Conditioning on the knife-edge majority therefore leads him to vote against the reform.

Nature's draw is almost certain to result in type-0's being in a majority at date 0; consequently the inefficient choice B is almost certain. A candidate committed to a fixed policy beats the unbiased candidate, even though the probability of idiosyncratic change is very low! If candidate B is elected, the probability that the policy chosen is bad for the everyone is 0.2; this gives a sizeable lower bound on the degree of inefficiency[5]. It might appear at first glance that what the rational voters should try to guard against is the common shock towards 1, when B proves bad for everybody ; but the logic of pivotal voting leads to a surprising conclusion.

In the framework of FR, pivotal and sincere voting yield the same results because voting is over alternative policies not over candidates: Whether or not a voter is strategic, it is a weakly dominant strategy for him to vote the alternative he prefers in expectation. The others' strategies are not relevant given that values are private. The voter is presented with a choice between lotteries whose outcomes are independent of other voters' types and how they vote. As explained in the introduction, our model also has private values, but the strategic element has bite because a pivotal voter can infer the types of other voters and therefore what candidate K is likely to do at the next date.

Inefficiency a la' Fernandez and Rodrik is possible only when the efficient policy hurts the current majority on average; we refer to this as the type I inefficiency. In terms of our model, an idiosyncratic swing in rankings must large enough to change the balance of power and to reduce the erstwhile majority to a minority. If p is small then there cannot be any inefficiency. We shall see that, in contrast,

[5]The majority is most likely to switch only when there is an common shock in favor of policy 1.

the inefficiency displayed by our pivotal voter is more pervasive— it can happen even when idiosyncratic shocks are extremely unlikely to change the balance of power (in expectation). (As noted earlier, the common shocks tilt all voters towards the unbiased candidate.) In terms of our results, our main proposition shows that the inefficiency persists even when the probability p of any voter changing sides is small, and there is a substantial probability of a common shock, which helps K. We refer to this as type-II inefficiency.

In summary, our work differs from FR on three counts — model, methodology, and implications. First, in our model voters choose between candidates rather than policy alternatives. Since one candidate offers a policy that is conditional on the electorate's final rankings of policies, a pivotal voter must take into account the voting strategies of others; as a consequence he behaves differently from the sincere voter. Finally, from a substantive point of view the pivotal argument implies that inefficiency remains even when the "idiosyncratic" shock (one that does not affect all voters the same way) is unlikely to precipitate a large change., i.e. p is small.

2.2 The model

We consider a simple two-period model with voters in $S = \{1, 2, ..., 2n + 1\}$, $(n > 1)$and a set of policies $A = \{0, 1\}$. At date-0, nature draws each voter's type $t_i^0 \in \{0, 1\}$ from a Bernoulli distribution with $\Pr\{t_i^0 = 0\} = q > 1/2$. The voter's type on a particular date specifies the policy he prefers on that date. Elections are held at date-0 as well. There are two candidates to choose from — B and K. Candidate B is known to derive private benefits from the policy 0 [6], while K is an unbiased candidate who promises to maximize social welfare. After the elections, each voter's type changes according to a stochastic process described below. Once each voter's

[6] $q > 0.5$ is without loss of generality as long as there is a candidate who derives private benefits from implementing the policy that is favored by the majority at date 0.

date-1 type t_i^1 is determined, K implements the ex-post social optimum policy if he has been elected at date-0; if B was elected, he chooses policy 0. If $a \in A$ is the policy implemented at date-1, the utility of voter i is given by

$$u_i(a;\ t_i^1) = \begin{cases} 1 \text{ if } a = \ t_i^1 \\ 0 \text{ otherwise} \end{cases}.$$

Figure 1 summarizes the temporal structure of the game.

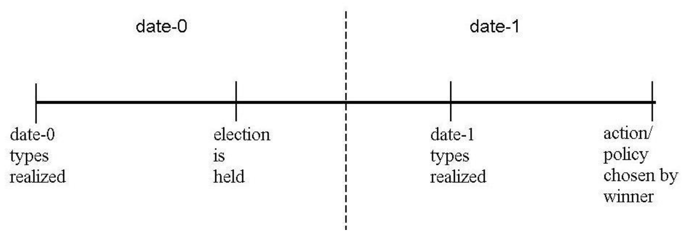

Figure 2.1: Timeline

Voter's types change over time as new information is revealed to them. Informative signals arrive according to the following process. With probability δ, a public signal favoring one of the policies arrives, forcing everyone to one side; all voters then prefer policy 0 with probability π or policy 1 with probability $1-\pi$. With probability μ, idiosyncratic private signals arrive, leading to independent changes in voters' rankings. These signals favor policy 0 with probability ϕ and policy 1 with probability $1-\phi$. If the signal favors policy 0, each voter of type $t_i^0 = 1$ (henceforth, type-0) switch to policy 0 with probability p independently of the others, while the voters of type-0 stick to their original preferences ; if it supports 1 then all 1 types stay put but each type-0 changes to 1 with probability p independently of the others. Lastly, no information arrives with probability $1-\delta-\mu$; in that case we have $t_i^1 = t_i^0 \forall i$. Later we shall make the simplifying assumption that $\pi = \phi$; this is in no way important for our results and merely permits cleaner algebra and succinct interpretations of the derived results. In a realistic case, one would expect δ to be small relative to μ-

it is more likely that individuals do not switch en-masse but rather in response to information that each voter chooses to interpret as either strong enough to switch to the other side or too weak to make a difference. All results continue to hold inspite of the relative magnitudes of δ and μ as long as $\delta < \mu$. In the rest of the paper, we refer to the δ-event as being a common shock, and the μ-event as an idiosyncratic shock. Figure 2 summarizes the signal structure.

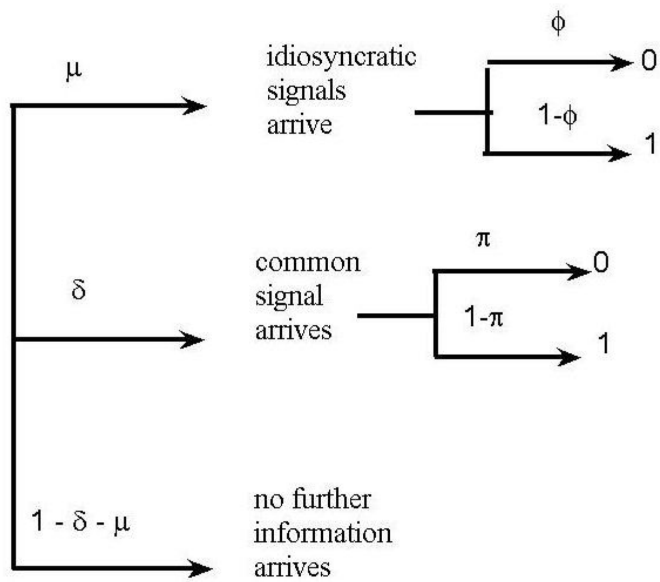

Figure 2.2: Schematic Representation of Changing Types

Each voter's type at date-0 is private information; everything else, including the rationality of voters and the stochastic process for the change of types from date-0 to date-1 is common knowledge. It also seems realistic to say that a voter does not know the exact types of the others, but has a general sense of the dispersion in opinion. Since our focus will be on symmetric pure strategy equilibria, assuming types are private information will allow voters to rationally condition on the pivotal scenario.

2.2.1 The sincere voter

We use the sincere voter to shed light on the forces that influence the decisions of a pivotal or sophisticated voter. The sincere voter does not condition on the state in which he is pivotal, but instead picks the candidate who, given the unconditional distribution of date-0 and date-1 types, is more likely to agree with him at date-1. He uses a weakly dominated strategy: when he is not pivotal his vote does not affect the result, and when he is indeed pivotal, his vote may not coincide with that of the rational pivotal voter. While falling short of rationality in one of many ways. the sincere voter provides a very useful benchmark against which to compare the results for the pivotal voter, and also facilitates comparison with previous work. An interesting point emerges from the comparison: When all voters are strategic the net outcome could be worse. As we shall see, the sincere voter can only generate one of the two forms of inefficiency discussed below. We should also like to remark at this point that our sincere voter is similar to Harsanyi's rule-utilitarian voter[31]. A rule-utilitarian is one who votes according to the rule that maximizes social utility if everyone else follows the same rule. This concept is further extended by Feddersen and Sandroni [22], and by Coate and Conlin [14] to include *group* rule-utilitarians, who choose the action that is best for the *group* when everybody in the group follows it. Our sincere voter chooses like the group-rule utilitarian voter.

Since $q > 0.5$, the vote of the sincere type-0 voter determines the outcome of the election in a large population. We find in Proposition 1 below that the type-0 voter supports K either when (1) type-0's are expected to be in a majority at date-1 , or (2) type-1's are expected to be in a majority at date-1, but the typical type-0 voter is very likely to switch preferences at the next date i.e. p is 'high'. The only case in which he votes B is the one where the majority is likely to be at 1 at the next date but any given voter is very likely to stay put i.e. p is 'low'. In other words, he prefers to commit and safeguard his interests today as he might not have enough support to do so at the next date.

Proposition 1 *When* $q \in (0.5, 1)$ *and* n *is large enough, sincere type-0 voters vote* B *if and only if* $q(1-p) < 1/2$ *and* $p < \frac{1}{2}(1 - \frac{\delta}{\mu})$.

Otherwise they vote K. The sincere type-1 voters always vote K. When type-0 voters support B, the probability that B wins goes to 1 *as* $n \to \infty$.

Proof: The expected utility of a sincere type-0 voter i when B is elected is given by

$$U_i(B, 0) = 1 - \mu(1-\phi)p - \delta(1-\pi)$$

; the second term corresponds to the loss incurred when the voter sways to an idiosyncratic 1-signal and the third term is the loss due to an common 1-signal. When K is elected, the expected utility is

$$U_i(K, 0) = 1 - \mu(1-\phi)p\Lambda_{01} - \mu(1-\phi)(1-p)\Lambda_{10} - \mu\phi\Pi_1$$

, where Λ_{01} is the probability that the date-1 majority is at 0 when type-0 voter i switches to 1 in response to an idiosyncratic 1-shock, Λ_{10} is the probability that the date-1 majority is at 1 when voter i ignores the idiosyncratic 1-signal and Π_1 is the probability that the majority stays at 1 despite an idiosyncratic 0-signal.[7] The three negative terms correspond to the three potential sources of loss under K. Note that losses can only be incurred under K when idiosyncratic signals arrive. Conditioning on an idiosyncratic 1-signal, the probability that an arbitrary voter supports policy 0 at date 1 is $q(1-p)$ for large n. If we define the random variable $X_{01} \sim \text{Binomial}(2n, q(1-p))$as the number of voters (barring one) who support policy 0 at date-1 following an idiosyncratic 1-signal, then $\Lambda_{01} = \Pr\{X_{01} > n+1\} = \Pr\{\frac{1}{2n}X_{01} > \frac{1}{2} + \frac{1}{2n}\}$. The Weak Law of Large Numbers guarantees that

[7] $\Lambda_{01} = \sum_{j=n+1}^{2n} \binom{2n}{j} \theta^j (1-\theta)^{2n-j}, \Lambda_{10} = \sum_{j=0}^{n-1} \binom{2n}{j} \theta^j (1-\theta)^{2n-j}, \Pi_1 = \sum_{j=0}^{n-1} \binom{2n}{j} \psi^j (1-\psi)^{2n-j}, \theta = q(1-p), \psi = q + (1-q)p$

$\lim_{n\to\infty} \Pr\left\{\left|\frac{1}{2n}X_{01} - q(1-p)\right| < \epsilon\right\} = 1$for any $\epsilon > 0$. If $q(1-p) > 0.5$, there exists an $n$large enough and an $\epsilon > 0$so that

$$\Pr\left\{\left|\frac{1}{2n}X_{01} - q(1-p)\right| < \epsilon\right\} \leq \Pr\left\{\frac{1}{2n}X_{01} > \frac{1}{2} + \frac{1}{2n}\right\} \equiv \Lambda_{01} \leq 1.$$

It follows that

$$q(1-p) > 0.5 \Rightarrow \lim_{n\to\infty} \Lambda_{01} = 1 \text{ and } \lim_{n\to\infty} \Lambda_{10} = 0 \text{ ; } q > 0.5 \Rightarrow \lim_{n\to\infty} \Pi_1 = 0$$

For large n and $q(1-p) > 0.5$, the type-0 voter then votes K, the unbiased candidate for all $\delta \geq 0$. When $q(1-p) < 0.5$, the date-1 majority is expected to be at 1 and for large n, $\Lambda_{01} \to 0$, $\Lambda_{10} \to 1$ and $\Pi_1 \to 0$. The type-0 voter then votes K if $p > \frac{1}{2}\left(1 - \frac{\delta(1-\pi)}{\mu(1-\phi)}\right)$. Under $\pi = \phi$, this reduces to $p > \frac{1}{2}\left(1 - \frac{\delta}{\mu}\right)$.

Let us now turn to the sincere 1-voter. His expected utility from B is given by

$$U_i(B,1) = (\delta\pi + \mu\phi p)$$

; he gets a utility of 1 iff he switches to 0 himself, in response to either an idiosyncratic or an common 0-signal. His utility from K is

$$U_i(K,1) = \delta + \mu\phi p\Pi_{00} + \mu\phi(1-p)\Pi_{11} + \mu(1-\phi)\Lambda_1$$

, where Π_{ab} is the probability that, following an idiosyncratic 0-signal, the majority is at a when our voter is at b; $\Lambda_1$is the probability that, following an idiosyncratic 1-signal, the majority is at 1. By arguments similar to the ones made for the sincere 0-voter above,

$$q > 0.5 \Rightarrow \lim_{n\to\infty} \Pi_{00} = 1 \text{ and } \lim_{n\to\infty} \Pi_{11} = 0$$

$$q(1-p) > 0.5 \Rightarrow \lim_{n\to\infty} \Lambda_1 = 0 \; ; \; q(1-p) < 0.5 \Rightarrow \lim_{n\to\infty} \Lambda_1 = 1.$$

When $q(1-p) < 0.5$, we can therefore reduce the decision for large n to one of comparing $\delta\pi + \mu\phi p$ with the quantity $(\delta + \mu(1-\phi) + \mu\phi p)$ that he gets from K. Since $\delta(1-\pi)+\mu(1-\phi) > 0$, he votes K. When $q(1-p) < 0.5$, he compares $\delta\pi+\mu\phi p$ from B with $\delta + \mu\phi p$ from K and votes K. Thus the sincere type-1 voter votes K for all values of $q \in (0.5, 1)$and all $p \in (0, 1)$.

When either (1) $q(1-p) > 0.5$, or (2) when $q(1-p) < 0.5$and $p > \frac{1}{2}\left(1-\frac{\delta}{\mu}\right)$, both the type-0's and 1's vote K and he wins with probability 1. When $q(1-p) < 0.5$ but $p < \frac{1}{2}\left(1-\frac{\delta}{\mu}\right)$, the type-0's vote B while the 1's vote K; since $q > 0.5$, B is the more likely winner; his exact probability of winning is $\Sigma_{k=n+1}^{2n+1}\binom{2n}{k}q^k(1-q)^{2n-k}$, which is greater than 0.5 for all n and tends to 1 as $n \to \infty$. ■

Figure 3 summarizes the behavior of the type-0 voter for large enough n. In regions II and IV, the type-0 voter elects K as he expects to remain in the majority if there is an idiosyncratic shock, and has nothing to lose by voting K; when there is an aggregative 1-signal, he is better off with K as B would still continue to implement the alternative 0, which gives B private benefits; with an common 0-signal, both B and K implement 0. In region I, the date-1 majority is expected to prefer policy 1, but since p is high each type-0 expects to switch and be in the subsequent majority. So he votes for K, who always picks the right alternative when there is an aggregative shock. Finally in region III, the sincere type-0 voter picks the socially suboptimal candidate B because the majority is likely to prefer 1 at the next date, but given that p is small he would probably stay put at 0.

The sincere voter is nothing but the representative agent of the group, one who maximizes the value of the group[8]. This is easy to see - all agents of a group are ex-ante identical, and the probabilities that appear in the decision of the sincere voter are

[8]The sincere voter's behavior is akin to the notion of group rule utilitarianism introduced by Harsanyi [31].

the expected proportions in the decision of the representative agent. The probability of switching, for example, is now to be interpreted as the expected proportion of type-0's who switch when there is an idiosyncratic 1-signal.

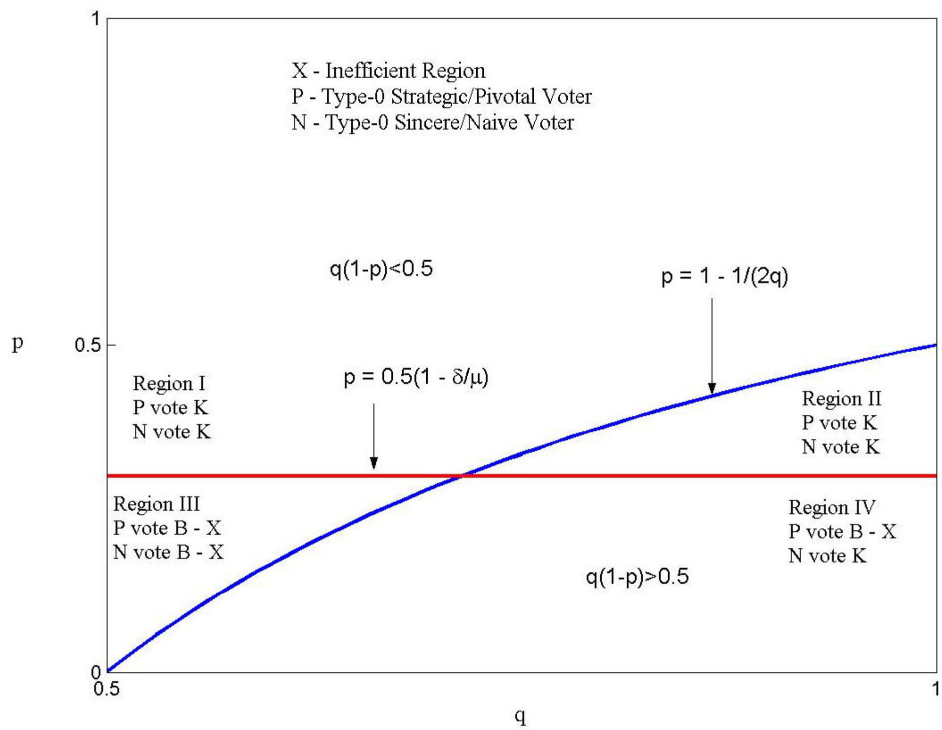

Figure 2.3: Voting Behavior of the Type-0 Voter

2.2.2 The pivotal voter

The situation is very different when we require the strategies to constitute a Nash equilibrium. We conjecture the following Nash equilibrium in pure strategies : each voter of type-0 votes B, and all type-1's vote K; then we solve for the range of parametric values where this is indeed the case.[9] Let $U_i(c \mid c_0, c_1;\ t_i^0,\ \text{piv})$ denote the utility of the pivotal voter of date-0 type t_i^0 when he votes for candidate c, the type-0's vote for candidate c_0 and the type-1's vote for candidate c_1. Consider a pivotal type-0 voter. When he is pivotal the utility of voting for B is the same as that for the sincere

[9]We ignore trivial equilibria in which all voters support the same candidate. Our focus is on symmetric equilibria in which all voters of a type vote the same way.

voter:

$$U_i(B \mid B, K;\ 0,\ \text{piv}) = 1 - \mu(1-\phi)p - \delta(1-\pi)$$

The utility of voting for K is now different :

$$U_i(K \mid B, K;\ 0,\ \text{piv}) = \delta + (1-\mu-\delta) + \mu\phi + \mu(1-\phi)\left\{(1-p)^{n+1} + p\right\}$$

Utility from K is 1 whenever there is an common signal, no signal, or an idiosyncratic 0-signal. If an idiosyncratic 1-signal arrives, the type-0 voter gets 1 if nobody switches or if he himself switches. The pivotal voter prefers B when

$$U_i(B \mid B, K;\ 0,\ \text{piv}) > U_i(K \mid B, K;\ 0,\ \text{piv}) \Leftrightarrow \delta(1-\pi) + \mu(1-\phi)p < \mu(1-\phi)(1-p)\{1-(1-p)^n\} \tag{2.1}$$

Under the assumption $\pi = \phi$, the condition reduces to

$$f(p, n) = 1 - 2p - (1-p)^{n+1} > \frac{\delta}{\mu}.$$

Figure 4 shows how $f(p, n)$ compares to a threshold of δ/μ for different values of n. Since $f(n+1, p) > f(n, p)$ for all p in $(0, 1)$,the range of values of p for which the committed candidate wins is growing with n; as $n \uparrow \infty$, the negative term $(1-p)^{n+1}$ goes rapidly to 0, and the condition reduces to $1 - 2p > \frac{\delta}{\mu} \Leftrightarrow p < \frac{1}{2}(1 - \frac{\delta}{\mu})$.[10] There exists a $p > 0$ satisfying the above if $\delta < \mu$, i.e. if the probability of an common switching signal is less than that of a idiosyncratic signal. The only requirement is that a signal that causes the entire population to switch to one side is less likely than a signal that causes voters to switch idiosyncratically, - surely a reasonable assumption.

[10] If we turn to the case when there is no common shock, i.e. $\delta = 0$, then 0's vote B if $(1-p)^{n+1} < 1 - 2p$. Using a quadratic Taylor series expansion, a sufficient condition for this is $1 - (n+1)p + \frac{(n+1)n}{2}p^2 < 1 - 2p$, or $\frac{(n+1)n}{2}p < n - 1$. i.e. $p < \frac{2(n-1)}{n(n+1)}$. For 7 voters, for example, this effect is observed for $p < 1/3$. Thus with $\delta = 0$, the bias towards commitment is very much a reality even with relatively few voters.

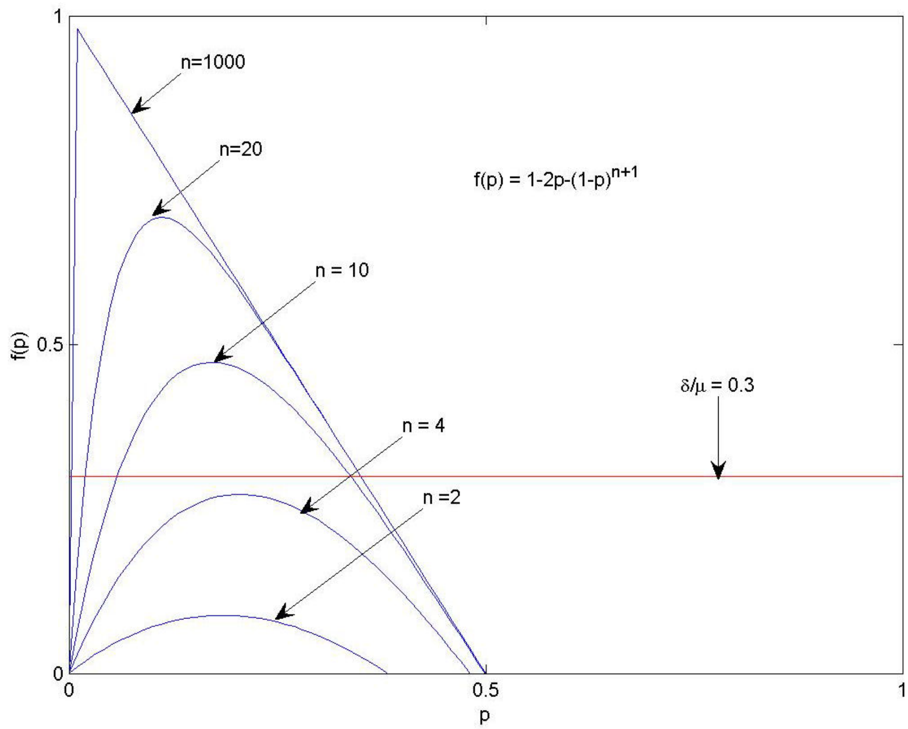

Figure 2.4: Small n behavior

So δ may be quite large, even close to but less than one-half, provided that it is less than μ.

For n large, (2.1) reduces to

$$\delta(1-\pi)+\mu(1-\phi)p<\mu(1-\phi)(1-p) \tag{2.2}$$

The LHS of (2.2) is the loss from voting for B and getting him elected - the first term is the loss when the entire population switches to 1 at date 1, the second is when voter i responds to an idiosyncratic signal and finds himself on the wrong side vis-a-vis B. The RHS is the loss when K is elected. Note that the pivotal voter reacts very differently to the possibility of an idiosyncratic switch depending on who is in power- B or K. When B is in power, the loss is when voter i himself switches whether or not others switch. With K in power, i no longer fears switching even if others don't; what he fears is staying put when others switch.

Proposition 2 *When $q \in (0.5, 1)$ and n is large enough, there exists a Nash equilibrium with pivotal voters in which all type-0's vote B and all type-1's vote K when $p < \frac{1}{2}(1 - \frac{\delta}{\mu})$. B wins with a probability that tends to 1 in large populations. When $p > \frac{1}{2}(1 - \frac{\delta}{\mu})$ all voters (type-0 and 1) prefer K to B.*

Proof: The argument above shows that for large enough n, the pivotal type-0 voters prefer B if $p < \frac{1}{2}(1 - \frac{\delta}{\mu})$ and the type-1 voters conform to the conjectured equilibrium strategy. Since $q > 0.5$, the Weak Law of Large Numbers ensures that almost surely, a majority of voters support 0 on election day and therefore B wins.

Finally to show that the above conjecture indeed gives us a NE, we show that the pivotal type-1's will vote K. Expected utility from B is

$$U_i(B \mid B, K;\ 1,\ \text{piv}) = \delta\pi + \mu\phi p$$

, while that from K is

$$U_i(K \mid B, K;\ 1,\ \text{piv}) = (1 - \mu - \delta) + \delta + \mu(1 - \phi) + \mu\phi(1 - p)^{n+1} + \mu\phi p.$$

Under $\mu = \phi$, type-1 's vote K if $\phi(\mu + \delta) < 1$, which necessarily holds. ■

For large n, there are two forms of inefficiencies illustrated above. The first, which corresponds to region III of Fig. 3 and is exhibited by both sincere and pivotal voters, was discussed in Section 2.1. The difference between the sincere and the pivotal voter is in region IV: the pivotal voter prefers the committed candidate even when, following an idiosyncratic shock, he expects to remain in the majority. What drives the fear of the pivotal voter is that his decision is conditional on himself being pivotal, thereby unraveling the effect of q. In contrast to the decision of the sincere voter for large n, his decision depends only on the value of p and not that of q. The interaction among pivotal voters enters through the size of the population: when n is large it is almost certainly the case that, starting from a pivotal situation, the pivotal 0-type voter will be in minority if he does not switch following an idiosyncratic signal. The proposition below summarizes this.

Proposition 3 *When* $p < \frac{1}{2}(1 - \frac{\delta}{\mu})$ *and* $p < 1 - \frac{1}{2q}$ *, sincere voting results in election of the unbiased candidate, while pivotal voting almost surely results in the election of the committed candidate. If B wins, he implements a suboptimum policy with a probability that is bounded below by the probability of an common 1-type shock* $\delta(1 - \pi) > 0$*. All sincere voters vote K, whereas the pivotal type-0 and 1 vote B and K respectively. Thus for large n, B wins with a probability arbitrarily close to unity if voters are fully rational; K wins if they are sincere.*

Remark *For all n, not necessarily large, and* $q > 0.5$ *the probability of an inefficient decision is bounded below by* $\delta(1 - \pi)/2$*. This follows since B wins with a probability bounded below by* 0.5 *for all n.*

Finally, returning to Figure 4 , we note that for small n, it is not the case that the pivotal type-0's prefer B for all values of p less than a certain threshold. For very small p the condition reduces to $\delta < 0$, which is impossible. What is the reason for this difference? Recall that the pivotal agent's fear is of being left behind - of not switching to the other side while at least one other person on his side defects and destroys a fragile majority. But the fewer the voters the less likely is this fear, and so the dominant fear is that of an common 1-shock that would render B undesirable for all voters.

2.3 Varying the Voting Rule

The previous sections proceeded under the assumption that the voting rule used at date-0 is the same as the decision rule used at date-1 by the unbiased candidate K. Recall that an interpretation of our framework, one that we mention earlier, is the choice between acting now or waiting; with this interpretation it is indeed natural to suppose that the voting rule at 0 and K 's rule at 1 are the same. But if we think of it as an electoral contest, one is naturally led to investigate the properties when the two rules are different. This section accordingly looks at an $m - rule$ at date 0, to be defined shortly. In this section, we examine the behavior of the pivotal voter for a range of such voting rules.[11] We shall continue to focus our attention on large electorates. Suppose now that the committed candidate B and unbiased candidate K contest in an election where B wins if he receives a fraction $m > \frac{1}{2}$ or more of the votes and K wins otherwise.[12]

As one would expect, increasing m from $\frac{1}{2}$ makes it difficult for B to win, and helps mitigate the inefficiency generated by the pivotal voter. However the inefficient equilibrium turns out to be robust in a large range of (m, p) values. To see this, let us re-examine the conjectured equilibrium in which type-0's vote B and type-1's vote K, and analyze the pivotal type-0 voter's decision. Recall that $q > 0.5$ means that in the symmetric equilibrium the vote of the type-0 voters is the deciding vote. Under the new voting rule, the pivotal voter conditions on the state of the world in which k voters are of type-0 , where $\frac{k}{2n+1} < m \leq \frac{k+1}{2n+1}$. If $m(1-p) > \frac{1}{2}$ and n is large, by the law of large numbers we know that it is highly likely the majority at date-1 will remain at 0. The pivotal voter is therefore not conditioning on a precarious majority and this allays his fear of being left behind while the majority switches to policy 1 at the next date. Under this condition we thus find that the inefficient equilibrium conjectured above fails to exist. By providing a buffer between the point the pivotal voter conditions upon and the simple majority, the $m - rule$ reduces the type-0 voter's incentive to protect himself by voting for the committed candidate B. When the above condition is not met the inefficient symmetric equilibrium survives.

More formally, the utility of the pivotal voter of type-0 when he votes B , given

[11] At the risk of being redundant, we should like to emphasize that while the voting rule has been altered, candidate K remains committed to implementing the policy that is preferred by the majority at date-1.

[12] A similar analysis with Bas the status quo may also be performed.

that all 0's vote B and K 's vote 1 following the dictates of the equilibrium, is identical to that in the previous section and is given by

$$U_i\left(B|B,K;0,\,piv\right)=1-\mu\left(1-\phi\right)p-\delta\left(1-\phi\right)$$

, where we have assumed $\phi=\pi$ for simplicity. The utility of the pivotal voter above from voting K can be written as

$$U_i\left(K|B,K;0,\,piv\right)=\delta+\left(1-\mu-\delta\right)+\mu\phi+\mu\left(1-\phi\right)\left\{p\Theta_1+(1-p)\Theta_0\right\}$$

, where Θ_a is the probability that the majority prefers policy $a\in\{0,1\}$ at date-1, conditional upon starting at date-0 from a situation in which $k+1$ voters are of type-0 and the rest of type-1. The four terms correspond to the situations in which the pivotal voter i of type $t_i^0=0$ gets a utility of 1 by voting B and thereby electing him. The first term δ is for an common shock, when all voters agree on the policy and B implements it ; the next is when no additional information arrives and i continues to be in a $m-$supermajority and thus in a simple majority at date-1; the third term is for the common shock towards 0 , when the majority for 0 is bolstered; the last terms is for the idiosyncratic shock towards 1— i gets 1 iff he switches and the majority swings to 1 or if he stays put and so does the majority. Depending on whether p is large or small relative to the value m , Θ_0 or $\Theta_1$is much larger than the other in the limit. Let us first consider the case when $p<1-\frac{1}{2m}$. By the Weak Law of Large Numbers, for any $\epsilon>0$,

$$\lim_{n\to\infty}P(|\text{fraction of popln. supporting policy 0 at next date }-m(1-p)|<\epsilon)=1$$

It follows then that if we choose a small enough ϵ then $\lim_{n\to\infty}\Theta_0=1$ and $\lim_{n\to\infty}\Theta_1=$

0. Therefore, the utility from voting for K converges to

$$\lim_{n\to\infty} U_i\left(K|B,K;0,\, piv\right) = 1-\mu\left(1-\phi\right)p$$

For the pivotal voter to prefer candidate B, it is therefore necessary for $\delta\left(1-\phi\right)<0$. Since this is not true, the conjectured equilibrium does not exist when the electorate is large and $p<1-\frac{1}{2m}$.

When $p>1-\frac{1}{2m}$, a similar argument gives $\lim_{n\to\infty}\Pi_0=0$ and $\lim_{n\to\infty}\Pi_1=1$. The utility from voting for K then converges to

$$\lim_{n\to\infty} U_i\left(K|B,K;0,\, piv\right) = 1-\mu\left(1-\phi\right)\left(1-p\right)$$

In this case, the pivotal voter prefers voting B when $\delta<\mu(1-2p)$, or equivalently $p<\frac{1}{2}(1-\frac{\delta}{\mu})$. Note that this constraint is identical to the one derived for the simple majority rule. Figure 5 shows how the pivotal *type* $-$ 0 voter's relative preference for each candidate varies with the parameter p for different $m-rules$. For each m, the voter prefers candidate B to K for values of p where the curve lies above 0. Note that as $m\to\frac{1}{2}$, the inefficiency is possible for smaller and smaller values of p. When $m=\frac{1}{2}$, any $p>0$ can give rise to the inefficient equilibrium for large enough electorates; this is the content of the previous section.

Proposition 4: *The inefficient equilibrium exists for large electorates when* $1-\frac{1}{2m}<p<\frac{1}{2}(1-\frac{\delta}{\mu})$*. The set of values of* p *which supports the inefficient equilibrium shrinks as* m *increases. If* $m>\frac{\mu}{\mu+\delta}$*, the inefficient equilibrium does not exist for any value of* p*.*

Proof : We have already verified that the pivotal *type* $-$ 0 voters will conform to the behavior of the inefficient equilibrium when p is in the range specified. Now we show

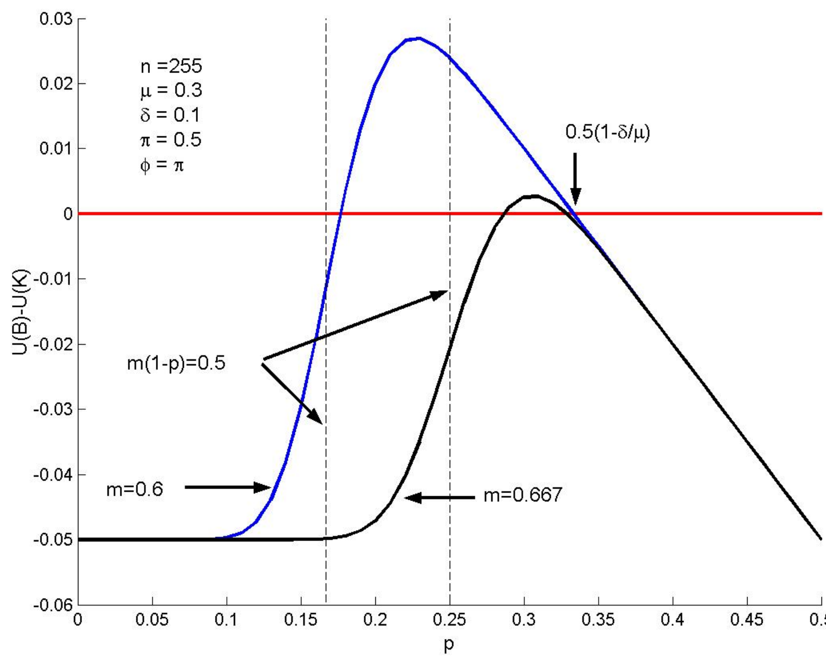

Figure 2.5: Simulation Results for m-Rules with $m > \frac{1}{2}$

that the pivotal $type-1'$s will vote K. Expected utility from B is

$$U_i(B \mid B, K;\ 1,\ \text{piv}) = \delta\pi + \mu\phi p$$

while that from K is

$$U_i(K \mid B, K;\ 1,\ \text{piv}) = \delta + \mu(1-\phi)\Pi_1 + \mu\phi p$$

Under $\phi = \pi$, and $p > 1 - \frac{1}{2m}$, $type - 1$'s vote K if $\phi < 1$, which necessarily holds.■

It follows immediately from the last inequality in the proposition above that the inefficient equilibrium cannot arise for the unanimity rule ($m = 1$). While our framework is not directly comparable to the information aggregation models, it might be interesting to note that this result contrasts with the inferiority of the unanimity rule documented previously.

2.4 Implications for candidate entry : a discussion

So far our work proceeded under the assumption that candidate preferences were common knowledge. If we now choose to move back one step in time, we might ask which candidates actually enter an election. Let $c > 0$ denote the cost of entering an election; this includes the cost of filing nomination, campaigning, etc. The candidate could be either one who cares about his legacy and when in office picks the alternative that the majority prefer, or a partisan who benefits from one of the two alternatives independently of the electorate's rankings. Let us consider a simple extension of the model where q has not been revealed when the candidates decide to enter. Suppose q is equally likely to be 0.25 and 0.75 and let b denote the benefits from being in office for any type of candidate. If there is a potential candidate of each type and $b > 2c$, then in equilibrium one 0-candidate and a 1-candidate will both choose to enter. When $q = 0.25$ the 1-candidate wins with a probability arbitrarily close to 1 for a large enough n ; with $q = 0.75$ the 0-candidate wins. The legacy candidate will therefore choose not to enter the fray as his expected gain from entry will fall short of the cost $c > 0$ — He cannot win the election irrespective of what value of q is drawn. As political satirist Frank Hubbard once said, “We'd all like to vote for the best man but he's never a candidate”. We see that extreme candidates do better than one who is completely unbiased and proposes to implement the ex-post majority's preferred policy. In a loose sense this hints that gradual resolution of uncertainty might have something to do with the entry of biased candidates.

2.5 Conclusion

Many papers look at information aggregation in voting and the role of pivotal voting. They almost always have the feature that voters' rankings of policies do not change between the time they vote and the time a policy is implemented. The role of the electoral process is to aggregate their private signals. We relax this framework and allow rankings to change; this leaves the electoral system prone to widespread inefficiency. This paper illustrates two forms of inefficiency. When voters who are in a majority today are more likely to be in a minority tomorrow, they oppose social-welfare improving policies. This requires a probability of idiosyncratic switching large enough to reduce the ex-ante majority to an ex-post minority. Perhaps a large range of electoral situations is better described by a model in which the probability of voters changing idiosyncratically is small. This, one might even assert, is the rule rather than the exception. We should hope that in such a case the inefficiency will be mitigated, if not eliminated. We argued above that this is not the case — In the unique (informative) symmetric Nash equilibrium, voters prefer to elect the ideologue rather than elect an unbiased candidate, who waits for all information to be revealed and thereafter takes the optimal decision. The key to understanding this paradoxical result is that the *pivotal* voter finds himself in a fragile majority that is easily overturned; even though such a situation is (unconditionally) unlikely, he bases his vote on this situation and commits to the alternative that he currently prefers. This continues to hold even if there is a large chance that everybody will dislike the committed candidate's choice due to an common shock.

Chapter 3

Efficient Information Aggregation in Elections with Differentially Informed Agents [1]

3.1 Introduction

The idea that voting is a means to aggregate private information and ensure a better collective choice is one that dates at least as far back as Condorcet. An extensive literature in political economy has since examined the veracity of the so called Condorcet Jury Theorem, according to which a group of imperfectly informed voters is more likely to select the "better" of two alternatives than any single agent acting alone.

In this work, we propose and analyze an election format which permits an electorate to exploit the fact that voters may be differentially informed. We consider a voting framework where voters can choose when to cast their ballot and the intermediate tally of votes is publicly observable. By strategically choosing the time at

[1] Co-authored with Konstantinos Rokas.

which they cast their vote, better informed voters can now communicate their signals to those less sure of themselves. We find that this added flexibility vastly improves information aggregation in common value environments.

The voting environment described above is not uncommon. Voting in legislatures and corporate boardrooms is often done by a show of hands and uncertain voters may delay casting their votes to gauge the level of consensus in the room. Perhaps the most notable instance of such an election is the manner in which the Democratic National Party conducts the Iowa caucus. The caucus site is partitioned into sections according to candidate and for the first thirty minutes after the caucus begins voters can decide which "preference group" to join. The proportion of voters in each group then forms the basis for the allocation of delegates to each candidate. We should point out however, that while these are examples of open voting, the mechanism we examine can be implemented without compromising the privacy of the voter. Efficiency can be achieved while preserving anonymity since the only information released publicly is the intermediate vote count.

Our findings are of particular interest to the ongoing debate on the effects of exit polls on election outcomes. A recent survey [43] shows that laws pertaining to the broadcast of exit polls on election day vary considerably across countries. While most arguments in favor of exit polls are on the grounds of transparency and a constitutionally mandated freedom of speech, we believe our efficiency results may be viewed as a novel reason why exit polls (or even official reports of intermediate results) should be instituted.

In a seminal contribution, Austen-Smith and Banks [5] show that in a common value setting, simultaneous elections aggregate information efficiently only under knife-edge conditions. In fact, they aggregate information efficiently under a simple majority rule if and only if all voters have the same quality of information. The inefficiency in these elections stems from the fact that in equilibrium, strategic voters who condition on being pivotal do not find it rational to vote according to their signal. In contrast, the election we consider is outcome equivalent to the first best (that is, the electorate chooses the same outcome as it would choose if all information were publicly available) regardless of the size of the electorate. Efficiency is obtained because under mild conditions, there exists an order of voting, which when followed in equilibrium (i) allows better-informed voters to communicate their signals without inadvertently deciding the election in favor of the wrong candidate and (ii) reveals all relevant information while leaving a sufficient number of voters to swing the election in either direction.

Information aggregation has also been studied in the context of sequential or roll-call elections where players vote according to a predetermined (exogenously specified) order. Fey [27], Wit and more recently Ali and Kartik [1] have examined the possibility of bandwagons (or informational cascades) developing in sequential elections. Dekel and Piccione[2] also study sequential voting games but illuminate a different issue. They prove that any symmetric responsive strategy profile is an equilibrium of the sequential election if and only if it is also an equilibrium of the simultaneous election. This body of work has focussed on inefficiency in sequential voting. In contrast, we will argue that our proposed election format is less vulnerable to these inefficiencies. Two natural equilibrium refinements that allow us to do so; first, we restrict attention to equilibria that survive the introduction of small costs of voting and second we focus on equilibria that are insensitive to the choice of off-equilibrium beliefs. These characteristics will play a central role in distinguishing efficient equilibria from others.

Our work also contributes to another strand of the literature on information ag-

[2]Our work was motivated in part by a quote from Dekel and Piccione, pg. 35 which reads "On the negative side, it [our result] completely demolishes any hope of obtaining strong conclusions about endogenous timing...". While we don't believe Dekel and Piccione meant to use the phrase "endogenous timing" as we have in this paper, we would like to thank them for leading us to this project.

gregation that begins with Feddersen and Pesendorfer [25]. Using pivotal voter arguments, (FP) argue that the unanimity rule, commonly used by juries to protect the innocent from being convicted, is in fact more likely to result in an incorrect judgement than the simple majority rule. In defense of the unanimity rule, Coughlan [15] shows that in a common value setting, adding a pre-play communication phase (such as a non-binding straw poll[3]) allows all information to be aggregated prior to the jury's vote and renders all voting rules outcome-equivalent. We present an alternate solution. Our election achieves first best regardless of the voting rule. We find that information can be aggregated through the voting process itself, without having to resort to a cheap-talk phase. This is an important finding for at least two reasons. First, pre-play communication is not always possible. Voters may be geographically separated. The electorate may be large. Communicating with every member of the jury can be costly[4]. Second, an experimental study conducted by Guarnaschelli and Palfrey [29] shows that voters do not use the information revealed in a straw poll correctly and election outcomes can vary with the voting rule. While convincing the reader that our election design performs better than a straw poll in the lab is outside the scope of this paper, we hope that in the light of [29], the reader will recognize the merit in examining alternative voting mechanisms. At the very least, one might expect that since information in our mechanism is signalled through the costly act of voting, the quality of aggregation will be different.

We examine several natural extensions. We find that our results generalize to scenarios where the voter is uncertain about the number of voters. We provide some insights about this election when there is diversity in voter preferences. Although progress in this direction has been slow, we are able to provide instances where the presence of partisans does not prevent our proposed design from achieving efficiency.

[3]The Republican National Party holds a non-binding straw poll in Ames, IA before the primaries.

[4]For a more extensive discussion of why relying on communication before voting is not always appropriate, we refer the reader to Nicola Persico [38].

Finally, we construct an example that shows how this election format can efficiently aggregate information even when voters receive multidimensional signals.

3.2 Model

We adopt the framework of Austen-Smith and Banks [5]. Let us consider an electorate with an odd number of voters indexed by $i = \{1, 2, ..., n \geq 3\}$. The underlying state of the world may be Left or Right, denoted by $\omega \in \{L, R\}$. The voters must collectively choose between two alternatives, l and r. We will focus on a common value setting where all voters prefer to pick the alternative that corresponds to the present state of the world. Hence, the utility function for every voter is given by $U(L, l) = U(R, r) = 1$ and $U(R, l) = U(L, r) = 0$. Unless otherwise specified, our voters dislike both types of errors equally.

Without loss of generality, we can assume that the prior probability of the state of the world being L satisfies $\pi_L \geq \frac{1}{2}$. For now, suppose each voter receives a binary-valued noisy private signal $s_i \in \{0, 1\}$ that conveys information about the state of the world, with $p_R = P(s_i = 1|\omega = R) \geq \frac{1}{2}$ and $p_L = P(s_i = 0|\omega = L) \geq \frac{1}{2}$; that is, 0-signals are more likely when $\omega = L$ and 1-signals are more likely when $\omega = R$. Later, we will consider an information structure where private signals will take one of $d+1$ values that will satisfy the monotone likelihood ratio property (MLRP). With binary signals, MLRP reduces to $\frac{p_L}{1-p_L} \geq \frac{1-p_R}{p_R}$, which follows immediately from $p_L, p_R \geq \frac{1}{2}$. We refer to the vector $s = [s_1, s_2, \ldots s_n]$ as the signal vector.

As discussed in the introduction, our focus in this work is the efficient aggregation of information in collective decision making. We define efficiency in an election as follows:

Definition 1 *An election format aggregates information efficiently if for every signal vector the equilibrium outcome when voters are privately informed is identical to the*

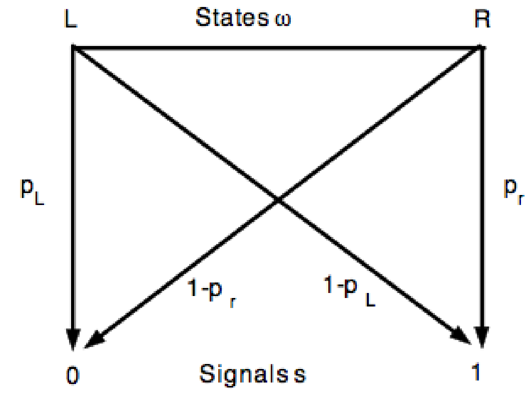

Figure 3.1: Signals and States

outcome when all information is public.

That is, the efficient outcome is the outcome preferred by each and every voter when all n signals are commonly known. We will adopt the Bayesian Nash Equilibrium (BNE) concept as our notion of equilibrium. A detailed discussion of the equilibrium concept appears later in this exposition.

We now define k_0 as the number of 0-signals in the vector s, so that $k_0 = \sum_{i=1}^{n}(1-s_i)$ where $s_i \in \{0,1\}$, and $k_0^*(n)$ as the minimum number of 0-signals necessary for l to be the utility maximizing choice. k_1 and k_1^* can be defined similarly for 1-signals. The quantity $k_0^*(n)$ is defined explicitly through the likelihood-ratio:

$$\begin{aligned}\beta(k_0,n) &= \frac{P(\omega = L|n_0 = k_0)}{P(\omega = R|n_0 = k_0)} \\ &= \frac{\pi_L p_L^{k_0}(1-p_L)^{n-k_0}}{(1-\pi_L)p_R^{n-k_0}(1-p_R)^{k_0}}\end{aligned}$$

The likelihood ratio is a function of the number of signals n and the number of 0-signals, k_0. We suppress the dependence on prior beliefs and the information structure for notational convenience. Then $k_0^*(n)$ is the unique $\widetilde{k} \in [1,n]$ that satisfies $\beta(\widetilde{k}-1,n) < 1 \leq \beta(\widetilde{k},n)$. Such a $\widetilde{k}$ need not exist. If $\beta(n,n) < 1$, alternative r is the efficient choice for all signal vectors and we define $k_0^*(n) = n+1$. Similarly, if

$\beta(0,n) > 1$, alternative l is the efficient choice for all signal vectors and we set $k_0^*(n) = 0$. It is clear that the aggregation problem is nontrivial only when $k_0^*(n) \in [1,n]$.

The next step is to describe the voting rule. Without loss of generality, we define the q-rule ($1 \geq q \geq \frac{1}{2}$) whereby alternative l is chosen if it receives $m = \lceil nq \rceil$ or more votes. Consider first a simultaneous election under this rule. In such an environment, it is rational for voters condition on the pivotal scenario to determine how they should vote. Conditioning on the pivotal state, each voter infers the distribution of private signals in the electorate and incorporates this information into her decision. We note that voters are sincere if they vote according to their private signal - $v(0) = l$ and $v(1) = r$. Does a rational voter choose to vote sincerely?

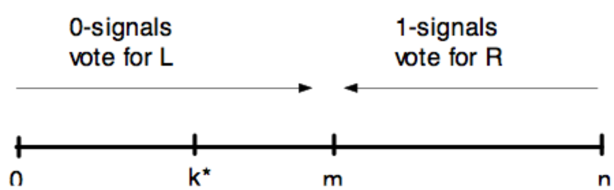

Figure 3.2: Informative voting

Figure 3.2 shows the information that a pivotal voter would infer in a hypothetical sincere voting equilibrium. If $k_0^* < m$, the pivotal type-1 voter would prefer to deviate and vote against her signal. Sincere voting can therefore be sustained in equilibrium only when $k_0^* = m$. This reasoning serves as a sketch of the proof for the result in [5]:

Proposition 1 ***(Austen-Smith and Banks, 1996):*** *Sincere voting in a simultaneous election is rational if and only if the q-rule is such that* $k_0^*(n) = m(n)$. *That is, the efficient outcome is achieved iff* $k_0^*(n) = m(n)$.

Two points are worth noting about this result. First, electoral rules are typically written into charters or constitutions well before any elections are held, whereas the quality of private information available to voters varies from one decision to the next. Second, the knife-edge condition for efficiency relies critically on a binary signal structure. Equivalent conditions do not exist in the finite-signal case. Hence simultaneous elections with finite electorates are generically inefficient. We will now examine an election format that successfully addresses these shortcomings.

3.3 Endogenous Timing

Consider the following framework - we partition time into a number of discrete periods $t = 1, 2, 3 \ldots$ and allow voters to choose the period in which to cast their vote. The intermediate tally of votes is broadcast at the end of each period. Voters are permitted to abstain. The election ends if no votes are cast for $T > 1$ consecutive periods[5]. The information structure is inherited from the model above.

The new format allows better informed voters to communicate their signals to others through their vote. Consider an extreme example in which $p_L = \frac{1}{2}$ and $p_R = 1$, so that the presence of a single voter with a 0-signal establishes with certainty that the state of the world is L. Suppose further that the simple majority voting rule is adopted.

As discussed in the previous section, a simultaneous election cannot aggregate this information structure efficiently except under knife-edge conditions on n and π_L. In our proposed election format, voters with 0-signals can vote for l in the first period while others wait. All relevant information regarding the signal vector is then revealed by the end of the first period and the efficient outcome can be ensured. There is no risk that such an order of voting will produce in an inefficient outcome as it only takes one 0-signal to guarantee that the state is L.

This reasoning holds in general. Consider any $n, \pi_L, p_L,$ p_R and q such that $0 < k_0^*(n) < m(n)$. Figure 3.3 illustrates graphically how the efficient outcome can be guaranteed if type-0 voters vote first. If the type-0 voters vote for l in the first

[5] The election format is defined so that number of periods in which voting occurs is determined endogenously; we show later that this will play an important role in equilibrium selection.

period, the signal vector is fully revealed at the end of the first period. If the election is determined after the first period, the outcome must be efficient. If not, the remaining voters are sufficient in number to swing the election in the direction of the efficient choice.

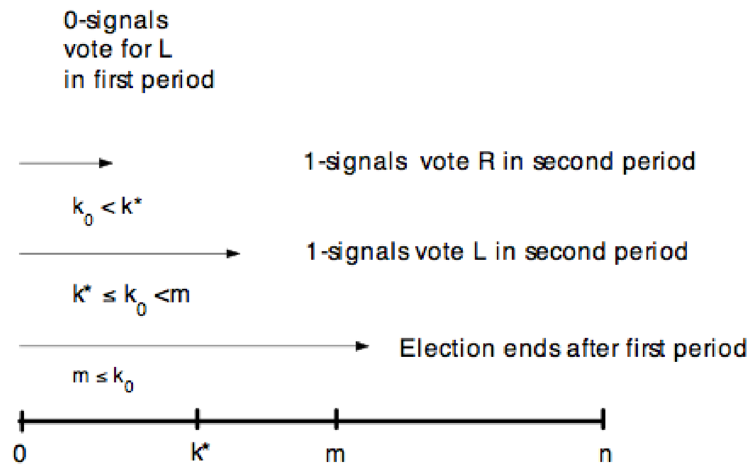

Figure 3.3: Sequential voting achieves efficiency

Notice that the order of voting is critical for efficient aggregation of information. As Figure 3.4 shows, if the type-1 voters vote for r in the first period, they could produce an inefficient outcome. This problem would occur whenever $n - m + 1 \leq k_1 < k_1^*$ or equivalently whenever $k_0^* \leq k_0 < m$.

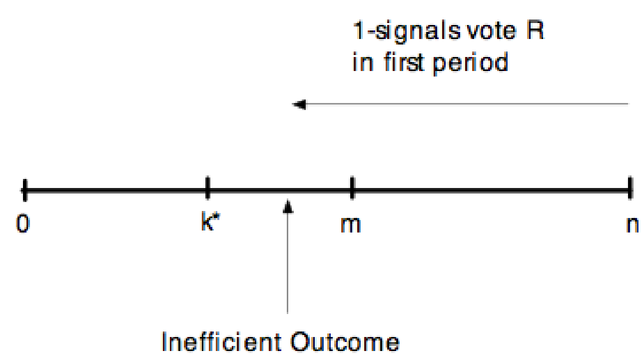

Figure 3.4: The order of voting matters

The following lemma shows that an efficient order of voting always exists. Who votes first depends on the information structure and on the voting rule. The key insight of the proof is that either $k_0^* \leq m$ or $k_1^* = n - k_0^* + 1 < n - m + 1$ must be true, effectively determining the optimal order of voting.

Lemma 1 *The following strategy achieves efficiency in a two-period common value election with binary signals:*

- *If $k_0^*(n) < m(n)$, then type-0 voters vote for l in the first period .*
- *If $k_0^*(n) > m(n)$, then type-1 voters vote for r in the first period.*

In both cases, the voters that didn't vote in the first period learn s perfectly and can guarantee the efficient outcome by voting in the second period. If $k_0^* = m$, the order of voting does not matter.

Proof. Let us consider what happens if the proposed strategy is followed when $k_0^* < m$. We observe that k_0 is revealed after the first period. If $k_0 \geq m$ then l has won the election and the outcome is efficient. If $k_0 < m$ then the election is not determined yet. If $m > k_0 \geq k_0^*$, type-1 voters select l in the second period and the election is determined unanimously in its favor. If $k_0 < k_0^*$, the type-1 voters choose r in the second period and win the election with $k_1 = n - k_0 > n - m$ votes. In either case, the efficient outcome is obtained.

Next, we note that $k_1^* = n - k_0^* + 1$,a relation that follows by definition. Hence, when $k_0^* > m$, $k_1^* < n - m + 1$ where the RHS is the minimum number of votes r must receive to win the election. Thereafter, an argument similar to the $k_0^* < m$ case applies. When $k_0^* = m$, we know that simultaneous voting achieves efficiency. ■

The lemma gives us an efficient order for every q-rule. This is in stark contrast to the result in [24] which highlights the inferiority of the unanimity rule in simultaneous elections.

An important insight regarding the order proposed above follows when we focus our attention on the simple majority rule. The symmetric voting rule allows us to directly relate the order of voting to the quality of information. If $k_0^* < \lceil 0.5n \rceil$, the efficient order recommends type-0 voters to vote first. In this scenario it follows that $k_1^* = n - k_0^* + 1 > k_0^*$. The key to efficiency is to convey the maximum amount of information per vote. In this scenario, type-0 voters can reveal all the relevant information using a smaller number of votes. With uninformative priors ($\pi_L = 0.5$), the condition $k_0^* < \lceil 0.5n \rceil$ reduces to $p_R > p_L \geq \frac{1}{2}$. When $p_R > p_L \geq \frac{1}{2}$, a type-0 voter is better informed (has a more accurate posterior belief) than a type-1 voter. For a simple majority voting rule, the efficient order therefore coincides with the more informed type voting first.

We have yet to establish that voters will follow the efficient ordering prescribed in Lemma 1. Before we do so, a few words regarding our equilibrium concept are in order. We will focus our attention on pure strategy, symmetric, Bayesian Nash equilibria (PSBNE) in this paper. Furthermore, while we will continue to study equilibria where voting is costless, we will limit our attention to those that are vanishing cost proof, a criterion defined as follows:

Definition 2 *An equilibrium σ_E is vanishing cost proof if for every sequence of costs $\{c_i\}$ converging to zero there exists a sequence of equilibria $\sigma_E(i)$ that converges to σ_E.*

We are now ready to state our result.

Proposition 2 *There exists a PSBNE that is efficient and vanishing cost proof. Furthermore, this equilibrium is insensitive to the choice of off-equilibrium beliefs.*

Proof. Suppose without loss of generality that $k_0^*(n) < m(n)$. First we note that no voter has a profitable deviation from the order prescribed in Lemma 1 when voting is costless. A type-0 voter prefers voting for l to abstaining in the first period since there is a positive probability that the number of type-0 voters is exactly $k_0^*(n)$. In a common-value setting, no voter has an incentive to deceive the electorate about his type. A type-0 voter cannot profit from voting for r in the first period (an off-equilibrium action) regardless of how off-equilibrium beliefs are chosen. Similarly, a type-1 voter prefers waiting to voting for l in the first period since there is a positive probability that $k_0 = k_0^*(n)$ and depending on off-equilibrium beliefs can at best be indifferent between voting for r in the first period and waiting until the second. Voters who wait until the second period do not benefit from deviating from the prescribed order.

When voting is costly, voters have an incentive to free ride. But for a small enough cost of voting, the benefit a type-0 voter receives from abstaining is surpassed by the expected loss from communicating the wrong message. Once k_0 is revealed, type-1 voters will mix between voting and abstaining in a proportion that makes each voter indifferent. Voting spills over into additional periods until the outcome is determined. It is easy to show that the probability of voting should increase as the cost of voting goes to zero. We have implicitly assumed here that $T = 1$ and therefore if nobody votes in any period, the election ends.

The fact that type-1 voters mix means that the election may not be determined after the second period. We can show that this does not give the type-0 voters adequate reason to deviate from voting in period 1 when the cost of voting is small. Consider the case where an off-equilibrium vote for l in period 2 is believed to come from a type-0 voter; this off-equilibrium belief maximizes type-0′s incentives to wait. Now let the likelihood that a type-0 voter is pivotal be denoted Δ and the probability that the election is not determined in the second period be denoted $\Phi(c)$. A type-0 voter prefers not to wait in period 1 if

$$(1-\Delta) + \Delta\Phi(c)(1-c) > 1-c$$
$$\Rightarrow 1-\Phi(c) < c\left(\tfrac{1}{\Delta} - \Phi(c)\right)$$

where the RHS represents the utility from voting in the first period. Since

$\lim_{c \to 0} \Phi(c) = 0$, taking limits of both sides shows that the inequality cannot hold for small enough c. This completes the proof. ■

We note that the election format proposed does not prevent voters from coordinating on a pre-determined order of voting based on voter identity. We can refine away any equilibria that arise in sequential voting environments (including those that result in bandwagons as in) on the grounds that they are not symmetric. On a more substantive note, these equilibria are also very sensitive to the choice of off-equilibrium beliefs (see [1]), a characteristic that sets them apart from the efficient equilibrium discussed above.

Another property of the efficient equilibrium worth highlighting is vanishing-cost proofness, a refinement that no simultaneous voting equilibria can survive in our proposed election format. We see this in the proof of the following result.

Proposition 3 *All vanishing-cost proof PSBNE are efficient.*

Proof. First we observe that no simultaneous voting equilibrium can be vanishing cost proof. In any equilibrium of a simultaneous election, the voter influences the outcome only when she is pivotal. For any positive cost of voting, the rational voter therefore prefers to wait until the next period before voting. Since the proposed election format does not have an exogenously specified final period, the option to wait is always available. It then follows that there is no vanishing-cost proof equilibrium in which all voters vote simultaneously. In fact, it is easy to see that if more than one voter is voting in the last period of any vanishing-cost proof equilibrium, the signal vector has already been revealed.

In a binary signal model, we are then left with symmetric, pure strategy profiles in which type-0's and type-1's vote in separate periods. A feature of an inefficient order is that there is some type of voter who can determine the election in favor of the wrong candidate with positive probability. Furthermore, in a binary model, such voter types do not benefit from signalling their type (as in Figure 3.4) and therefore prefer to abstain if voting is costly. This completes the proof. ■

A limitation of the above result is that it only holds in the binary signal model. Proposition 2, however, highlights another important property of the efficient equilibria - robustness to the choice of off-equilibrium beliefs. It is this feature that will distinguish the efficient equilibria in general finite signal environments.

3.4 Finite Signals

The existence of an efficient equilibrium can be generalized to more than two signals if we permit voting to occur over a larger number of periods. Assume that each voter i receives a noisy private signal $s_i \in \{0, 1, ..., d\}$. The signals satisfy the monotone likelihood ratio property:

Assumption 1 *For every s,s', if $s > \acute{s}$, then* $\frac{P(\acute{s}|\omega=L)}{P(s|\omega=L)} \geq \frac{P(\acute{s}|\omega=R)}{P(s|\omega=R)}$

Let $\beta(k_0, k_1, ..., k_d, n)$ be the likelihood-ratio with $\sum_{i=0}^{d} k_i = n$. We can generalize the definition of $k_0^*(n)$ above to the unique $\widetilde{k}_0 \in [1, n]$ such that:

$$\beta(\widetilde{k}_0, 0, 0, ..., n - \widetilde{k}_0, n) \geq 1 > \beta(\widetilde{k}_0 - 1, 0, 0, ..., n - \widetilde{k}_0 + 1, n)$$

It is easy to see that this definition also defines k_d^* as $k_d^* = n - k_0^* + 1$. As before if $\beta(0, 0, 0, ..., n, n) > 1$, $k_0^*(n) = 0$ and if $\beta(n, 0, 0, ..., 0, n) < 1$, $k_0^*(n) = n + 1$. In order to preserve the interpretation of $k_0^*(n)$ as being the minimum number of 0-signals such that l is the efficient choice, we define $k_0^*(n)$ in terms of signal vectors of the form $[k_0, 0, 0, \ldots, k_d]$. It follows from the MLRP that if $\beta(\widetilde{k}_0, 0, 0, ..., n - \widetilde{k}_0, n) \geq 1$, then $\beta(\widetilde{k}_0, k_1, ..., k_d, n) \geq 1$ for all possible signal vectors of the form $s = \left[\widetilde{k}_0, k_1, \ldots, k_d\right]$.

We then have:

Proposition 4 *There exists an efficient order of voting if the noisy signals satisfy MLRP and voting can be conducted over any number of periods. This efficient order can be sustained in a vanishing-cost proof PSBNE.*

Proof. The proof for finite signals follows recursively. The argument for binary signals is given in Lemma 1. We now begin with $d+1$ signals and an initial prior π_L. The MLRP provides a natural ranking of signals. We use it to define $k_0^*(n)$ as being the minimum number of 0-signals that make l the efficient choice, regardless of what the remaining $n - k_0^*(n)$ signals are. Similarly, we define $k_d^*(n)$ and note that $k_0^*(n) + k_d^*(n) = n + 1$ must hold. It therefore follows that either $k_0^*(n) \leq m(n)$ or $k_d^*(n) \leq n - m(n) + 1$ is true. If $k_0^*(n) \leq m(n)$, the type-0 voters vote for l in the first period and if $k_d^*(n) \leq n - m(n) + 1$, type-$d$ voters for r in the first period. If the winner is determined at the end of the first period, the outcome must be efficient. If the election is still open, the aggregation of signals proceeds as follows. Without loss of generality, let us assume that the type-0's voted first and there were $k_0 < m(n)$ of them. At this point, $n - k_0$ voters remain, candidate l requires $m(n) - k_0$ votes to win and d types of voters are yet to vote. If $k_0 \geq k_0^*(n)$, all remaining voters vote for l in the second period and the election is determined efficiently. If not, we compute a conditional $k_1^*(n - k_0; k_0)$ which represents the minimum number of 1-signals which when combined with k_0 signals of type-0 identifies l as the efficient choice, regardless of the remaining $n - k_0 - k_1^*(n - k_0; k_0)$ signals. That is, $k_1^*(n - k_0; k_0)$ is the unique $\widetilde{k}_1 \in [1, n - k_0]$ such that:

$$\beta(k_0, \widetilde{k}_1, 0, ..., n - k_0 - \widetilde{k}_1, n) \geq 1 > \beta(k_0, \widetilde{k}_1, 0, ..., n - k_0 - \widetilde{k}_1, n)$$

Since $k_0 < k_0^*(n)$, we know that $\beta(k_0, 0, 0, ..., n - k_0, n) < 1$. If $\beta(k_0, n - k_0, 0, ..., 0, n) < 1$, then we define $k_1^*(n - k_0; k_0) = n - k_0 + 1$; in this case information aggregation is complete and r is the efficient choice. The definition of $k_1^*(n - k_0; k_0)$ implicitly defines a $k_d^*(n-k_0;k_0)$ and as before the relationship $k_1^*(n-k_0;k_0)+k_1^*(n-k_0;k_0) = n-k_0+1$ must hold. This implies that either $k_1^*(n-k_0;k_0) < m(n)-k_0$ or $k_d^*(n-k_0;k_0) \leq n-m(n)+1$ must be true. The aggregation problem now looks identical to the first iteration with $n-k_0$, $m(n)-k_0$, $k_1^*(n-k_0;k_0)$ and $k_d^*(n-k_0;k_0)$ replacing n, $m(n)$, $k_0^*(n)$ and $k_d^*(n)$ respectively. The next voter type to vote and the candidate for whom they should vote can now be specified as before. This recursive process is repeated until the election is determined or all $d+1$ signals have been aggregated. In any case, no matter when the election ends, the outcome is guaranteed to be efficient. The argument demonstrating that the efficient order can be sustained in a vanishing-cost PSBNE follows that in Proposition 2 and is omitted. ■

The above proposition proves the existence of an efficient order when private signals can take one of $d+1$ values and time can be partitioned into $d+1$ periods. The key insight in this argument is that at each stage we can compute thresholds k_i^* for each of the two extreme signals and determine which one of them should vote next. This allows us to maximize the information communicated per vote cast. The efficient order we present is however not unique. It is not difficult to see that for some signal distributions, it is not necessary to follow the ranking induced by the MLRP. Relatively less informative signals may vote earlier than more informative signals without jeopardizing social welfare. In fact, it is also possible to allow more than one type to vote in the same period while preserving efficiency.

3.5 Inefficient Equilibria

In the finite signal model, it is possible to construct vanishing cost proof PSBNE that are inefficient. The following example shows one such equilibrium with $d+1=3$ signals.

Consider an electorate with $n=9$ voters and an election with a simple majority rule. Voters receive one of three signals $s_i \in \{0, 1, 2\}$. Both states are ex-ante equally likely. We use the notation p_{JK} to denote the probability $P(s_i = K|\omega = J)$. Then $p_{L0} = 0.51$, $p_{L1} = 0.246$, $p_{L2} = 0.244$, $p_{R0} = 0.003$, $p_{R1} = 0.006$, $p_{R2} = 0.991$. In state $\omega = R$, an overwhelming majority of voters receive a signal 2, whereas the electorate has a more diffuse set of types when the state is L . Conditional on signals 0 or 1, the state is more likely to be L.Following the notation of the previous section, $k_0^*(n) = 2$. i.e, $\beta(k_0 = 1, k_1 = 0, k_2 = 8, n = 9) < 1 \leq \beta(2, 0, 7, 9)$. Recall that this notation implies that if 2 or more signals were of type-0, it would efficient for the electorate to choose alternative l. Similarly $k_1^* = 3$, i.e. $\beta(0, 2, 7, 9) < 1 \leq \beta(0, 3, 6, 9)$.

We now propose an inefficient equilibrium in two periods. In period 1, voters of type-0 and type-1 vote in favor of l. If the number of voters who vote in the first period is greater or equal to 3, the type-2 voters vote for l in the second period. Otherwise, they vote for R. If this proposed strategy profile were an equilibrium, an inefficiency would arise whenever the electorate had 2 type-0 voters and the remaining voters were type-2. This follows from the fact that $k_0^* = 2$. We now proceed to confirm that what we have proposed is indeed an equilibrium.

Consider first the type-2 voter's problem. Suppose he observes $k \geq 3$ votes in the first period. It follows from $k_1^* = 3$ and the fact that 0-signals are stronger signals in favor of l that the type-2 voter should vote for L. But what if $k = 2$? One of three scenarios is possible - the signal profile of the electorate can be (i) $(k_0 = 2, k_1 = 0, k_2 = 7)$ or (ii) $(1, 1, 7)$ or (iii) $(0, 2, 7)$. In cases (i) and (iii) the type-2 voter votes l and r respectively. In case (ii), she prefers r because $\beta(1, 1, 7, 9) < 1$. Moreover, it can easily be shown that $P((2, 0, 7)|s_i = 2) < P((1, 1, 7)|s_i = 2) + P((0, 2, 7)|s_i = 2)$. Therefore the type-2 voter does not vote for l unless $k \geq 3$ and adheres to the proposed equilibrium.

Next consider the type-1 voter. Her vote only matters when she is pivotal, i.e. when the number of type-0's and type-1's is exactly equal to 3. Since $\beta(0, 3, 6, 9)$

≥ 1 and the likelihood ratio in the pivotal scenario must be greater than or equal to $\beta(0,3,6,9)$, the type-1 voter has no incentive to deviate from the proposed equilibrium. If the type-1 voter adheres to the equilibrium, so does the type-0 voter whose signal in favor of l is stronger. The arguments from Proposition 2 can be applied to argue that this equilibrium is vanishing cost proof.

To sustain this equilibrium, we must also specify off-equilibrium beliefs. Suppose first that any vote for r in the first period is associated with a type-2 voter who could not wait until the next period for some exogenous reason. In this case, no type-0 or type-1 voter will deviate and the inefficient equilibrium can be sustained. The inefficiency arises because types 0 and 1 have no way of distinguishing themselves from each other and they strictly prefer not to mimic type 2 voters. The off equilibrium path however can be used as a mechanism for types 0 and 1 to separate. If any vote in favor of r in the first period is associated with a type-1, for instance, the above equilibrium breaks down. In this case, the type-1 voter prefers to identify himself by voting insincerely in favor of r and thereby potentially averting an inefficient outcome.

We observe that while vanishing cost proof PSBNE can be inefficient, all efficient PSBNE are robust to the choice of off-equilibrium beliefs. Although this is not a standard refinement used in the signaling literature, we consider it worth highlighting. Typical refinements such as the intuitive criterion or divinity do not eliminate inefficiency in this framework. An alternative approach would be to employ Farrell's [21] notion of "neologism proofness". The inefficient equilibrium constructed above exists because voters are limited to a language that is not rich enough to communicate their types. If the election is augmented so that voters have the option to communicate cheap messages, inefficient PSBNE will cease to exist. We do not pursue this alternative so as to distinguish our results from the logic of Coughlan's work.

3.6 Partisans

The next question one might ask is how does the proposed mechanism perform when voters have heterogenous preferences. In the context of our model, the most natural way to introduce interdependent preferences would be to give each voter a multi-dimensional type comprised of a preference class and a noisy private signal about the underlying state. A voter's preference class is most readily thought of (and indeed has been defined in this manner in the literature) as the relative disutility she suffers from each of the two possible errors that are likely to be made in the collective decision. Differences in relative disutilities between preference classes translate into differences in threshold values of $k_0^*(n)$. A voter who suffers much greater disutility when r is chosen in state L than when l is chosen in state R will require much stronger evidence in favor of R before voting for alternative r than a voter who is unbiased. Unfortunately, such modeling has proved intractable and general results have been difficult to come by. However, we are able to provide some special cases where efficiency obtains in environments where voters don't share the same objectives.

Suppose we define l-partisans as voters who prefer policy l in every state of the world. A similar definition can be used for r-partisans. This is a special case of the first modeling of preference heterogeneity - one in which the l-partisans' disutility when choosing r in state L is infinite relative to the disutility of choosing l in state R. Suppose we partition the electorate into l-partisans, r-partisans and unbiased voters. If the number of l-partisans and r-partisans is known, the efficient outcome can be achieved. It is worth taking note at this point of how efficiency was defined earlier in the paper. If the number of l-partisans is large enough to sway the election in favor of l in every state of the world, that outcome is not a violation of efficiency. The reasoning behind the efficient outcome when the composition of the electorate is common knowledge is the following - the partisans would like to mimic the most informed voters. Since the non-partisan voters know exactly how many partisans there are, they can always glean the number of informative votes from the intermediate tally and ensure efficient aggregation. A formal argument is provided below for the binary signal model for the sake of simplicity. The result is easily extended to the general case.

Proposition 5 *If voter preferences are common knowledge, there exists a vanishing cost proof PSBNE that achieves the efficient outcome. This equilibrium is not robust to the choice of off-equilibrium beliefs.*

Proof. Consider an electorate with n voters, k_l partisans supporting l and k_r partisans supporting r. We define $k_0^*(n-k_l-k_r)$ as the number of 0 signals that would identify l as the efficient choice for the unbiased voters. Since $k_0^*(n-k_l-k_r)+k_1^*(n-k_l-k_r) = n-k_l-k_r+1$, either $k_0^*(n-k_l-k_r)+k_l < m(n)$ or $k_1^*(n-k_l-k_r)+k_r \leq n-m(n)+1$, must be true. Consider the scenario in which $k_0^*(n-k_l-k_r)+k_l < m(n)$; in this case type-0 voters vote for l in the first period and type-1 voters follow in the next period. We now require that in equilibrium, type-1 voters commit to subtracting k_l from k_0, the number of votes cast in the first period before they decide how to vote. We also assume that off the equilibrium path, type-1 voters believe that any voter who votes for r in the first period must be an r-partisan. In light of this strategy, the l-partisans strictly prefer to vote for l in the first period. The type-1 voters do not benefit by deviating from their commitment in a common value setting. The r-partisans are indifferent about the time at which they vote as they are not relevant to the aggregation process. No partisan benefits from voting insincerely. The equilibrium does not survive if voters assign positive probability to off-equilibrium path votes being informative. The efficient order ensures that information is correctly aggregated whenever partisans are not large enough in number to swing the election in their favor. This happens whenever $k_l > m(n)$ or $k_r > m(n)$. ■

A second scenario, one that has been used in the literature, but is not particularly convincing to the authors is the case where partisan voters act like automatons. That is, they vote for their preferred alternative, but do not act strategically in choosing the time to vote; in other words, they do not actively attempt to deceive unbiased voters. In this case, unbiased voters can coordinate among themselves and vote in periods where partisan voters are not present.

3.7 Uncertainty in Population Size

In large elections, the precise number of voters is often unknown. In this section we examine an environment where the number of voters n is a random variable and all voters share common beliefs about the size of the electorate. As before, alternative l is chosen if it receives $m(n)$ or more votes. We will restrict our results to simple majority rule, where $m(n) = \lceil \frac{n}{2} \rceil$, but our results hold for any q-rule. Uncertainty in n translates to uncertainty in $m(n)$. We will assume $s_i \in \{0, 1\}$ and that $k_0^*(n)$ is the minimum number of 0-signals necessary for l to be the efficient choice, given that there are n voters. $k_1^*(n)$ is similarly defined.

To simplify our calculations, we will assume that $n = 2\lambda + 1$, where λ is a random variable that comes from a probability distribution $f(\lambda)$, with $\lambda \in [1, \infty) \bigcap \mathbb{Z}$. This setting has the convenient property that $m(n) = \lceil \frac{n}{2} \rceil = \lambda + 1$ for all the values of the random variable λ.

We now modify our model - time is continuous, but votes are declared publicly at discrete intervals. That is, a voter can vote at any time $t \in [0, \infty)$, but intermediate results are broadcast only at $t = 1, 2, 3, ...$ and so on. A voter can choose in which period to vote, but cannot choose the exact time - the exact time of her vote is determined randomly. No two voters vote at the same time. A voter can observe the results up until the end of the last period as well as the number of voters ahead of her in the current period. For example, a voter that chooses to vote in period 5 will be assigned a time to vote $t \in [5, 6)$ at random. She will observe the number of voters who have voted before her in the current period and the results announced at $t = 5$. The voter can always choose not to vote or postpone her vote for later - either at some random time within the same period or to some future period of her choosing. Our main result will be make use of the following lemmas:

Lemma 2 *If $\frac{1}{2} \leq p_L \leq p_R$ then for any n, $k_0^*(n+2) - k_0^*(n) \leq 1$. Similarly, if $\frac{1}{2} \leq p_R \leq p_L$ then for any n, $k_1^*(n+2) - k_1^*(n) \leq 1$.*

Proof. Let $\frac{1}{2} \leq p_L \leq p_R$, so $\frac{p_L}{1-p_R} \geq \frac{p_R}{1-p_L} \geq 1$. For fixed n, let $k_0^*(n+2) = k_0^*(n) + \mu$ and $k_0^*(n) = k_0^*$. From the definition of $k_0^*(n)$, we have

$$\frac{p_L^{k_0^*-1}(1-p_L)^{n-k_0^*+1}}{p_R^{n-k_0^*+1}(1-p_R)^{k_0^*-1}} < \frac{1-\pi_L}{\pi_L} \leq \frac{p_L^{k_0^*}(1-p_L)^{n-k_0^*}}{p_R^{n-k_0^*}(1-p_R)^{k_0^*}} \quad (1)$$

From the definition of $k_0^*(n+2)$, we have

$$\frac{p_L^{k_0^*+\mu-1}(1-p_L)^{n-k_0^*+3-\mu}}{p_R^{n-k_0^*+3-\mu}(1-p_R)^{k_0^*+\mu-1}} < \frac{1-\pi_L}{\pi_L} \leq \frac{p_L^{k_0^*+\mu}(1-p_L)^{n-k_0^*+2-\mu}}{p_R^{n-k_0^*+2-\mu}(1-p_R)^{k_0^*+\mu}}$$

$$\Rightarrow \frac{p_L^{\mu-1}(1-p_L)^{3-\mu}}{p_R^{3-\mu}(1-p_R)^{\mu-1}} \frac{p_L^{k_0^*}(1-p_L)^{n-k_0^*}}{p_R^{n-k_0^*}(1-p_R)^{k_0^*}} < \frac{1-\pi_L}{\pi_L} \leq \frac{p_L^{k_0^*+\mu}(1-p_L)^{n-k_0^*+2-\mu}}{p_R^{n-k_0^*+2-\mu}(1-p_R)^{k_0^*+\mu}} \quad (2)$$

If $\mu \geq 2$ then $(\frac{p_L}{1-p_R})^{\mu-1} \geq (\frac{p_R}{1-p_L})^{3-\mu}$, so $\frac{p_L^{\mu-1}(1-p_L)^{3-\mu}}{p_R^{3-\mu}(1-p_R)^{\mu-1}} \geq 1$ and (1), (2) give us

$$\frac{1-\pi_L}{\pi_L} \leq \frac{p_L^{k_0^*}(1-p_L)^{n-k_0^*}}{p_R^{n-k_0^*}(1-p_R)^{k_0^*}} \leq \frac{p_L^{\mu-1}(1-p_L)^{3-\mu}}{p_R^{3-\mu}(1-p_R)^{\mu-1}} \frac{p_L^{k_0^*}(1-p_L)^{n-k_0^*}}{p_R^{n-k_0^*}(1-p_R)^{k_0^*}} < \frac{1-\pi_L}{\pi_L}$$

which is a contradiction. Hence, $\mu \leq 1$ and $k_0^*(n+2) - k_0^*(n) \leq 1$ for any n. The proof for $\frac{1}{2} \leq p_R \leq p_L$ is similar. ■

Lemma 3 *If $\frac{1}{2} \leq p_L \leq p_R$, then for all integers n of the form $n = 2\lambda + 1$ with $\lambda \in [1, \infty) \bigcap \mathbb{Z}$ we have $m(n) \geq k_0^*(n)$. Similarly, if $\frac{1}{2} \leq p_R \leq p_L$, then for all integers n of the form $n = 2\lambda + 1$ with $\lambda \in [1, \infty) \bigcap \mathbb{Z}$ we have $m(n) \geq k_1^*(n)$.*

Proof. Let $\frac{1}{2} \leq p_L \leq p_R$. Then $m(n) = m(2\lambda + 1) = \lambda + 1$ and, from Lemma 2, $k_0^*(n) = k_0^*(2\lambda + 1) \leq \lambda + 1$, so $k_0^*(n) \leq m(n)$. The proof is similar for $\frac{1}{2} \leq p_R \leq p_L$. ■

Proposition 6 *There exists an equilibrium strategy that achieves the efficient outcome in an election with an unknown number of voters if votes may only be cast one at a time and intermediate results are announced at discrete intervals.*

Proof. Without loss of generality, let $\frac{1}{2} \leq p_L \leq p_R$. Assume that all 0-signal voters vote in the first period for l, and that there are k_0 of them. (If $\frac{1}{2} \leq p_R \leq p_L$ then 1-signal voters will vote in the first period for r.) If $k_0 \geq m(n)$ then the election is over and the efficient outcome has been achieved, since $m(n) \geq k_0^*(n)$. (Lemma 3) Suppose $k_0 < m(n)$. After the first period, all the remaining voters observe k_0, but they don't know n. Let $\bar{n}$ be defined as the smallest odd $\bar{n} = 2\bar{\lambda} + 1$ with $k_0^*(\bar{n}) > k_0$. From Lemma 2, we know that $k_0 < k_0^*(2\bar{\lambda} + 1) \leq \bar{\lambda} + 1$. Notice that $\bar{n}$ is known by all the remaining voters after the first period. Now, if the remaining $n - k_0$ voters knew n, then they would like to all vote for r if $n \geq \bar{n}$ and for l if $n < \bar{n}$. In our equilibrium we will achieve the same result without requiring from the voters to know n. The idea is as follows: the first $\bar{\lambda} - k_0 \geq 0$ remaining voters vote for l and the ones remaining after that (if any) vote for r. Then the efficient result will always be achieved, regardless of the value of n, as r will be the outcome of the election if and only if $n \geq 2\bar{\lambda}+1 = \bar{n}$. This outcome can always be implemented as voters can always see how many others voted before them and vote or postpone voting accordingly. ■

3.8 Multidimensional Signals

Efficient decision-making typically involves evaluating choices according to a variety of factors. In this section, we consider environments where voters receive multidimensional signals. If these signals can be collapsed into a one dimensional sufficient statistic that satisfies the MLRP, proposition 4 establishes the efficiency of our proposed election format. While such a mapping is sufficient it is not necessary for efficient aggregation. The following example demonstrates how two-dimensional signals can be aggregated as long as MLRP is satisfied on each dimension.

Suppose that each of the n voters receives a multidimensional signal $s_i = (s_i^1, s_i^2) \in \{0,1\}^2$. Consider an information aggregation rule $E : \{s_1, s_2, ..., s_n\} \longrightarrow \{l, r\}$ as follows, for given $S_1(n)$, $S_2(n)$:

$$E(s_1, s_2, ..., s_n) = \begin{cases} l, & \text{if } \sum_{i=1}^n s_i^1 < S_1 \text{ and } s^2 = \sum_{i=1}^n s_i^2 < S_2 \\ r, & \text{otherwise} \end{cases}$$

This aggregation rule can interpreted as one in which the alternative r is preferred as long as it meets a minimum threshold on at least one of the two dimensions. In the following figure, the aggregation rule is represented by the solid line that partitions the signal space into the efficient regions for l and r.

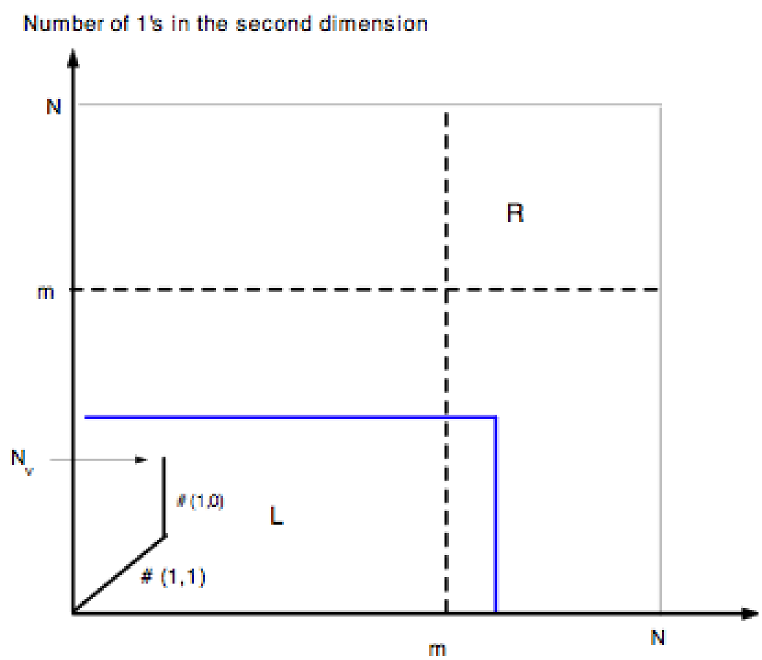

Figure 3.5: Efficient information aggregation with multidimensional signals

Notice that this aggregation rule cannot be implemented using a one-dimensional sufficient statistic. To see this, let $s_i = (0, 1)$ and $s_j = (1, 0)$. It is impossible to rank s_i and s_j; which of the two signals is stronger evidence in favor of alternative l, for instance, depends on the private signals of the remaining voters. For example, an s_i is stronger evidence in favor of r when the remaining signals, when aggregated, produce a point close to the horizontal portion of the solid line, whereas an s_j is stronger evidence on the vertical boundary delineated by the solid line. Nevertheless, we show below that our election format can achieve the efficient outcome.

Proposition 7 *The proposed election format aggregates the rule E efficiently.*

Proof. Assume that $S_2 < m$ as in Figure 3.5. The proof for $S_2 \geq m$ will follow by appropriately adjusting the following argument. We will allow the voters with signal $s_i = (1, 1)$ to vote in the first period for r. Suppose there are $k_{(1,1)}$ such voters.

• If $k_{(1,1)} \geq m$, then the election is determined, but since $k_{(1,1)} \geq m > S_2$ the efficient outcome r has been achieved.

• If $k_{(1,1)} < m$, then we will allow voters with signal $s_i = (1, 0)$ to vote in the second period for r. Suppose there are $k_{(1,0)}$ such voters. If $k_{(1,1)} + k_{(1,0)} > S_2$, then the efficient outcome is revealed to be r and voters in the next period can vote to ensure that r is achieved. If, on the other hand, $k_{(1,1)} + k_{(1,0)} \leq S_2$ then the remaining voters are of type $(0, 0)$ and $(0, 1)$ and the election has not yet been determined, since $k_{(1,1)} + k_{(1,0)} \leq S_2 < m$. But this is just equivalent to the single dimensional problem presented in Lemma 1. ■

3.9 Conclusion

We find that allowing voters to choose when to cast their vote can lead to substantial improvements in the aggregation of information in common value settings. Our results are of particular relevance for small electorates where simultaneous voting is known to be inefficient. In future work, we would like to develop a better understanding of the performance of our electoral framework when voters have heterogenous preferences. While recent papers by Austen-Smith and Feddersen [6] and also by Gerardi and Yariv suggest that efficiency may be difficult to implement in interdependent value settings, a complete characterization of the optimal mechanism for information aggregation in such environments is yet to be found.

Bibliography

[1] Ali, Nageeb, and Navin Kartik. 2007. “Social Learning in Elections”, UCSD, Working Paper.

[2] Aoyagi, Masaki (2007), ”Efficient Collusion in Repeated Auctions with Communication”, *Journal of Economic Theory*, 134, 61-92.

[3] Athey, Susan and David Miller (2003), ”Efficiency in Repeated Trade with Hidden Valuations”, Harvard University Working Paper.

[4] Athey, Susan and Ilya Segal (2007), ”An Efficient Dynamic Mechanism”, Harvard University Working Paper.

[5] Austen-Smith, David, and Jeffrey Banks. 1996. “Information Aggregation, Rationality, and the Condorcet Jury Theorem”, *American Political Science Review,* 90(1):34-45.

[6] Austen-Smith, David, and Timothy Federssen. 2006. “Deliberation, Preference Uncertainty and Voting Rules”, *American Political Science Review,* 100(2):209-217.

[7] Battaglini, Marco. 2005. “Sequential Voting with Abstention”, *Games and Economic Behavior*, 51, 445-463.

[8] Battaglini, Marco (2005), ”Long Term Contracting with Markovian Consumers”, *American Economic Review*, 95(3), 637:658.

[9] Bergemann, Dirk and Juuso Valimaki (2003), "Dynamic Common Agency", *Journal of Economic Theory*, 111, 23-48.

[10] Bergemann, Dirk and Juuso Valimaki (2007), "Dynamic Marginal Contribution Mechanism", Cowles Foundation working paper.

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

[12] Callander, Steven. 2007. "Political Motivations" *Review of Economic Studies*. Forthcoming.

[13] Cavallo, Ruggiero, David Parkes and Satinder Singh (2006), "Optimal Coordinated Planning Among Self-Interested Agents with Private States", *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence*, Cambridge.

[14] Coate, Stephen, and Michael Conlin. 2004. “A Group Rule-Utilitarian Approach to Voter Turnout: Theory and Evidence” *American Economic Review,* 94(5): 1476-1504.

[15] Coughlan, Peter J. 2000. “In Defense of Unanimous Jury Verdicts: Mistrials, Communication, and Strategic Voting”, *American Political Science Review*, 94(2): 375-393.

[16] Cremer, Jacques and Richard Mclean (1988), "Full Extraction of Surplus in Bayesian and Dominant Strategy Auctions", *Econometrica*, 56, 1247-1257.

[17] d'Aspremont, Claude and Louis-Andre Gerard-Varet (1979), "Incentives and incomplete information", *Journal of Public Economics*, 11, 25-45.

[18] Dekel, Ed and Michele Piccione. 2000. “Sequential Voting Procedures in Symmetric Binary Elections”, *Journal of Political Economy,* 108: 34-55.

[19] Doepke, Matthias and Robert Townsend (2006), "Dynamic Mechanism Design with Hidden Income and Hidden Actions", *Journal of Economic Theory*, 126, 235-285.

[20] Farquharson, Robin. 1969. *Theory of Voting*. New Haven: Yale University Press.

[21] Farrell, Joseph. 1993. "Meaning and Credibility in Cheap Talk Games", *Games and Economic Behavior*, 5(4): 514-531.

[22] Feddersen, Timothy, and Alvaro Sandroni. 2006. "A Theory of Participation in Elections" *American Economic Review*, 96(4): 1271-1282.

[23] Feddersen, Timothy, and Wolfgang Pesendorfer. 1996. "The Swing Voter's Curse", *American Economic Review*, 86: 408-424.

[24] Feddersen, Timothy, and Wolfgang Pesendorfer. 1997. "Voting Behavior and Information Aggregation in Elections with Private Information", *Econometrica*, 65(5): 1029-1058.

[25] Feddersen, Timothy, and Wolfgang Pesendorfer. 1998. "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts", *American Political Science Review*, 92: 23-35.

[26] Fernandez, Raquel, and Dani Rodrik. 1991. "Resistance to Reform: Status Quo Bias in the Presence of Individual-Specific Uncertainty" *American Economic Review, 81*(5): 1146-1155.

[27] Fey, Mark. 2000. "Informational Cascades and Sequential Voting", University of Rochester, Working Paper.

[28] Fudenberg, Drew, David Levine and Eric Maskin (1995),"The Folk Theorem with Imperfect Public Information," *Econometrica*, 62 (5), 997-1039.

[29] Guarnaschelli, Serena, Richard D. McKelvey, and Thomas R. Palfrey. 2000. “An Experimental Study of Jury Decision Rules” *American Political Science Review*, 94(2): 407-423.

[30] Hansson, Ingemar, and Charles Stuart. 1984. “Voting Competitions with Interested Politicians: Platforms do not Converge to the Preferences of the Median Voter,” *Public Choice*, 44(3): 431-441.

[31] Harsanyi, John. 1980. “Rule Utilitarianism, Rights, Obligations and the Theory of Rational Behavior,” *Theory and Decision*, *12*(2): 115-133.

[32] Kartik, Navin, and R. Preston McAfee. 2007. “Signaling Character in Electoral Competition”, *American Economic Review*, 97(3): 852-870.

[33] Kocherlakota, Narayana (2004), ”Figuring Out the Impact of Hidden Savings on Optimal Unemployment Insurance”, *Review of Economic Dynamics*, 7, 541-554.

[34] Krishna, Vijay (2002), *Auction Theory*, Academic Press, USA.

[35] Mezetti, Claudio (2004), ”Mechanism Design with Interdependent Valuations: Efficiency, ” *Econometrica*, 72(5), 1617.

[36] Mezetti, Claudio (2007), ”Mechanism Design with Interdependent Valuations: Surplus Extraction, ” *Economic Theory*, 31, 473-488.

[37] Neeman, Zvika (2004), ”The relevance of private information in mechanism design”, *Journal of Economic Theory*, 117, 55-77.

[38] Persico, Nicola. 2000. “Committee Design with Endogenous Information”, *Review of Economic Studies*, 71(1):165-191.

[39] Piketty, Thomas. 2000. “Voting as Communicating”, *Review of Economic Studies*, 67(1): 169-191.

[40] Puterman, Martin, (1994), *Markov Decision Processes*, Wiley: Hoboken, NJ.

[41] Rochet, Jean-Charles (1987), "Necessary and Sufficient Condition for Rationalizability in a Quasilinear Context", *Journal of Mathematical Economics*, 16, 191-200.

[42] Samuelson, William, and Richard Zeckhauser. 1988. "Status quo bias in decision making*". Journal of Risk and Uncertainty,* 1(1): 7-59.

[43] Spangenberg, Frits. 2003. "The Freedom to Publish Opinion Poll Results" Foundation for Information.

[44] Williams, Noah, (2006), "On Dynamic Principal-Agent Problems in Continuous Time", Working Paper, Princeton University.

[45] Williams, Noah (2008), "Persistent Private Information", Working Paper, Princeton University.