

Pivotal Players and the Characterization of Influence¹

Nabil I. Al-Najjar

*Department of Managerial Economics and Decision Sciences,
J. L. Kellogg Graduate School of Management, Northwestern University,
2001 Sheridan Road, Evanston, Illinois 60208
al-najjar@nwu.edu*

and

Rann Smorodinsky

*Davidson Faculty of Industrial Engineering and Management, Technion,
Haifa 32000, Israel
rann@ie.technion.ac.il*

Received May 13, 1998; revised October 8, 1999

A player's influence relative to a mechanism and opponents' strategies is the maximum difference his action can make in the expected value of a collective outcome. A player is α -pivotal if his influence exceeds a threshold α . We provide tight bounds on the number of pivotal players and on average influence. These bounds are uniform over all mechanisms and action profiles and are achieved at mechanisms that take the form of a majority rule. We illustrate our analysis with an example of provision of a public good where each individual compares the private cost of his contribution with its influence on the collective outcome. *Journal of Economic Literature* Classification Numbers: D62, D89, H41. © 2000 Academic Press

1. INTRODUCTION

Strategic behavior often hinges on players' beliefs about their influence on a collective outcome. This dependence surfaces in many contexts, ranging from the provision of public goods and allocations in the presence of externalities, voting, implementation, auctions, and repeated games. In these settings, only players who believe they will be pivotal take into account the impact of their actions on a collective outcome; non-pivotal players ignore such strategic considerations and behave myopically.

¹ We thank Greg Greiff, Peter Klibanoff, Ehud Lehrer, Wolfgang Pesendorfer, Eilon Solan, and seminar participants at Tel-Aviv, Technion, Jerusalem, Northwestern, Texas A&M, and the University of Montreal for their comments. Rann Smorodinsky wishes to acknowledge the Technion V.P.R. Fund for financial support.

This paper provides a systematic treatment of the notions of pivotalness and influence in general environments with N players, each with a random signal \tilde{t}_n . Signals, which can be correlated and asymmetrically distributed, may represent privately known types (as in implementation and mechanism design problems) or noisy outcomes of unobserved actions (as in problems with moral hazard). An outcome function F is an arbitrary mapping of the vector of signals into a collective outcome. We define a player's *influence* as the maximum change in the expected value of F caused by a change in his signal; a player is α -*pivotal* if his influence exceeds some threshold α .

Our main results provide tight upper bounds on the number of α -pivotal players and average influence in the population. These bounds are uniform over all mechanisms and profiles, and achieved by mechanisms that take the form of a majority rule. We show that for any general, non-anonymous mechanism and asymmetric distribution there is an anonymous mechanism and a symmetric distribution in which the number of α -pivotal players and average influence increase. This enables us to reduce the general problem to a straightforward calculation of pivot probabilities in simple majority rules.² In particular, we show that the number of α -pivotal players is bounded independently of N , with the bound achieved by applying a simple majority rule to an "oligarchy" of K players.

This analysis of influence and pivotalness is useful in settings where a player's decision hinges on a comparison between the private cost of an action with his impact on a collective outcome. An example illustrating our analysis is asymptotic inefficiency results in public good economies with private information (Rob [10] and Mailath and Postlewaite's [6]; see Section 4). In that context, an individual with private valuation for a public good compares the cost of his contribution with its expected impact on the probability of provision. We provide a simple, transparent proof of a finite-population upper bound on the probability of provision, from which asymptotic inefficiency immediately follows.

The idea that, in an environment with uncertainty, there may be bounds on the number of potentially pivotal players is not new. As indicated earlier, Mailath and Postlewaite's analysis builds on the observation that, in a large economy with independent private valuations "few agents can be pivotal" (p. 363). In the appendix of their paper, they consider a setting with N independent signals and a random variable f which may be correlated with these signals (the outcome and the signals may be continuous); they show that the ratio of agents whose signals are highly

² To put this reduction in perspective, suppose there are $N = 10$ players each having three signals and that the outcome is binary (i.e., $F \in \{0, 1\}$). Then there are 6 majority rules, 2^{66} anonymous rules, and $2^{3^{10}} = 2^{59,049}$ general non-anonymous rules.

correlated with f goes to zero at the rate of $N^{-1/2}$. Fudenberg, Levine and Pesendorfer ([4], Lemma A), in an environment with independent, discrete signals, derive a bound on average influence on the probability of a subset of outcomes that converges to zero at the rate $N^{-1/2}$. They use this to study games with many small players, each taking an unobserved action that might influence the action of a large player. Our results differ from these papers in many respects, as discussed in Section 3.4.

A related literature centers on the work of Green [5] on the foundation of price-taking behavior in a competitive industry with repeated interaction. The usual argument given to justify price taking behavior (namely that there are many firms) is not sufficient because a collusive outcome may be sustained via folk theorem-type constructions, regardless of the number of players. Green suggested that imperfect observability of demand makes it difficult to trigger punishments when a firm deviates from the collusive agreement, and that this problem worsens as the number of firms increases. Sabourian [11] generalized this argument in several directions, showing that in repeated games with noisy anonymous outcomes the only equilibria are those consisting of playing a sequence of stage game Nash equilibrium. In the language of the present paper, collusion in Green/Sabourian-type models cannot be sustained because few players will believe their actions will be pivotal relative to the continuation of future cooperation. In the present paper, we derive tight bounds on general non-anonymous mechanisms, with no counterpart in the arguments of Green or Sabourian. In a companion paper (Al-Najjar and Smorodinsky [2]) we extend Green/Sabourian model along these lines.³

Our notion of influence generalizes that of *pivot probability* in the voting literature (see, for instance, Chamberlain and Rothchild [3], Palfrey and Rosenthal [8], and Myerson [7]). This literature focuses on a common, discrete set of possible signals and anonymous mechanisms. Our analysis shows that the problem of characterizing influence for potentially complex, non-anonymous mechanisms can be reduced to the corresponding problem in a symmetric voting environment in which the outcome is determined according to majority rule.

2. THE MODEL

2.1. Signals and Mechanisms

Consider an environment with N players, each with a random signal \tilde{t}_n taking values in a set T_n (t_n denotes the realized signal). A signal t_n may

³ Other recent examples of settings in which notions of pivotalness and influence play a key role include Swinkels's [13] work on large, multi-unit auctions; Segal's [12] model of contracting between a principal and N agents; and Al-Najjar [1] model of free-riding among opponents in challenging the reputation of a central authority.

represent player n 's type, a piece of information correlated with his unknown type, a signal correlated with an unobserved action he has taken, etc. The specific interpretation of the signal t_n will therefore depend on the intended application.

We focus on the case where each T_n is finite (in Section 3.3 we consider the case of continuous signal sets). The profile of random signals is denoted $\tilde{\mathbf{t}} = (\tilde{t}_1, \dots, \tilde{t}_N)$ and takes values in $T = T_1 \times \dots \times T_N$. We use $\tilde{\mathbf{t}}_{-n}$ and \mathbf{t}_{-n} to denote, respectively, profiles of random signals and vectors of signal realizations for all players other than player n , and write $T_{-n} = T_1 \times \dots \times T_{n-1} \times T_{n+1} \times \dots \times T_N$.

Signals are generated according to a joint probability distribution P on T , with E denoting expectations with respect to P . To simplify the exposition, assume for the moment that the signals are independent, so P is the product of its marginals on T_{-n} and T_n . Section 3.2 extends the results to environments with correlated signals.

Let $\Delta(T_n)$ denote the simplex representing all probability distributions on T_n , and $\tilde{t}_n \in \Delta(T_n)$ the distribution of \tilde{t}_n . For $\varepsilon > 0$, let $\Delta_\varepsilon(T_n)$ denote the set of distributions $\tilde{t}_n \in \Delta(T_n)$ in which each signal has probability at least ε , and let

$$\Delta_\varepsilon^N = \Delta_\varepsilon(T_1) \times \dots \times \Delta_\varepsilon(T_N).$$

Theorems 1 and 2 will require the following assumption:

ASSUMPTION A1. There is $\varepsilon > 0$ such that $\tilde{t}_n \in \Delta_\varepsilon(T_n)$ for all n .

Each vector of signals \mathbf{t} is mapped into a collective outcome in $[0,1]$ by an *outcome function*, or mechanism, $F: T \rightarrow [0, 1]$.⁴ Here, $F(\mathbf{t})$ can denote either the value of an aggregate outcome, e.g., the level of pollution, output of team production, a principal's reward, etc., or the probability of a binary outcome, e.g., the probability that a public project is undertaken.

2.2. Influence

Denote the expected impact of player n on the outcome as his signal changes from t to t' by:

$$V_n(F, \tilde{\mathbf{t}}; t, t') = E(F | t_n = t) - E(F | t_n = t')$$

⁴ Our results go through (with appropriate re-scaling of the bounds) if the interval $[0,1]$ is replaced with any bounded subset of real numbers.

We define player n 's *influence* relative to F and $\tilde{\mathbf{t}}$ to be:

$$V_n(F, \tilde{\mathbf{t}}) = \max_{t, t'} V_n(F, \tilde{\mathbf{t}}; t, t').$$

This definition of influence is a natural extension of the concept of pivot probability to much broader settings than those considered in the voting literature (e.g., Chamberlain and Rothchild [3], and Myerson [7], among many others). In particular, F is not necessarily binary or anonymous, and players can have arbitrary sets of signals. This extension makes it possible to unify the treatment of many problems not traditionally studied in terms of pivot probability.

To motivate the use of this definition in applications, consider first a problem of designing a mechanism F to determine a public outcome (e.g., the provision of a public good; see Section 4). The influence of player n , V_n , is defined assuming that this player controls his own signal, but uncertain about the realization of other players' signals.⁵ If player n 's true type is t , incentive compatibility requires that this player does not benefit from reporting some other type t' . A mechanism can condition the outcome on the players' reports in a very complex way. Thus, the impact of player n 's report of t' depends on the signals \mathbf{t}_{-n} reported by the remaining individuals. But, since individual n is unsure about the realization of \mathbf{t}_{-n} , the relevant object in his decision problem is the change in the expected impact of his report on the collective outcome. This is the notion of influence captured in the definition above.

Another type of applications are those where a player n takes an unobserved action that generates a signal $t \in T_n$. Player n in this case controls the probability with which his signal is observed, and takes into account the impact of his action on an outcome F only if his influence V_n is large. Thus, in settings with unobserved actions, our concept of influence also captures the idea of how strong an impact a player's action has on an outcome.

2.3. Majority Rules and Pivot Probabilities

Our main results characterize the maximum number of pivotal players and average influence in terms of the pivot probabilities under simple majority rule in the following simple environment: Each of N players has three signals (Yes, No, Abstain) with probabilities $(\varepsilon, \varepsilon, 1 - 2\varepsilon)$, $0 < \varepsilon < \frac{1}{3}$; the outcome is determined by $F(\mathbf{t}) = 1$ if $\#\{t_n : t_n = \text{Yes}\} > \#\{t_n : t_n = \text{No}\}$ and 0 otherwise. Obviously, all agents have the same influence, equal to the probability that they cast the decisive vote:

⁵ In some applications, this definition is stronger than necessary because it might ignore useful model-specific information which could rule out some changes as being clearly suboptimal for the player. The definition can be modified accordingly by defining the max on a more restricted set of signal pairs.

$$R_{\varepsilon, N} = \sum_{k=0}^{N-1} P(K) \rho_k,$$

where $P(K)$ is the probability that K players out of the remaining $N - 1$ cast a Yes/No vote and $\rho_K \equiv \max_{k=0, \dots, K} \binom{K}{k} 0.5^k 0.5^{K-k}$ is the maximal binomial probability of K independent Bernoulli trials.

The exact value of $R_{\varepsilon, N}$ for low values of N can be derived from the definition of the binomial distribution. For large N , we have the following approximation:⁶

$$R_{\varepsilon, N} \simeq \frac{1}{\sqrt{\varepsilon\pi}} \frac{1}{\sqrt{N}}.$$

3. NON-PIVOTALNESS THEOREMS

3.1. Main Results

Given a threshold $\alpha \in (0, 1)$, player n is α -pivotal relative to F and $\tilde{\mathbf{t}}$, if $V_n(F, \tilde{\mathbf{t}}) \geq \alpha$. The number of α -pivotal players is

$$K(F, \tilde{\mathbf{t}}, \alpha) = \#\{n : V_n(F, \tilde{\mathbf{t}}) \geq \alpha\}.$$

This measure is useful when there is a critical threshold such that a player completely ignores the impact of his actions if his influence drops below this threshold.

Let K_α^* be the smallest integer K satisfying $R_{\varepsilon, K} \geq \alpha$. Note that K_α^* is completely determined by ε and α , but otherwise independent of N , $\tilde{\mathbf{t}}$, and F . Our main result states that the number of pivotal players is bounded, and describes when the maximum is achieved:

THEOREM 1. *For any $0 < \alpha < 1$, N , $\tilde{\mathbf{t}} \in \Delta_\varepsilon^N$ and F , the number α -pivotal players $K(F, \tilde{\mathbf{t}}, \alpha)$ is bounded by K_α^* .*

⁶ Chebyshev's inequality implies that with high probability the number of voting players is approximately $K = 2\varepsilon N$. Conditional on this value of K , the Yes's and no's are i.i.d. with mean 0.5, so the application of a majority rule means that a player's influence depends on the odds that there are exactly $k = \varepsilon N$ Yes's. The estimate now follows from Stirling's formula. The error in Stirling's formula is bounded by a function which converges very rapidly to zero. Thus, the error in our approximation of ρ also improves rapidly in N . For example, for $N = 10$, the error in Stirling's formula is no more than 0.8%, so our approximation is inaccurate by no more than: $\max\{|(1 - 0.008/(1 + 0.008)^2) - 1|, |(1 + 0.008/(1 + 0.008)^2) - 1|\} \leq 0.025$, or less than 2.5%. Similarly, for $N = 100$, the error in our estimate of ρ is no more than 0.2%.

This bound is achieved in a symmetric environment by an outcome function that applies a majority rule relative to K_α^ players and ignores the signals of the remaining players.*⁷

An alternative measure of aggregate influence is average influence:

$$V(F; \tilde{\mathbf{t}}) = \frac{1}{N} \sum_{n=1}^N V_n(F, \tilde{\mathbf{t}}),$$

This measure is useful in dealing with problems, like the provision of a public good, where the influence of a large number of agents is aggregated (see Section 4 for an example). The next theorem says that average influence is maximized using a majority rule applied to a symmetric profile.

THEOREM 2. *For any $\varepsilon > 0$, $\tilde{\mathbf{t}} \in \Delta_\varepsilon^N$ and F ,*

$$V(F, \tilde{\mathbf{t}}) \leq R_{\varepsilon, N}. \quad (*)$$

*This bound is tight: It is achieved in a symmetric environment and an outcome function that takes the form of a majority rule.*⁸

It is easy to see that $NR_{\varepsilon, N}$ is unbounded, so (*) is in principle consistent with an unbounded number of α -pivotal players for any fixed α . Theorem 1 is therefore not a consequence of the bound (*) in Theorem 2. Rather, our approach is to first characterize mechanisms maximizing average influence, then use this characterization to prove the two theorems above.

3.2. Correlated Signals

Theorems 1 and 2 do not hold if the signals are correlated:

⁷ Without noise, it is possible to find mechanisms that make every player 1-pivotal, so average influence equals 1 regardless of how large N is. In a public good setting with publicly known valuations, this can be exploited to induce individuals to contribute their true valuations via “discontinuous,” pivot mechanisms that make building the project contingent on every player reporting the truth.

⁸ Although our exposition is restricted to binary outcomes, one can easily extend the analysis to a finite number of L outcomes by replicating the argument for the one outcome case. Observe that the distance in the L -dimension simplex, which is contained R^{L-1} , is bounded by the sum of the distances along the axes. Thus, the influence for L outcomes is bounded by $L - 1$ times the bound for the binary case. Of course, the mechanism maximizing influence in this case is no longer a majority rule.

EXAMPLE. There are $2N$ players each with two signals (0 or 1). The distribution P is obtained by first selecting at random a subset of N players, then assigning to members of that set signal 1 and 0 to the remaining players. Note that for every n , $P(t_n = 1) = P(t_n = 0) = \frac{1}{2}$, so for $\varepsilon < \frac{1}{2}$, $\tilde{\mathbf{t}}$ belong to $\mathcal{A}_\varepsilon^N$. Consider the mechanism $F(\mathbf{t}) = 1$ if the signals of exactly N players are 1 and $F(\mathbf{t}) = 0$ otherwise. Then every player is fully pivotal, and average influence is equal to 1 regardless of how large N is.

The problem in this example is that the information contained in players' signals is non-exclusive, in the sense of Postlewaite and Schmeidler [9]. They define non-exclusive information to mean that the signal of any individual player can be inferred by pooling the signals of the remaining $N - 1$ players.⁹

To extend our results to environments with correlated signals, we need to rule out correlations that allow a perfect prediction of a player's signal from the pooled information contained in the signals of the remaining players. To this end, we assume that signals are generated as the outcome of a two-stage lottery in the following sense: Let Θ be a finite set of aggregate parameters and P a joint distribution on $T \times \Theta$, with $P(t | \theta)$ denoting the conditional probability of \mathbf{t} given θ (we assume, without loss of generality, that $P(\theta) > 0$ for all θ).

For $\theta \in \Theta$, let $P(\cdot | \theta)$ denote the conditional probability on T given θ . We make the following assumption:

ASSUMPTION A2 (Conditional independence). For every $\theta \in \Theta$, the signals $(\tilde{t}_1, \dots, \tilde{t}_N)$ are independent given θ , and their distribution $P(\cdot | \theta)$ belongs to $\mathcal{A}_\varepsilon^N$.

Roughly, assumption A2 says that: (1) θ is a sufficient statistic summarizing all that a player hopes to infer about \mathbf{t}_{-1} from his private signal t_n ; and (2) even if player n knew θ , there is still enough residual randomness about other players' signals.

Let $P(\mathbf{t}_{-n} | t_n)$ and $P(\theta | t_n)$ denote the posterior distributions of player n on T_{-n} and Θ respectively. The conditional independence assumption implies that $P(\mathbf{t}_{-n} | t_n, \theta) = P(\mathbf{t}_{-n} | \theta)$. In particular, for any random variable f that is measurable with respect to \mathbf{t}_{-n} , we have $E(f | t_n, \theta) = E(f | \theta)$. We will apply this fact later to the random variables $F(\mathbf{t}_{-n}, t_n = t)$, for those values of t that achieve player n 's maximum influence.

We introduce two notions of influence, each corresponding to a different assumption about the information available to a player. Under the first scenario, player n is first informed that his signal is t_n , then determines the

⁹ Formally, information is non-exclusive if for any player n and any vector of signals \mathbf{t} that has positive probability, $P(t_n = 1 | t_{-n})$ is either zero or one.

pair of signals which generates the greatest expected impact on F . Note that when signals are correlated, knowing t_n may be useful in increasing influence because the player can use that information to better forecast other players' signals. In this setting, a natural analogue of our definition in the case of independence is to require that player n computes his influence as before, using the posterior belief $P(\mathbf{t}_{-n} | t_n)$. That is:

$$\text{Player } n\text{'s influence at } t_n: \quad V_n(F; t_n) = \max_{t \in T_n} E(F(\mathbf{t}_{-n}, t) | t_n) \\ - \min_{t \in T_n} E(F(\mathbf{t}_{-n}, t) | t_n),$$

$$\text{Average influence at } \mathbf{t}: \quad V(F; \mathbf{t}) = \frac{1}{N} \sum_n V_n(F; t_n)$$

$$\text{Expected average influence:} \quad V(F) = \sum_{\mathbf{t}} P(\mathbf{t}) V(F; \mathbf{t}).$$

Note that under this definition, different players will typically have different posteriors about the distribution of signals of other players.

In our second definition, player n is informed of both his signal t_n and the aggregate parameter θ . He then determines the pair of actions that yields the greatest expected influence. Our assumption of conditional independence implies that his private signal t_n is superfluous because all information relevant to forecasting the signals of others is contained in θ . Thus, we can define player n 's influence at (t_n, θ) in terms of the posterior $P(\cdot | \theta)$:

$$V_n(F; \theta) = \max_{t_n \in T_n} E(F(\mathbf{t}_{-n}, t_n) | \theta) - \min_{t_n \in T_n} E(F(\mathbf{t}_{-n}, t_n) | \theta)$$

and define

$$\text{Average influence at } \theta: \quad V(F; \theta) = \frac{1}{N} \sum_n V_n(F; \theta)$$

$$\text{Expected average influence:} \quad V^\theta(F) = \sum_{\theta} P(\theta) V(F; \theta).$$

Note that in this case players agree on the probability distribution used in computing expectations. Thus, conditional on θ , this setting is covered by our result for the independent signal case applied to the distribution $P(\cdot | \theta)$. Our bound for the case of independent signals applies to $V(F; \theta)$ and $V^\theta(F)$.

THEOREM 3. Under assumption A2,

- (i) $\sum_{\mathbf{t}} P(\mathbf{t}) V(F; \mathbf{t}) \leq \sum_{\theta} P(\theta) V(F; \theta)$;
- (ii) For any $\varepsilon > 0$,

$$\sum_{\theta} P(\theta) V(F; \theta) \leq R_{\varepsilon, N}.$$

3.3. Continuum of Signals

We extend the analysis to the case of continuous signal spaces. This allows us to consider important applications (e.g., a continuum of possible valuations) as well as clarify how the bounds of Theorems 1 and 2 change as each player has a finite but increasingly large number of possible signals.

A key parameter in our bound is the probability ε of the least likely signal. With M signals per player, this probability cannot exceed $\frac{1}{M}$. As M increases, the minimum probability decreases so our bound becomes weaker. In the limit when all players have a continuum of possible signals it is easy to design a mechanism F such that each player has maximum influence of 1:

EXAMPLE. \tilde{t}_n is uniformly distributed on $T_n = [0, 1]$ for every n . $F(\mathbf{t}) = 0$ if exactly one signal is 0, $F(\mathbf{t}) = 1$ if one signal is 1, and $F = 0.5$ otherwise. Using our earlier definition, every player has influence equal to 1 as he moves from signal 0 to signal 1. On the other hand, $F(\mathbf{t}) = 0.5$ with probability 1, so F is in fact independent of the signal of any individual player, suggesting that no player has any influence in this case.

The formal definition we provide below captures the observation that a player's influence should reflect the change he can cause in the outcome over a wide range (but not necessarily all) signals. It is convenient to be explicit about the underlying probability space on which the random signals are defined. Specifically, we will assume that each agent's signal is defined on a non-atomic probability space (Ω, Σ, P) . Player n 's signal is a random variable $\tilde{t}_n: \Omega \rightarrow T_n$, where T_n is an arbitrary set. It is clear that what is relevant for this player's influence is the σ -algebra \mathcal{G}_n generated by \tilde{t}_n rather than the signal space T_n itself. A mechanism F is a random variable $F: \Omega \rightarrow [0, 1]$ and the ability of player n to influence F based on his signal t is represented by the conditional expectation $E(F | \mathcal{G}_n)(t)$.

The finite-signal case can be viewed as the special case where each \mathcal{G}_n is a finite partition in which each atom has probability at least ε . In this case, the conditional expectation $E(F | \mathcal{G}_n)(t)$ takes only finitely many possible values. Our earlier definition of influence of player n , $V_n(F)$, is then just the difference between the highest and the lowest values of the conditional

expectation of F viewed as a function of t_n . With a continuum of signals and using this definition, player n can have maximum influence of 1 even though F may be independent of his signal with probability 1.

One way to eliminate this problem can be roughly explained as follows: Fix $0 < \varepsilon < 1$, and remove a set of signals A_n^+ of measure ε on which $E(F | \mathcal{G}_n)(t)$ assumes its highest values. In the example discussed earlier, this would be any subset containing the point $t_n = 1$. Similarly, we remove a set of signals A_n^- of measure ε on which $E(F | \mathcal{G}_n)(t)$ assumes its lowest values, excluding the point $t_n = 0$. Let $A_n = T_n - (A_n^+ \cup A_n^-)$. That is, A_n represents the signal space after removing the two extreme sets A_n^+ and A_n^- . We then define influence as the difference between the highest and lowest value of $E(F | \mathcal{G}_n)(t)$ as t ranges over the set A_n . With this definition, a player has small influence if he cannot change by much the conditional expectation of the outcome by moving his signal over a set of large measure.

To make this definition precise, fix versions of the conditional expectations $E(F | \mathcal{G}_n)(t)$ and a parameter $0 < \varepsilon < 1$. We define the influence of player n relative to F and ε :

$$V_n(F, \varepsilon) = \inf_{\{A \in \Sigma : P(A) < \varepsilon\}} \sup_{t \notin A} E(F | \mathcal{G}_n)(t) - \sup_{\{A \in \Sigma : P(A) < \varepsilon\}} \inf_{t \notin A} E(F | \mathcal{G}_n)(t).$$

The complicated appearance of this definition stems from the need that it: (1) coincides with that in Section 3.1 when a player has a finite number of signals; and (2) does not depend on the particular version of conditional expectations selected.¹⁰ Average influence is defined in the usual way:

$$V(F, \varepsilon) = \frac{1}{N} \sum_{n=1}^N V_n(F, \varepsilon).$$

We call an agent (ε, α) -pivotal if $V(F, \varepsilon) \geq \alpha$.

THEOREM 4. Fix $\varepsilon > 0$ and N . Then for any set of signals $\{\mathcal{G}_1, \dots, \mathcal{G}_N\}$ and any mechanism F ,

$$V(F, \varepsilon) \leq R_{\varepsilon, N}.$$

¹⁰ That is, independent of the values of $E(F | \mathcal{G}_n)(t)$ over sets of signals of measure zero. To see this, let $x(t)$ and $y(t)$ be two versions and let E be the set of measure zero on which they disagree. If $\inf_{\{A : P(A) < \varepsilon\}} \sup_{t \notin A} y(t) > \inf_{\{A : P(A) < \varepsilon\}} \sup_{t \notin A} x(t)$, then there must be a set B , with $P(B) < \varepsilon$ such that $\inf_{\{A : P(A) < \varepsilon\}} \sup_{t \notin A} y(t) > \sup_{t \in B^c} x(t)$. Since E has measure zero, $P(B \cup E) < \varepsilon$, so $\inf_{\{A : P(A) < \varepsilon\}} \sup_{t \notin A} y(t) \leq \sup_{t \in B^c \cap E^c} y(t) = \sup_{t \in B^c \cap E^c} x(t) \leq \sup_{t \in B^c} x(t)$, which is a contradiction.

3.4. *Related Results*

We now relate our non-pivotalness theorems to similar ideas in the literature. The proof in the Appendix of Mailath and Postlewaite [6] exploits the fact that a random variable f , viewed as a point in a linear space, cannot have high covariance with many members of an orthonormal basis for that space. Interpreting members of the basis as the random types of agents, they conclude that not many agents can accurately predict f based on their signals. More specifically, in an environment where outcomes and signals may be continuous, they show that the ratio of agents whose signals are highly correlated with f goes to zero at the rate of $N^{-1/2}$. Fudenberg, Levine and Pesendorfer ([4], Lemma A), in an environment with independent, discrete signals, derive a bound on average influence on the probability of a subset of outcomes. They show that this bound converges to zero at the rate of $N^{-1/2}$.

The present paper differs from this literature in that it provides exact, finite-population bounds applicable in more general environments, as well as a general methodology explaining what sort of mechanism achieve these bounds. Our main result, Theorem 1, replaces asymptotic bounds on the fraction of pivotal players found in the literature by exact bound on the *number* of such players. The difference is crucial in applications, such as information aggregation, where the number of pivotal agents—rather than their ratio—is what matters. Similarly, our treatment of environments with correlated signals (Theorem 3) has no counterpart in these works. (Our analysis also handles the continuous outcome/signal case not covered by Fudenberg *et al.*) Finally, our analysis clarifies the role of anonymity by showing that characterizing influence for potentially complex, non-anonymous mechanisms reduces to looking at simple, anonymous mechanisms taking the form of majority rule. This leads to the surprising conclusion that majority rules already contain the maximal number of pivotal players; so, no further gain can be made by considering more complex mechanisms.

4. EXAMPLE: PROVISION OF A PUBLIC GOOD WITH PRIVATE INFORMATION

We provide a simple, transparent bound on average contributions in finite public good economies under private information. From this we derive Mailath and Postlewaite's [6] asymptotic inefficiency results. The analysis also covers, with obvious modifications, Rob's [10] of model of externalities.

The key observation underlying the analysis is that incentive compatibility requires the expected contribution of any individual to be bounded by his influence on the probability of provision. Our analysis of influence may then be applied to derive the desired result. Aside from its conceptual simplicity, this argument readily provides numerical bounds on the severity of the free-rider problem in finite populations.

Consider a public good economy \mathcal{E}_N with N individuals, and a public project costing C_N . As in Mailath and Postlewaite, the per capita cost of the project is bounded away from zero: there is $\beta > 0$ such that $C_N \geq \beta_N$ uniformly in N . All uncertainty is defined on a probability space (Ω, Σ, P) . Individual n has a privately known valuation for the public project, which we model as a random variable $t_n(\omega)$ with values in $[t_n^-, t^+]$. We normalize $t_n^- = 0$ for all N , assume that valuations are independent, and that there is an $\varepsilon > 0$ such that $P(t_n = t_n^-) \geq \varepsilon$ uniformly across all agents n and economies \mathcal{E}_N .

There are two possible collective outcomes corresponding to whether the project is built or not. A voluntary contribution mechanism is a pair (δ, c) , where $\delta: \Omega \rightarrow \{0, 1\}$ is a random variable indicating whether the public good is provided, and c is a vector of contributions (c_1, \dots, c_N) , where $c_n: \Omega \rightarrow R$ denotes the amount contributed by individual n . We restrict attention to direct revelation mechanisms in which each agent truthfully reports his type. Individual n 's payoff under (δ, c) in state ω when he truthfully reports his type is:¹¹

$$u_n(\omega) = t_n(\omega) \delta(\omega) - c_n(\omega).$$

We require (δ, c) to be (ex-ante) *budget balanced*:

$$C_N E\delta \leq \sum_n E c_n, \quad (BB)$$

and ex ante *individually rational*:

$$E(u_n | t_n)(\omega) \geq 0 \quad \text{for all } n, P\text{-a.s.} \quad (IR)$$

Our analysis of influence enters through the incentive compatibility constraint: there is a subset $\Omega' \subset \Omega$, with $P(\Omega') = 1$ such that, for all $\omega, \omega' \in \Omega'$, and n :

$$t_n(\omega) E(\delta | t_n)(\omega) - E(c_n | t_n)(\omega) \geq t_n(\omega) E(\delta | t_n)(\omega') - E(c_n | t_n)(\omega'). \quad (IC)$$

¹¹ This formulation covers mechanisms which allow for reimbursing (portions of the) contributions if the project is not built. For example, the case of full reimbursement can be expressed by requiring $c_n(\omega) = 0$ if $\delta(\omega) = 0$.

By (IR), we have $E(c_n | t_n = 0) \leq 0$, so (IC) implies

$$E(c_n | t_n)(\omega) \leq t_n(\omega)[E(\delta | t_n)(\omega) - E(\delta | t_n)(\omega')] \leq t_n(\omega) V_n(\delta).$$

That is, agent n 's expected contribution per dollar valuation for the public good is bounded by the influence of his report on the probability of provision. This influence is, of course, bounded by $V_n(\delta)$. The inefficiency result below then follows from our bounds on influence:

PROPOSITION.

(i) For every $0 < \eta \leq \varepsilon$, the probability of provision $E\delta$ satisfies:

$$\sup_{(\delta, c)} E\delta \leq \frac{t^+}{\beta} [R_{\eta, N} + \eta]$$

where the sup is taken over all mechanisms (δ, c) satisfying IR, IC, and BB;

(ii) $\lim_{N \rightarrow \infty} \sup_{(\delta, c)} E\delta = 0$.

Part (ii) is the counterpart of Mailath and Postlewaite's conclusion of asymptotic inefficiency (the main difference is that the rate of convergence here is $N^{-1/2}$ instead of $N^{-1/4}$). Part (i) provides a finite-population bounds on the probability of provision, making it possible to assess the severity of the free-riding problem in settings where asymptotic arguments may be inappropriate.¹²

To compute numerical bounds on the probability of provision $E\delta$, it is more convenient to work with the following corollary to the proof of the Proposition:

COROLLARY. Suppose that T_n is finite and there is $\varepsilon > 0$ such that for every n , $t_n \in T_n P(t_n) \geq \varepsilon$. Then

$$\sup_{(\delta, c)} E\delta \leq \frac{\max_n Et_n}{\beta} R_{\varepsilon, N}$$

where the sup is taken over all mechanisms (δ, c) satisfying IR, IC, and BB.

Proof. From the definition of influence, for every n , $E(\delta | t_n)(\omega) - E(\delta | t_n = 0) \leq V_n(\delta)$ for all n and t_n . This implies $E(c_n,) \leq \sum_{t_n} t_n [E(\delta | t_n)(\omega) - E(\delta | t_n = 0)] P(t_n) \leq \max_n Et_n V_n(\delta)$. From (BB), we have $E(\delta) \leq \sum Ec_n / C_N \leq (\max_n Et_n / \beta) V(\delta) \leq (\max_n Et_n / \beta) R_{\varepsilon, N}$.

With $C_N = N$, $\beta = 1$, $\varepsilon = .10$ and $E_n t_n = 2$ for all n , each individual on average values the public good twice as much as the expected per capita

¹² The proof in fact yields a somewhat sharper bound $(\max_n Et_n / \beta) R_{\eta, N} + \frac{t^+}{\beta} \eta$.

cost. The following table provides the value of $R_{e,N}$ and the maximum probability of provision for different values of N :

N	$R_{e,N}$ (exact)	$R_{e,N}$ (estimate)	$E\delta$ at most
30	0.320	0.325	64%
100	0.177	0.178	36%
200	0.125	0.126	25%
500	0.079	0.079	16%
1000	0.056	0.056	12%
10,000	—	0.017	3.5%

The bound on the probability of provision declines with N slowly at the rate $N^{-1/2}$. Note that the Weak Law of Large Numbers implies that for moderately large values of N , with high probability, the sum of individual valuations for the public good is approximately twice as large as its cost, in which case provision is efficient.

APPENDIX

We begin with some definitions used throughout the proof. An environment is *symmetric* if all players have the same signal sets; i.e., $T_n = T_m$ for any pair of players n and m . In such environments, a profile \mathbf{t} is *symmetric* if agents' signals are identically distributed: $\tilde{t}_n = \tilde{t}_m$ for all n, m . In symmetric environments it makes sense to talk about anonymous mechanisms that ignore the agents' names. Specifically, F is *anonymous relative to a subset of players K* if $F(\mathbf{t}) = F(\sigma(\mathbf{t}))$ for any permutation σ of the names of players in K .¹³ We call F *anonymous* if $K = N$. With M signals per player, anonymity can be equivalently defined in terms of the vector $d(\mathbf{t}) = (d_1(\mathbf{t}), \dots, d_M(\mathbf{t}))$ of empirical frequencies of the signals (that is, $d_m(\mathbf{t})$ is the number of times the m th signal is observed, divided by N). It is easy to verify that F is anonymous if and only if $F(\mathbf{t})$ depends on \mathbf{t} only through $d(\mathbf{t})$. A special class of anonymous mechanisms is that of majority rules. Formally, F is a *majority rule* if there are two signals m and m' such that for any \mathbf{t} ,

$$d_m(\mathbf{t}) > d_{m'}(\mathbf{t}) \Rightarrow F(\mathbf{t}) = 1, \quad \text{and}$$

$$d_m(\mathbf{t}) \leq d_{m'}(\mathbf{t}) \Rightarrow F(\mathbf{t}) = 0.$$

¹³ That is, permutations $\sigma: \{1, \dots, N\} \rightarrow \{1, \dots, N\}$ such that $\sigma(n) = n$ for all $n \notin K$.

Let M_n denote the cardinality of T_n . For a player n let t_n^+ and t_n^- denote the pair of signals at which his maximum influence is achieved (that is, $V_n(F, \tilde{\mathbf{t}}) = V_n(F, \tilde{\mathbf{t}}; t_n^+, t_n^-)$). It is also convenient to define his *conditional influence* given the signal of player n' to be:

$$V_n(F, \tilde{\mathbf{t}} | t_{n'} = t_{n'}^m) = \sum_{\mathbf{t}_{-n}} P(\mathbf{t}_{-n} | t_{n'} = t_{n'}^m) [F(\mathbf{t}_{-n}, t_n = t_n^+) - F(\mathbf{t}_{-n}, t_n = t_n^-)]$$

With this notation, average influence $V(F, \tilde{\mathbf{t}})$ can be expressed in terms of the probability distribution of player n 's signals as:

$$N V(F; \tilde{\mathbf{t}}) = V_n(F, \tilde{\mathbf{t}}; t_n^+, t_n^-) + \underbrace{\sum_{m=1}^{M_n} P(\tilde{t}_n = t_n^m) \sum_{n' \neq n}^N V_{n'}(F, \tilde{\mathbf{t}} | t_n = t_n^m)}_{A_m}$$

The second part of this expression, $\sum_{m=1}^{M_n} P(\tilde{t}_n = t_n^m) A_m$, is a linear function of agent n 's distribution $P(\tilde{t}_n = t_n^m)$ and is therefore maximized at some signal which we denote t_n^{\max} .

A signal distribution \tilde{t}_n is ε -extremal (or, simply, *extremal* if ε is clear from the context) if it coincides with a vertex of the trimmed simplex $\Delta_\varepsilon(T_n)$. That is, \tilde{t}_n is ε -extremal if it puts probability ε on all but possibly one signal. Denote the set of extremal distributions for player n by $ext \Delta_\varepsilon(T_n)$, and profiles of extremal distributions by $ext \Delta_\varepsilon^N$.

PROPOSITION. A.1. Fix $\varepsilon > 0$, F and $\tilde{\mathbf{t}} \in \Delta_\varepsilon^N$. Then there is a restricted signal set \hat{T}_1 of three distinct signals for player 1, a distribution $\hat{t}_1 \in ext \Delta_\varepsilon(\hat{T}_1)$ and a mechanism $\hat{F}: \hat{T}_1 \times T_{-n} \rightarrow [0, 1]$ such that

$$V(F, \tilde{\mathbf{t}}) \leq V(\hat{F}, \hat{t}_1 \times \tilde{\mathbf{t}}_{-n})$$

Proof. Consider the subset of signals $\hat{T}'_1 = \{t_1^+, t_1^-, t_1^{\max}\}$ in which some of the signals may be repeating (this will necessarily be the case if player 1 has only two signals). Consider the distribution \hat{t}'_1 which assigns to each of t_1^+, t_1^- probability ε and probability $1 - 2\varepsilon$ to t_1^{\max} . Define $\hat{F}': \hat{T}'_1 \times T_{-n} \rightarrow [0, 1]$ as the natural restriction of the original F to the restricted signals space \hat{T}'_1 . Since individual 1's influence depends on the value taken by F at $\{t_1^+, t_1^-\}$ only, his contribution to total influence is unaffected by this change in the signal space and the distribution. From the definition of t_1^{\max} we also have that shifting weight $1 - 2\varepsilon$ to it (weakly) increases the total contribution of other players to total influence. Thus, $V(F, \tilde{\mathbf{t}}) \leq V(\hat{F}', \hat{t}'_1 \times \tilde{\mathbf{t}}_{-n})$.

We now convert \hat{T}'_1 to a set of three distinct signals \hat{T}_1 . For later use, it will be notationally convenient to choose $\hat{T}_1 = \{0, 1, 2\}$ to be a standard signal space common to all players. If the signals $\{t_1^+, t_1^-, t_1^{\max}\}$ are all distinct, then identify t_1^{\max} with 1, t_1^+ with 2, and t_1^- with 0, and define $\hat{F}: \hat{T}_1 \times T_{-n} \rightarrow [0, 1]$ to coincide with \hat{F}'_1 under this identification of signals. We now turn to the various cases in which $\{t_1^+, t_1^-, t_1^{\max}\}$ fail to be distinct.

Assume first that $t_1^+ = t_1^{\max}$, so the probability of this combined signal is actually $1 - \varepsilon$. In this case, identify t_1^{\max} with 1, and t_1^- with 0. We then split from signal 1 a new signal 2 that carries probability ε (so signal 1 is now left with probability $1 - 2\varepsilon$). Define the mechanism \hat{F}_1 so that signal 2 is redundant in the sense that \hat{F} treats signal 2 in exactly the same way as signal 1. That is, the new mechanism $\hat{F}: \hat{T}_1 \times T_{-n} \rightarrow [0, 1]$ is defined by

$$\hat{F}(0, \mathbf{t}_{-1}) = F(t_1^-, \mathbf{t}_{-1}),$$

$$\hat{F}(1, \mathbf{t}_{-1}) = F(t_1^+, \mathbf{t}_{-1}),$$

$$\hat{F}(2, \mathbf{t}_{-1}) = F(t_1^+, \mathbf{t}_{-1}).$$

Note that this does not affect the influence of any individual. The proof for the case in which $t_1^- = t_1^{\max}$ is similar. Finally, if $t_1^+ = t_1^-$, then the mass of 2ε assigned to t_1^{\max} can be split over two signals which F treats in the same manner, and this can be done without reducing influence. This completes the proof.

Call $(\{\hat{T}_n\}_{n=1}^N, \tilde{\mathbf{t}})$ a *standard environment* if: (1) all agents have the same signal sets $\hat{T}_n = \{0, 1, 2\}$; (2) each agent has an extremal distribution such that signal 1 has probability $1 - 2\varepsilon$. In such environment, call a mechanism F *regular* if each agent's maximum influence is achieved when his signal changes from 0 to 2.

PROPOSITION. A.2. *Fix $\varepsilon > 0$, F and $\tilde{\mathbf{t}} \in \Delta_\varepsilon^N$. Then there is a regular mechanism \hat{F} in the standard environment $(\{\hat{T}_n\}_{n=1}^N, \hat{\mathbf{t}})$ such that*

$$V(F, \tilde{\mathbf{t}}) \leq V(\hat{F}, \hat{\mathbf{t}}).$$

Proof. Apply Proposition A.1 relative to player 1 to obtain a new signal space $\hat{T}_1 \times T_2 \times \cdots \times T_N$, a vector of random signals $\hat{\mathbf{t}}^1 = \hat{t}_1 \times \tilde{\mathbf{t}}_{-n}$, and a mechanism \hat{F}_1 so that $V(F, \tilde{\mathbf{t}}) \leq V(\hat{F}_1, \hat{\mathbf{t}}^1)$. Repeating this process for player 2 relative to \hat{F}_1 and $\hat{\mathbf{t}}^1$ yields a new mechanism \hat{F}_2 and profile $\hat{\mathbf{t}}^2$ such that average influence does not decrease. Continuing in this manner for all remaining players, we obtain a sequence of pairs $(\hat{F}_n, \hat{\mathbf{t}}^n)$ along which $V(\hat{F}_n, \hat{\mathbf{t}}^n)$ is increasing. The claim is proved by setting $\hat{F} = \hat{F}_N$ and $\hat{\mathbf{t}} = \hat{\mathbf{t}}^N$. The new distribution $\hat{\mathbf{t}}$ is symmetric and extremal by construction.

PROPOSITION. A.3. *In the standard environment $(\{\hat{T}_n\}_{n=1}^N, \tilde{\mathbf{t}})$, for any regular mechanism F there is an anonymous regular mechanism F' such that $V(F, \tilde{\mathbf{t}}) = V(F', \tilde{\mathbf{t}})$.*

Proof. Let σ be any permutation of the set of players' names, and $\tilde{\mathbf{t}}$ be any symmetric distribution (references to $\tilde{\mathbf{t}}$ are dropped for notational simplicity). Define the new mechanism F^σ by $F^\sigma(\mathbf{t}) = F(\sigma(\mathbf{t}))$. We show that $V_n(F^\sigma) = V_{\sigma^{-1}(n)}(F)$ for every n :

$$\begin{aligned} V_n(F^\sigma) &= E(F^\sigma \mid t_n = 2) - E(F^\sigma \mid t_n = 0) \\ &= E(F \mid t_{\sigma^{-1}(n)} = 2) - E(F \mid t_{\sigma^{-1}(n)} = 0) \\ &= V_{\sigma^{-1}(n)}(F). \end{aligned}$$

Thus,

$$V(F^\sigma) = \frac{1}{N} \sum_{n=1}^N V_n(F^\sigma) = \frac{1}{N} \sum_{n=1}^N V_{\sigma^{-1}(n)}(F) = V(F).$$

Let p be the uniform probability distribution on the set of all permutations σ over players' names. Define

$$F'(\mathbf{t}) = \sum_{\sigma} p(\sigma) F^\sigma(\mathbf{t}).$$

Obviously, $F'(\mathbf{t})$ is symmetric and

$$V(F') = \sum p(\sigma) V(F^\sigma) = \sum p(\sigma) V(F) = V(F).$$

PROPOSITION. A.4. *Fix the standard environment $(\{\hat{T}_n\}_{n=1}^N, \tilde{\mathbf{t}})$ and let F_m denote the majority rule relative to signals 0 and 2. Let F be any regular mechanism F , then*

$$V(F, \tilde{\mathbf{t}}) \leq V(F_m, \tilde{\mathbf{t}}) \leq R_{\varepsilon, N}.$$

For the proof we will need some additional notation. As before, it is convenient to think of signals 0 and 2 as “No” and “Yes” respectively, and signal 1 as representing “all-other-signals,” or “Abstain”. Fix any set of $N - 1$ players, and let $K = 0, \dots, N - 1$ be the random variable denoting the number of non-abstaining players out of this set, and let $P_{N-1}(K)$ denote its probability. Since the signal sets and the profiles are symmetric, the identity of the $N - 1$ players is irrelevant.

Proof. Let k be the random variable denoting the number of players saying “Yes” (i.e., whose signal is 2) out of $N - 1$ players, and let

$P_{N-1}(k | K)$ denote its conditional probability (or simply $P(k | K)$ when N is clear from the context).

It is convenient to express F as a function $F(k, K)$ of the number of Yes's k and the number of voting players K . For given K and k , if player n changes his signal from 0 to 2 (a change that gives him maximum influence), then there will be $K + 1$ non-abstaining players and $k + 1$ Yes's, and the outcome changes from $F(k, K + 1)$ to $F(k + 1, K + 1)$. Note that while K and k represent the relevant uncertainty from the perspective of player n , as far as F is concerned the total number of non-abstaining players is $K + 1$.

Fix F that satisfies the assumptions of the proposition (we drop references to $\tilde{\mathbf{i}}$ for notational simplicity); define

$$V_n^K = \sum_{k=0}^K P(k | K) [F(k + 1, K + 1) - F(k, K + 1)];$$

and note that player n 's influence is just:

$$V_n(F) = \sum_{K=0}^{N-1} P(K) V_n^K.$$

Thus, V_n^K represents player n 's influence conditional on there being K non-abstaining players out of the remaining $N - 1$ players.

It is convenient to rewrite V_n^K as:

$$\begin{aligned} V_n^K(F) &= -P(0 | K) F(0, K + 1) + F(k = 1, K + 1) \\ &\quad \times [P(k = 0 | K) - P(k = 1 | K)] \\ &\quad + \dots \\ &\quad + F(k, K + 1) [P(k - 1 | K) - P(k | K)] \\ &\quad + \dots \\ &\quad + F(K - 1, K + 1) [P(K - 2 | K) - P(K - 1 | K)] \\ &\quad + F(K + 1, K + 1) P(K | K). \end{aligned}$$

We view $V_n^K(F)$ as a function of the values $F(k, K + 1)$, $k = 0, \dots, K$ for fixed probability weights and with the constraints that $0 \leq F(k, K + 1) \leq 1$, for $k = 0, \dots, K$. Since the constraint set is bounded, the maximum of V_n^K is achieved, and at which point the first order conditions must be satisfied. Since $P(0 | K) > 0$ and $P(1 | K) > 0$, we must have $F(0, K + 1) = 0$ and $F(k, K + 1) = 1$.

For $k = 1, \dots, K$ we have

$$\frac{P(k | K)}{P(k - 1 | K)} > 1 \Rightarrow F(k, K + 1) = 0 \quad \text{and}$$

$$\frac{P(k | K)}{P(k - 1 | K)} < 1 \Rightarrow F(k, K + 1) = 1.$$

That is, for a fixed K , the highest value of V_n^K is achieved when $F(\cdot, K + 1)$ is zero as long as the binomial probability $P(k | K)$ is increasing in k , and 1 as long as $P(k | K)$ is decreasing in k .

Conditional on K , the signals of the K non-abstaining players are independently and identically distributed with probability 0.5 for 0 and 0.5 for 2. Thus, $P(k | K)$ is a Bernoulli distribution of K independent and identically distributed random variables with probability of success 0.5. This is a symmetric distribution whose maximum value is ρ_κ , by definition.

We show that the optimal mechanism F_m has the form of a simple majority rule. If $K = 0$, then $V_n^K = F(1, 1) - F(0, 1)$, a number which is maximized at the “majority rule” $F(1, 1) = 1$ and $F(0, 1) = 0$. For $K \geq 1$, note that

$$\frac{P(k | K)}{P(k - 1 | K)} = \frac{\binom{K}{k}}{\binom{K}{k - 1}} = \frac{K - k + 1}{k}.$$

Assume first that $K = 2L$ for some non-negative integer L (i.e., K is even). In this case, $P(k | K)$ achieves its unique maximum at $k = L$, so F_m must be of the form

$$k \leq L \Rightarrow F_m(k, K + 1) = 0$$

$$k > L \Rightarrow F_m(k, K + 1) = 1$$

which is a majority rule. If $K = 2L + 1$ for some integer L (i.e., K is odd), then $P(k | K)$ achieves its maximum at both $k = L$ and $k = L + 1$. In this case, we can set F_m to be the majority rule:

$$k \leq L + 1 \Rightarrow F_m(k, K + 1) = 0$$

$$k > L + 1 \Rightarrow F_m(k, K + 1) = 1$$

which again has the form of a majority rule.

We have therefore shown that conditional on K , the mechanism which maximizes influence is a simple majority rule for which influence is ρ_κ .

Averaging over $K=0, \dots, N-1$ with the probability distribution $P(K)$ yields the desired result.

Proof of Theorem 2. The proof follows by combining Propositions A.1 through A.4. Using Propositions A.1 and A.2 we can reduce any general problem to one in a standard environment and a regular F without reducing average influence. Proposition A.3 then shows that any such F is equivalent (i.e., yields the same influence for each player) to a mechanism which is anonymous. Finally, Proposition A.4 shows that simple majority rule is the anonymous rule with the maximum influence and computes the bound. Furthermore, the influence of such a mechanism is exactly $R_{\varepsilon, N}$, so the bound we obtain is indeed tight.

Proof of Theorem 1. The bound is clearly achieved as described in the statement of the theorem, and no anonymous mechanism on a symmetric environment can exceed this bound. Suppose, by way of contradiction, that there is $N > K_\alpha^*$ such that there is a profile $\tilde{\mathbf{t}}$ and a mechanism F with $L > K_\alpha^*$, α -pivotal players. For convenience, reorder the players so that the α -pivotal players are the first L and let $V^L(F, \tilde{\mathbf{t}}) \geq \alpha$ denote their average influence. Also, write any vector of signals \mathbf{t} as $(\mathbf{t}_L, \mathbf{t}_{-L})$. Now $V^L(F, \tilde{\mathbf{t}})$ is maximized at some vector $\tilde{\mathbf{t}}_{-L}$ of signals of players outside L . Define the new mechanism F_L by setting $F_L(\mathbf{t}) = F(\mathbf{t}_L, \tilde{\mathbf{t}}_{-L})$ for every \mathbf{t} . Clearly, $V_L(F_L, \tilde{\mathbf{t}}) \geq V_L(F, \tilde{\mathbf{t}})$. That is, F_L ignores what players outside L do, yet still increases the average influence of players in L (note, however, that there may be less than L α -pivotal players under F_L). We can now interpret F_L as a mechanism in a problem with $L > K_n^*$ players. For this problem, we know from the proof of Theorem 2 that there is an anonymous F' and a symmetric $\tilde{\mathbf{t}}'$ in a symmetric environment with L players such that $V(F', \tilde{\mathbf{t}}') \geq V(F_L, \tilde{\mathbf{t}}) \geq \alpha$. This together with the symmetry of $\tilde{\mathbf{t}}'$ and anonymity of F' imply that there are L α -pivotal players, which contradicts the definition of K_α^* , and the assumption that $L > K_\alpha^*$.

Proof of Theorem 3:

$$\begin{aligned} N V(F) &= N \sum_{\mathbf{t}} P(\mathbf{t}) V(F; \mathbf{t}) = \sum_{\mathbf{t}} P(\mathbf{t}) \sum_n V_n(F; t_n) \\ &= \sum_n \sum_{t_n} \sum_{\mathbf{t}_{-n}} P(\mathbf{t}_{-n}; t_n) V_n(F; t_n) \\ &= \sum_n \sum_{t_n} V_n(F; t_n) \sum_{\mathbf{t}_{-n}} P(\mathbf{t}_{-n}; t_n) \\ &= \sum_n \sum_{t_n} P(t_n) V_n(F; t_n). \end{aligned}$$

Let t_n^+ and t_n^- denote the signals of player n at which his influence is achieved when his actual signal is t_n . Also, let t_θ^+ and t_θ^- be the signals at which player n maximizes his influence if he knew that the aggregate state is θ . That is,

$$\begin{aligned} & \max_{t_n \in T_n} E_{\mathbf{t}_{-n}}(F(\mathbf{t}_{-n}, t_n) \mid \theta) - \min_{t_n \in T_n} E_{\mathbf{t}_{-n}}(F(\mathbf{t}_{-n}, t_n) \mid \theta) \\ & = E_{\mathbf{t}_{-n}}(F(\mathbf{t}_{-n}, t_\theta^+) \mid \theta) - E_{\mathbf{t}_{-n}}(F(\mathbf{t}_{-n}, t_\theta^-) \mid \theta). \end{aligned}$$

Then,

$$\begin{aligned} & \sum_{n=1}^N \sum_{t_n} P(t_n) V_n(F; t_n) \\ & = \sum_n \sum_{t_n} P(t_n) [E_{\mathbf{t}_{-n}}(F(\mathbf{t}_{-n}; t_n^+) \mid t_n) - E_{\mathbf{t}_{-n}}(F(\mathbf{t}_{-n}; t_n^-) \mid t_n)] \\ & = \sum_n \sum_{t_n} P(t_n) \sum_{\theta} P(\theta \mid t_n) \\ & \quad \times [E_{\mathbf{t}_{-n}}(F(\mathbf{t}_{-n}; t_n^+) \mid t_n, \theta) - E_{\mathbf{t}_{-n}}(F(\mathbf{t}_{-n}; t_n^-) \mid t_n, \theta)] \\ & \leq \sum_n \sum_{t_n} P(t_n) \sum_{\theta} P(\theta \mid t_n) \\ & \quad \times [E_{\mathbf{t}_{-n}}(F(\mathbf{t}_{-n}; t_\theta^+) \mid t_n, \theta) - E_{\mathbf{t}_{-n}}(F(\mathbf{t}_{-n}; t_\theta^-) \mid t_n, \theta)] \\ & = \sum_n \sum_{t_n} P(t_n) \sum_{\theta} P(\theta \mid t_n) [E_{\mathbf{t}_{-n}}(F(\mathbf{t}_{-n}; t_\theta^+) \mid \theta) - E_{\mathbf{t}_{-n}}(F(\mathbf{t}_{-n}; t_\theta^-) \mid \theta)] \quad (*) \\ & = \sum_n \sum_{t_n} P(t_n) \sum_{\theta} P(\theta \mid t_n) V_n(F; \theta) \\ & = \sum_n \sum_{\theta} V_n(F; \theta) \sum_{t_n} P(\theta \mid t_n) P(t_n) \\ & = \sum_{\theta} P(\theta) \sum_n V_n(F; \theta), \end{aligned}$$

where the conditional independence assumption was used to conclude that taking conditional expectation relative to t_n in the expression in square brackets in (*) is superfluous when θ is known.

Proof of Theorem 4. For each n , fix a version of the conditional expectation $E(F \mid \mathcal{G}_n)(t)$. The proof consists of constructing a sequence of mechanisms $F = \hat{F}_0, \hat{F}_1, \dots, \hat{F}_N = \hat{F}$ along which average influence increases, and such that, under \hat{F} , all players have just three signals. We then apply the bound obtained in Theorem 2 for the finite signal case to \hat{F} .

Fix $0 < \eta < \varepsilon$. For $n = 1, \dots, N$ and any mechanism F , define the functions

$$W_{-n}(F, t_n = t) = \sum_{m < n}^N V_m(F, \eta | t_n = t) + \sum_{m > n}^N V_m(F, \varepsilon | t_n = t)$$

$$W_{-n}(F) = \sum_{m < n}^N V_m(F, \eta) + \sum_{m > n}^N V_m(F, \varepsilon).$$

That is, W_{-n} represents the sum of the η -influences of players before player n and the ε -influences of players after n .

For $1 \leq n \leq N$, let \hat{F}_{n-1} be the mechanism constructed in the previous step (for $n = 1$, this is just the original mechanism F). Let A^+ be a set of measure η such that $E(F | \mathcal{G}_n)(t) \geq E(F | \mathcal{G}_n)(t')$ for any $t \in A^+$ and $t' \notin A^+$. Such set A^+ exists since P is non-atomic by assumption. Similarly, define A^- to be a set of measure η for which $E(F | \mathcal{G}_n)(t) \leq E(F | \mathcal{G}_n)(t')$ for any $t \in A^-$ and $t' \notin A^-$. Let $t_n^* \in \operatorname{argmax}_{t \in [0, 1]} W_{-n}(\hat{F}_{n-1}, t)$, $t_n^+ \in \operatorname{argmax}_{t \in A^+} W_{-n}(\hat{F}_{n-1}, t)$ and $t_n^- \in \operatorname{argmax}_{t \in A^-} W_{-n}(\hat{F}_{n-1}, t)$. If the maximum in these definitions is not achieved, then the same argument would go through by taking a sequence of signals for which $E(F | \mathcal{G}_n)$ converges to the supremum without changing the basic idea of the proof.

Define a new mechanism:

$$\hat{F}_n(t_n, \mathbf{t}_{-n}) = \begin{cases} F(t_n^+, \mathbf{t}_{-n}) & \text{if } t_n \in A^+ \\ F(t_n^-, \mathbf{t}_{-n}) & \text{if } t_n \in A^- \\ F(t_n^*, \mathbf{t}_{-n}) & \text{otherwise.} \end{cases}$$

Note that:

$$E(F | \mathcal{G}_n)(t_n^+) \geq \sup_{t \notin A^+} E(F | \mathcal{G}_n)(t) \geq \inf_{\{A : \lambda(A) < \varepsilon\}} \sup_{t \notin A} E(F | \mathcal{G}_n)(t)$$

and

$$E(F | \mathcal{G}_n)(t_n^-) \leq \inf_{t \notin A^-} E(F | \mathcal{G}_n)(t) \leq \sup_{\{A : \lambda(A) < \varepsilon\}} \inf_{t \notin A} E(F | \mathcal{G}_n)(t).$$

Thus,

$$\begin{aligned} V_n(F, \varepsilon) &\leq E(F | \mathcal{G}_n)(t_n^+) - E(F | \mathcal{G}_n)(t_n^-) \\ &= E(\hat{F}_n | \mathcal{G}_n)(t_n^+) - E(\hat{F}_n | \mathcal{G}_n)(t_n^-) \\ &= V_n(\hat{F}_n, \eta). \end{aligned}$$

By the choice of t_n^- , t_n^+ and t_n^* , we have $W_{-n}(\hat{F}_{n-1}) \leq W_{-n}(\hat{F}_n)$, and by the above argument we also have $V_n(\hat{F}_{n-1}, \varepsilon) \leq V_n(\hat{F}_n, \eta)$. We therefore conclude that:

$$W_{-n}(\hat{F}_{n-1}) + V_n(\hat{F}_{n-1}, \varepsilon) \leq W_{-n}(\hat{F}_n) + V_n(\hat{F}_n, \eta).$$

Continuing this process, we have:

$$NV(F, \varepsilon) = W_{-1}(F) + V_1(F, \varepsilon) \leq W_{-N}(\hat{F}_N) + V_N(\hat{F}_N, \eta) = N V(\hat{F}, \eta).$$

Note that under \hat{F}_n , each player $m \leq n$ effectively has only three signals: t_m^+ , t_m^- , and t_m^* , each with probability at least η . Therefore, Theorem 2 applies to \hat{F} , from which we conclude that

$$V(F, \varepsilon) \leq V(\hat{F}, \eta) \leq R_{\eta, N}.$$

The conclusion of the theorem now follows by taking $\eta \rightarrow \varepsilon$ and noting that $R_{\eta, N}$ is continuous in η .

Proof of the Proposition. To prove part (i), let $A_n = \{\omega : E(\delta | t_n)(\omega) \geq E(\delta | t_n = 0)\}$ denote the set of states at which individual n 's announcement of a type higher than 0 increases the conditional probability of provision. Note that (IR) implies that $E(c_n | t_n = 0) \leq 0$, so (IC) in the special case of reporting $t = 0$ can be rewritten as:

$$E(c_n | t_n)(\omega) \leq t_n(\omega) [E(\delta | t_n)(\omega) - E(\delta | t_n = 0)].$$

This implies $E(c_n | t_n)(\omega) \leq 0$, for $\omega \in A_n^c$, so

$$\begin{aligned} E(c_n) &= \int_{\Omega} E(c_n | t_n)(\omega) dP \leq \int_{A_n} E(c_n | t_n)(\omega) dP \\ &\leq \int_{A_n} t_n [E(\delta | t_n)(\omega) - E(\delta | t_n = 0)] dP. \end{aligned}$$

The assumption that $P(t_n = 0) \geq \varepsilon$, implies that $\sup_{\{A : \lambda(A) < \eta\}} \inf_{t \notin A} E(\delta | t_n) \leq E(\delta | t_n = 0)$. From the definition of influence, there is a sequence of sets $B_k \subset \Omega$, $k = 1, 2, \dots$ such that $P(B_k) < \eta$ and $\sup_{t \notin B_k} E(\delta | t_n)(\omega) - \frac{1}{k} \leq \inf_{\{B : \lambda(B) < \eta\}} \sup_{t \notin B} E(\delta | t_n)$. Thus, for every k , we have

$$V_n(\delta, \eta) \geq E(\delta | t_n)(\omega) - E(\delta | t_n = 0) - \frac{1}{k}, \quad \text{for all } \omega \in B_k^c.$$

These inequalities imply that for every k :

$$\begin{aligned} E(c_n) &\leq \int_{A_n \cap B_k^c} t_n [E(\delta | t_n)(\omega) - E(\delta | t_n = 0)] dP \\ &\quad + \int_{A_n \cap B_k} t_n [E(\delta | t_n)(\omega) - E(\delta | t_n = 0)] dP \\ &\leq t^+ \left[V_n(\delta, \eta) + \frac{1}{k} \right] + t^+ \eta. \end{aligned}$$

Since this is true for any k , we have $E(c_n) \leq t^+ [V_n(\delta, \eta) + \eta]$. From (BB), we have

$$E(\delta) \leq \frac{\sum E c_n}{C_N} \leq \frac{\sum_n t^+ [V_n(\delta, \eta) + \eta]}{\beta N} \leq \frac{t^+}{\beta} [V(\delta, \eta) + \eta] \leq \frac{t^+}{\beta} [R_{\eta, N} + \eta].$$

To prove part (ii), take a sequence of $\eta \rightarrow 0$ and $N = N(\eta)$ such that $R_{\eta, N} \rightarrow 0$.

REFERENCES

1. N. I. Al-Najjar, "A Reputational Model of Authority," MEDS Department, Kellogg GSM, North-western University, 1996.
2. N. I. Al-Najjar and R. Smorodinsky, Large non-anonymous games, *Games Econ. Behav.*, in press.
3. G. Chamberlain and M. Rothschild, A note on the probability of casting a decisive vote, *J. Economic Theory* **25** (1980), 152–162.
4. D. Fudenberg, D. Levine, and W. Pendorfer, When are non-anonymous players negligible? *J. Economic Theory* **79** (1998), 46–71.
5. E. Green, Noncooperative price taking in large dynamic markets, *J. Economic Theory* **22** (1980), 155–182.
6. G. J. Mailath and A. Postlewaite, Asymmetric information bargaining problems with many agents, *Rev. Econ. Stud.* **57** (1990), 351–167.
7. R. Myerson, "Population Uncertainty and Poisson Games," MEDS Department, Kellogg GSM, CMSEMS discussion paper 1102, Northwestern University, 1994.
8. T. Palfrey and H. Rosenthal, Voter participation and strategic uncertainty, *Amer. Polit. Sci. Rev.* **79** (1983), 62–78.
9. A. Postlewaite and D. Schmeidler, Implementation in differential information economies, *J. Econ. Theory* **39** (1986), 14–33.
10. R. Rob, Pollution claim settlements under private information, *J. Econ. Theory* **47** (1989), 307–333.
11. H. Sabourian, Anonymous repeated games with a large number of players and random outcomes, *J. Econ. Theory* **51** (1990), 92–110.
12. I. Segal, "Contracting with Externalities," Berkeley, 1997.
13. J. Swinkels, "Asymptotic Efficiency for Discriminatory Private Value Auctions with Aggregate Uncertainty," John M. Olin School of Business, Washington University, 1994.