

Center for Mathematical Studies in Economics and Management Science
Northwestern University, Evanston, Illinois 60208

Discussion Paper No. 1102R

POPULATION UNCERTAINTY AND POISSON GAMES

by

Roger B. Myerson*

September 1994, revised June 1997

Abstract. A general class of models is developed for analyzing games with population uncertainty. Within this general class, a special class of Poisson games is defined. It is shown that Poisson games are uniquely characterized by properties of independent actions and environmental equivalence. The general definition of equilibrium for games with population uncertainty is formulated, and it is shown that the equilibria of Poisson games are invariant under payoff-irrelevant type splitting. An example of a large voting game is discussed, to illustrate the advantages of using a Poisson game model for large games.

Acknowledgments. I have benefited from many discussions on this topic with John Hillas and Dov Samet. Support from the National Science Foundation grant SES-9308139 and from the Dispute Resolution Research Center is gratefully acknowledged.

*Department of Managerial Economics and Decision Sciences, J. L. Kellogg Graduate School of Management, Northwestern University, Evanston, IL 60208-2009.
E-mail: myerson@nwu.edu

POPULATION UNCERTAINTY AND POISSON GAMES

by Roger B. Myerson

1. Introduction

Most models in game theory begin by specifying a parameter that denotes the set of players in the game. Admitting such a parameter in the model is not as innocuous as it may seem, because game-theoretic methodology generally assumes that everything in the model is common knowledge among the players. So when a model specifies the set of players, it implicitly assumes that this set is common knowledge among all players. Thus every player must know the exact number of players in the game and must be aware of the individual identity of every other player.

In large games, however, it may be more realistic to admit that players have some uncertainty about how many others are in the game. Furthermore, some models with such population uncertainty may actually be easier to analyze than similar models without population uncertainty, particularly when uncertainty about the number of players is assumed to fit a Poisson probability distribution. That is, instead of assuming that beliefs about the number of players can be described by a probability distribution that assigns probability one to a single number, it may be technically simpler and more realistic to assume that these beliefs can be described by a Poisson probability distribution.

The goals of this paper are to develop a general mathematical framework for analysis of games with population uncertainty, and to show the special properties of the Poisson games within the general class of population-uncertainty games. In Section 2, a general model of

games with population uncertainty is formulated. Then in Section 3 Poisson games are defined as a special subclass of games with population uncertainty. Section 4 shows that a property of independent actions is uniquely satisfied by such Poisson games. Section 5 shows that another property of environmental equivalence is similarly satisfied only by Poisson games. The general definition of equilibrium for games with population uncertainty is formulated in Section 6, and it is shown that the equilibria of Poisson games are invariant under payoff-irrelevant type splitting. In Section 7, an example of a large voting game is discussed to illustrate the advantages of using a Poisson game model for large games.

2. Games with population uncertainty

In this section, we develop a general class of models for studying games where it is not appropriate to assume that the set of players is common knowledge. In such games, people may share an awareness of the various possible types of players that might be in the game, but the actual numbers of players of each type are generally uncertain and can only be described probabilistically. Thus, to construct a model of a game with population uncertainty, we begin by specifying some set T which represents the set of possible types of players. We assume throughout this paper that T is a nonempty finite set.

Having population uncertainty means that, for each type t in T , the number of players in the game whose type is t is a random variable. These random variables together form a vector, called the type profile, which lists the number of players in the game who have each type.

For any set S , we let $Z(S)$ denote the set of profiles $(w(s))_{s \in S}$ in \mathbb{R}^S such that, for every s

in S . $w(s)$ is a nonnegative integer. Also, we let $\Delta(S)$ denote the set of all probability distributions on the set S .

With this notation, the set $Z(T)$ denotes the set of possible type profiles in the game. (Notice that $Z(T)$ is a countable set when T is finite.) Then the population uncertainty can be described by a probability distribution Q which is in the set $\Delta(Z(T))$. That is, for each vector y in $Z(T)$, $Q(y)$ denotes the probability that, for every t in T , $y(t)$ is the number of players of type t in the game.

Next, our game model must specify the set of possible actions that each player can choose in the game. To simplify notation, let us assume that the set of possible actions for a player is the same, regardless of his (or her) type, and let C denote this set of possible actions. We assume throughout this paper that C is a finite set containing at least two distinct alternatives.

The utility payoff to each player can depend on the player's type, the player's action, and on the numbers of other players who choose each of the possible actions in C . A vector that lists these numbers of players for each possible action in C is called an action profile and is in the set $Z(C)$. So utility payoffs can be mathematically specified by a utility function of the form $U:Z(C) \times C \times T \rightarrow \mathbb{R}$. Then $U(x,b,t)$ denotes the utility payoff to a player whose type is t and who chooses action b , when x is the profile of actions chosen by the other players in the game (that is when, for each c in C , there are $x(c)$ other players who choose action c , not counting this player in the case of $c = b$). We assume in this paper that $U(\bullet, \bullet, \bullet)$ is a bounded function.

(Having a type-dependent utility function implies that there is no real loss of generality in our assumption that all types choose actions from the same set C . Any game in which the set of feasible actions depended on a player's type could be represented by a model in which C is the

union of the feasible sets for all types. Then the utility function could be specified so that any player who chooses an action that is "infeasible" for his type would get the lowest possible payoff, and so no player of any type would ever choose an infeasible action in equilibrium.)

These formal structures (T, Q, C, U) together define our general model of a game with population uncertainty. For related models of random-player games, see also Milchtaich (1997).

In a traditional game where the set of players is assumed to be common knowledge, we describe players' perceptions of each others' strategic behavior by strategy profiles that predict a distinct randomized strategy for each player in the game. In our games with population uncertainty, however, players' perceptions about each others' strategic behavior cannot be formulated as a strategy profile that assigns a randomized strategy to each specific individual in the game, because a player is not aware of the specific identities of all the other players. Instead, we assume here that each player is only aware of the possible types of the other players, and so players can only form perceptions about how the strategic behavior of other players is likely to depend on their types. It cannot be commonly perceived that two different individuals of the same type would behave differently because, in our model with population uncertainty, two players of the same type have no commonly known attributes by which others can distinguish them.

Thus, in a game with population uncertainty, the conclusion of our game-theoretic analysis will be to specify some strategy function $\sigma: T \rightarrow \Delta(C)$ such that

$$\sigma = (\sigma(c|t))_{c \in C, t \in T}$$

$$\sigma(c|t) \geq 0, \forall c \in C, \forall t \in T$$

$$\sum_{c \in C} \sigma(c|t) = 1, \forall t \in T.$$

Here, for each possible action c and each possible type t , $\sigma(c|t)$ represents that conditional probability that a player will choose action c if his type is t . That is, going to a model of population uncertainty requires us specify a probability distribution over actions for each type of player, rather than for each individual player. In effect, population uncertainty forces us to treat players symmetrically in our game-theoretic analysis.

The model of population uncertainty is actually general enough to subsume Harsanyi's (1967-8) basic model of Bayesian games with consistent priors. For example, consider a consistent Bayesian game in which N is the set of players and, for each player i in N , T_i is the set of possible types for player i . Then we can construct an equivalent game model with population uncertainty in which the set of types T is the disjoint union of all these T_i sets. The probability distribution Q in this population-uncertainty game assigns positive probability only to type profiles such that, for each i in N , there is exactly one player whose type is in the set T_i , and the probabilities of these type profiles are the same as the probabilities in the original Harsanyi Bayesian game.

The additional generality in our games with population uncertainty is found when we allow a positive probability of two or more players having the same type. Then, as we have seen, the population-uncertainty model stipulates a symmetry in others' beliefs about the behavior of these players of the same type. This symmetry requirement was not entailed by Harsanyi's basic Bayesian model.

3. Poisson games

Among the set of all games with population uncertainty, the special case that will interest us the most in this paper is the case of Poisson games. In such Poisson games, the number of players is a Poisson random variable with some given mean, and the players' types are independent identically distributed random variables.

A Poisson random variable with mean λ is a random variable that can equal any nonnegative integer k with probability

$$p(k|\lambda) = e^{-\lambda} \lambda^k / k!$$

(Here $e = 2.71828\dots$ and $k! = 1 \times 2 \times \dots \times k$, where $0! = 1$.) To explain why the Poisson distribution is particularly useful for modeling population uncertainty, we begin by recalling two well-known properties of the Poisson distribution. (See also Haight, 1967, and Johnson and Kotz, 1969.)

First, any sum of independent Poisson random variables is also a Poisson random variable. This fact is called the aggregation property of the Poisson distribution. This aggregation property is proven by noting that, for any k ,

$$e^{-(\lambda+v)} (\lambda+v)^k / k! = \sum_{j=0}^k (e^{-v} v^j / j!) (e^{-\lambda} \lambda^{k-j} / (k-j)!).$$

Second, suppose that the size of some population is a Poisson random variable with mean v , and each member of this population is independently assigned some characteristic in a set S according to some given probability distribution $(\theta(s))_{s \in S}$. If $\tilde{W}(s)$ denotes the number of individuals in this population who have characteristic s , then these random variables $(\tilde{W}(s))_{s \in S}$ are mutually independent, and each $\tilde{W}(s)$ has a Poisson distribution with mean $v\theta(s)$. This fact is called the decomposition property of the Poisson distribution. To show why this decomposition

property is true. let us suppose for simplicity that $S = \{0,1\}$, and so $\theta(0) = 1 - \theta(1)$. Then the probability that j individuals have characteristic 0 and k individuals have characteristic 1 is (from the Poisson and binomial distributions):

$$\begin{aligned} & (e^{-v} v^{j+k} / (j+k)!) ((j+k)! / (j!k!)) \theta(0)^j \theta(1)^k \\ &= (e^{-v\theta(0)} (v\theta(0))^j / j!) (e^{-v\theta(1)} (v\theta(1))^k / k!). \end{aligned}$$

So let us consider a game with population uncertainty in which the total number of players is a Poisson random variable with mean n . Suppose that each player's type is determined independently according to some fixed probability distribution $r = (r(t))_{t \in T}$. These parameters n and r characterize the population uncertainty in a Poisson game. with

$$(3.1) \quad Q(y) = \prod_{t \in T} (e^{-nr(t)} (nr(t))^{y(t)} / y(t)!),$$

for any y in $Z(T)$. Thus we may denote a Poisson game by the parameters (T, n, r, C, U) . (See Myerson, 1994, for a treatment of extended Poisson games in which the parameters n and r may be functions of a random state of the world.)

Given such a Poisson game, let $\tilde{Y}(t)$ denote the random number of players of type t . One application of the Poisson decomposition property implies that each $\tilde{Y}(t)$ is a Poisson random variable with mean $nr(t)$, and these random variables $(\tilde{Y}(t))_{t \in T}$ are mutually independent.

In this Poisson game, suppose now that each player independently chooses an action in the set C according to the strategy function σ , choosing action c with probability $\sigma(c|t)$ if his type is t . A second application of the Poisson decomposition property then implies that numbers of players of each type t who choose each action c are also independent Poisson random variables, each with mean $nr(t)\sigma(c|t)$.

Now let $\tilde{X}(c)$ denote the number of players (of all types) who choose action c . Then the Poisson aggregation property implies that $\tilde{X}(c)$ is a Poisson random variable with mean

$$\sum_{t \in T} nr(t)\sigma(c, t),$$

and these random variables $(\tilde{X}(c))_{c \in C}$ are mutually independent. So in a Poisson game, once we specify the expected number of players choosing each action, the full joint distribution of the action profile $(\tilde{X}(c))_{c \in C}$ is completely determined.

In the next section, we show that this convenient property of having the numbers of players choosing each action be independent random variables can occur only in games with such Poisson structure.

4. Independent actions

Consider again a general game with population uncertainty (T, Q, C, U) , as in Section 2. In a play of this game, let $\tilde{X}(c)$ denote the number of players who choose action c , for each c in C . Let $\tilde{X} = (\tilde{X}(c))_{c \in C}$ denote the action profile of all these numbers. These numbers $\tilde{X}(c)$ are random variables whose probability distribution depends on the given type distribution Q and on the strategy function σ that predicts players' behavior as a function of their types. To be specific, for any vector x in $Z(C)$, the probability that \tilde{X} equals x is

$$P(\tilde{X}=x|\sigma) = \sum_{w \in F(x)} Q(\bar{w}) \prod_{t \in T} \left(\bar{w}(t)! \prod_{c \in C} \left(\frac{\sigma(c, t)^{w(c, t)}}{w(c, t)!} \right) \right)$$

where $F(x) = \{w \in Z(C \times T) \mid \sum_{t \in T} w(c, t) = x(c), \forall c \in C\}$,

$$\bar{w}(t) = \sum_{c \in C} w(c, t), \text{ and } \bar{w} = (\bar{w}(t))_{t \in T}.$$

The game (T, Q, C, U) has the independent-actions property iff. for every strategy function σ , the random variables $\tilde{X}(c)$ for all c in C are independent random variables. Without population uncertainty, the independent-actions property could not be satisfied because the sum of the random variables $\tilde{X}(c)$ would have to equal the known number of players, and so any one component of the action profile \tilde{X} could be computed from the others. But the independent-actions property can be satisfied for one special class of games with population uncertainty: the games in which the number of players is a random variable drawn from a Poisson distribution. (This theorem is a game-theoretic version of a result of Daboni, 1959.)

Theorem 1. Suppose that the game (T, Q, C, U) satisfies the independent-actions property. Then for any strategy function σ , each $\tilde{X}(c)$ must be a Poisson random variable with mean $\sum_{t \in T} \sigma(c|t) \sum_{y \in Z(T)} Q(y) y(t)$, and the total number of players in the game must be a Poisson random variable.

Proof. Let σ be a strategy function, and let a and b be any two actions in C . Let σ^* be a strategy function such that, for every type t :

$$\sigma^*(b|t) = .5\sigma(b|t), \quad \sigma^*(a|t) = \sigma(a|t) + .5\sigma(b|t), \quad \text{and all other } \sigma^*(c|t) = \sigma(c|t).$$

That is, let σ^* be derived from σ by specifying that, whenever a player would have chosen action b under σ , the player instead randomizes between actions a and b with equal probability.

For any k , consider the ratio of the probability of having one player choose action a and k players choose action b under strategy function σ^* over the probability of having zero players choose action a and k players choose action b under strategy function σ^* ; that is:

$$P(\tilde{X}(a)=1 \cap \tilde{X}(b)=k | \sigma^*) / P(\tilde{X}(a)=0 \cap \tilde{X}(b)=k | \sigma^*).$$

By the independent-actions property for σ^* , this ratio must be independent of k , because learning the number of players who choose b would not change the conditional probability of having 1 or 0 players who choose action a . But by the way that σ^* is derived from σ (and using the fact that independent actions also applies for σ), this ratio is

$$\begin{aligned} & \frac{P(\tilde{X}(a)=1 | \sigma) P(\tilde{X}(b)=k | \sigma) .5^k + P(\tilde{X}(a)=0 | \sigma) P(\tilde{X}(b)=k+1 | \sigma) (k+1) .5^{k-1}}{P(\tilde{X}(a)=0 | \sigma) P(\tilde{X}(b)=k | \sigma) .5^k} \\ & = P(\tilde{X}(a)=1 | \sigma) / P(\tilde{X}(a)=0 | \sigma) + .5 P(\tilde{X}(b)=k+1 | \sigma) (k+1) / P(\tilde{X}(b)=k | \sigma) \end{aligned}$$

Thus, the quantity

$$P(\tilde{X}(b)=k+1 | \sigma) (k+1) / P(\tilde{X}(b)=k | \sigma)$$

must be independent of k . Call this quantity $\lambda(b)$. So for all k ,

$$P(\tilde{X}(b)=k+1 | \sigma) = P(\tilde{X}(b)=k | \sigma) \lambda(b) / (k+1).$$

Then by induction,

$$P(\tilde{X}(b)=k | \sigma) = P(\tilde{X}(b)=0 | \sigma) \lambda(b)^k / k!. \quad \forall k \in \{0, 1, 2, \dots\}.$$

To make these probabilities sum to 1, we must have

$$P(\tilde{X}(b)=0 | \sigma) = e^{-\lambda(b)}$$

and so $\tilde{X}(b)$ is a Poisson random variable with mean $\lambda(b)$.

This mean $\lambda(b)$ must equal the expected number of players who choose the action b , which is $\sum_{t \in T} \sigma(b|t) \sum_{y \in Z(T)} Q(y) y(t)$, because $\sum_{y \in Z(T)} Q(y) y(t)$ is the expected number of players of type t , and $\sigma(b|t)$ is the probability that any player of type t will choose action b .

As recalled in Section 3, the sum of independent Poisson random variables is also a Poisson random variable. Thus, the total number of players in the game, which equals

$\sum_{c \in C} \tilde{X}(c)$, must be a Poisson random variable.

Q.E.D.

5. A player's environment in a game with population uncertainty

Consider again a general game with population uncertainty (T, Q, C, U), as defined in Section 2. This game model describes what a game theorist who is not in the game should think about the structure of the game. We assume that each player understands this game-theoretic model and knows his (or her) own type, and so also knows whatever can be inferred about the game from the fact of that he (or she) is a player of this type in the game.

When a player assesses what he knows about the game, he actually assesses a model of his environment, where a player's environment is defined to include everyone in the game except himself. In a game with no population uncertainty, for example, if there are 450 players in the game, then each player's environment includes just 449 players.

In contrast, consider a game with population uncertainty where the number of players may be either 300 or 600, each with probability 1/2. A game-theorist looking at the game from the outside would say that the expected number of players is 450. But suppose now that you have just learned that you are a player in this game. Assuming that your likelihood of being recruited into this game was no different from any other possible player, you would be twice as likely to be in the game if the number of players was 600 than if the number of players was 300. So learning that you are in the game should cause you to update the conditional probability of 600 players to 2/3. That is, each player in the game should infer that the probability of having 599 other players in his environment is 2/3, and the probability of having 299 other players is

1/3. So the expected number of other players in any player's environment is $599 \times 2/3 + 299 \times 1/3 = 499$, which is greater than 450.

Thus a player might see his environment as smaller than an external game-theorist's perception of the whole game, because the player's environment excludes himself. But a player might also see his environment as larger than an external game-theorist's perception of the whole game, because the player may take his participation in the game as evidence in favor of a larger population of players. We will show that these two effects exactly cancel out only in the case of Poisson games.

For any possible type-profile y in $Z(T)$, let $|y|$ denote the total number in the type profile y ; that is,

$$|y| = \sum_{t \in T} y(t).$$

Given our general game with population uncertainty (T, Q, C, U) , let n denote the expected number of players

$$n = \sum_{y \in Z(T)} Q(y) \cdot |y|.$$

Let us assume that this expected value n is a finite positive number. Then for each t in T , let $r(t)$ be defined so that $nr(t)$ is the expected number of players of type t ; that is

$$r(t) = \sum_{y \in Z(T)} Q(y) y(t)/n.$$

Given any type profile y in $Z(T)$ and any type t in T , let $y+[t]$ denote the type profile that differs from w only in that the number of players of type t is $y(t)+1$; that is,

$$(y+[t])(s) = y(s) \text{ if } s \neq t, \text{ and } (y+[t])(t) = y(t)+1.$$

Consider a player in the game who knows that his type is t . Let M be a large positive integer; and let us temporarily perturb the story behind our population-uncertainty game by

supposing that, before the game began, there was a set of M candidates who were to be recruited into the game before any other players. Then for any y in $Z(T)$ and any t in T , any one of these M candidates should figure that the prior probability of his being recruited as a type- t player and having y be the type profile of the other players in his environment is

$$\begin{aligned} Q(y+[t]) \left(\frac{\min\{|y|+1, M\}}{M} \right) \left(\frac{y(t)+1}{|y|+1} \right) \\ = Q(y+[t]) (y(t)+1) \min\{1, M/(|y|+1)\} / M \end{aligned}$$

Such a candidate's overall probability of his being recruited as a player of type t is then

$$\begin{aligned} \sum_{y \in Z(T)} Q(y+[t]) (y(t)+1) \min\{1, M/(|y|+1)\} / M. \\ = \sum_{z \in Z(T)} Q(z) z(t) \min\{1, M/|z|\} / M. \end{aligned}$$

So when such a candidate has been recruited as a type- t player, the conditional probability of having type profile y for the other players in his environment is

$$\frac{Q(y+[t])(y(t)+1) \min\{1, M/(|y|+1)\}}{\sum_{z \in Z(T)} Q(z) z(t) \min\{1, M/|z|\}}$$

In our analysis of games with population uncertainty, we actually do not want to assume that the beliefs of some players in the games are influenced by the memory of having been in some prior pool of candidates. We want to compute players' beliefs that depend only on the fact of being in the game. Thus, in the above conditional-probability formula, we must eliminate the bias that is due to the fictitious construction of a prior pool of M candidates. But taking the limit as M goes to infinity will eliminate this informational distinction between simply being a player and being a player who was in the prior pool of M candidates, because the probability of every player having been a "candidate" will go to one as M goes to infinity. For any type profile y in $Z(T)$, the above conditional probability converges, as M goes to infinity, to the limit $q(y|t)$ such

that

$$q(y|t) = Q(y+[t]) (y(t)+1)/(nr(t)).$$

So this limit $q(y|t)$ represents the conditional probability that a player of type t should assess for the event that the other players in his environment have type profile y .

The game (T, Q, C, U) has environmental equivalence iff a player of any type would assess the same probability distribution for the type profile of the others in his environment as the external game theorist would assess for the type profile of the whole game. That is, there is environmental equivalence iff

$$q(y|t) = Q(y), \quad \forall y \in Z(T), \quad \forall t \in T.$$

This condition is equivalent to

$$Q(y) nr(t)/(y(t)+1) = Q(y+[t]), \quad \forall y \in Z(T), \quad \forall t \in T.$$

By induction, this system of equations is satisfied if and only if

$$Q(y) = Q(\bar{0}) \prod_{t \in T} ((nr(t))^{y(t)}/y(t)!), \quad \forall y \in Z(T).$$

Then to satisfy the condition $\sum_{y \in Z(T)} Q(y) = 1$, we must have

$$Q(\bar{0}) = \prod_{t \in T} (e^{-nr(t)}) = e^{-n},$$

which gives us the basic equation for a Poisson game (3.1). Thus we have proven the following theorem.

Theorem 2. A game with population uncertainty satisfies environmental equivalence if and only if it is a Poisson game.

6. Equilibrium

With the formulation of the type-conditional probability distribution $q(\bullet|\bullet)$ from the preceding section, we can now formulate the basic definition of equilibrium for games with population uncertainty.

For any type s in T , for any action b in C , and for any strategy function $\sigma:T \rightarrow \Delta(C)$, the expected utility for a player of type s who chooses action b when all other players are expected to behave according to the strategy function σ is

$$\bar{U}(b|s,\sigma) = \sum_{w \in Z(C \times T)} q(\bar{w}|s) \prod_{t \in T} \left(\bar{w}(t)! \prod_{c \in C} \left(\frac{\sigma(c|t)^{w(c,t)}}{w(c,t)!} \right) \right) U(\hat{w}, b, s)$$

where $\bar{w}(t) = \sum_{c \in C} w(c,t)$, $\bar{w} = (\bar{w}(t))_{t \in T}$, $\hat{w}(c) = \sum_{t \in T} w(c,t)$, $\hat{w} = (\hat{w}(c))_{c \in C}$.

This definition can be quite difficult to apply in practice, because it involves summing over all possible profiles in $Z(C \times T)$, each of which lists how many players of every type are choosing every action. For each such profile w , the probability of the corresponding type profile \bar{w} is determined by the type-contingent probability distribution $q(\bar{w}|s)$; and then the probability that the other players whose type is t (for each t) behave as specified by $(w(c,t))_{c \in C}$ is determined by a multinomial distribution with parameters $\bar{w}(t)$ and $(\sigma(c|t))_{c \in C}$. These probabilities are multiplied by the corresponding utility payoff and summed in the above formula. When the given utility function U is bounded, this function $\bar{U}(b|s,\sigma)$ is a continuous function of σ .

An equilibrium is a strategy function such that every type assigns positive probability only to actions that maximize the expected utility for players of this type: that is, for every action c in C and for every type t in T ,

$$\text{if } \sigma(c|t) > 0 \text{ then } \bar{U}(c|t,\sigma) = \max_{b \in C} \bar{U}(b|t,\sigma).$$

Under the assumption that C and T are finite sets and $U:Z(C) \times C \times T \rightarrow \mathbb{R}$ is a bounded function, the existence of such equilibria for our game with population uncertainty is a straightforward consequence of the Kakutani fixed-point theorem applied to the compact set of all strategy functions $\Delta(C)^T$. Thus, we can assert the following theorem.

Theorem 3. Any game with population uncertainty (T, Q, C, U) as above (where T and C are finite and U is bounded) must have at least one equilibrium.

In the special case where our population-uncertainty game is a Poisson game, the expected utility formula can be simplified to the following:

$$\bar{U}(b, s, \sigma) = \sum_{x \in Z(C)} \prod_{c \in C} \left(\frac{e^{-n\tau(c)} (n\tau(c))^{x(c)}}{x(c)!} \right) U(x, b, s)$$

where $\tau(c) = \sum_{t \in T} r(t) \sigma(c|t)$.

This formula uses the fact that (from Section 3), in a Poisson game, the number of players choosing each action c is an independent Poisson random variable with mean $\sum_{t \in T} nr(t)\sigma(c|t)$, which is rewritten as $n\tau(c)$ in the above formula. Here $\tau(c)$ is the marginal probability that a randomly sampled player will choose action c when σ describes how each player's action is expected to depend on his type. The environmental equivalence property of Poisson games is also used here to imply that, from the perspective of any player of any type, the number of other players (not including himself) who choose each action c is also an independent Poisson random variable with this same mean $n\tau(c)$.

We have emphasized that, in any game with population uncertainty, all players of the

same type must have the same predicted behavior in equilibrium, because the type of a player includes all the behaviorally-relevant attributes of the player that are recognized by others. This restriction begs the question of how equilibria might change if we subdivided types more finely in our model. For Poisson games, the answer is that utility-irrelevant subdivisions of types cannot substantively change the set of equilibria.

To see why, consider a Poisson game $\Gamma = (T, n, r, C, U)$. Let s be some type in T , and let $\Gamma' = (T', n, r, C, U)$ be a new Poisson game that differs from Γ only in that this type s has replaced by two types s_1 and s_2 that have probabilities summing to the probability of s and have the same utility as type s ; that is,

$$T' = (T \setminus \{s\}) \cup \{s_1, s_2\},$$

$$r(s_1) + r(s_2) = r(s),$$

$$U(x, c, s_1) = U(x, c, s_2) = U(x, c, s), \quad \forall x \in Z(C), \forall c \in C.$$

Let σ' be any equilibrium of the new game Γ' . Then we can create an equivalent strategy function σ for the original game Γ by

$$\sigma(c|s) = (r(s_1)\sigma'(c|s_1) + r(s_2)\sigma'(c|s_2))/r(s), \quad \forall c \in C,$$

$$\sigma(c|t) = \sigma'(c|t), \quad \forall c \in C, \forall t \in T \setminus \{s\}.$$

The marginal probability distribution on C generated by these two strategy functions is exactly the same: that is

$$\tau(c) = \sum_{t \in T} r(t)\sigma(c|t) = \sum_{t \in T'} r(t)\sigma'(c|t), \quad \forall c \in C.$$

Because the optimality criteria depend only on these marginal probabilities on C and because the utility functions are the same in Γ and Γ' , σ is an equilibrium of Γ .

Conversely it is easy to see that, for any equilibrium σ of the given Poisson game Γ , we

can create an equivalent equilibrium σ' of the game Γ' in which the new types s_1 and s_2 use the same randomized strategy as the old type s in Γ . That is, if

$$\sigma'(c|s_1) = \sigma'(c|s_2) = \sigma(c|s), \quad \forall c \in C,$$

$$\sigma'(c|t) = \sigma(c|t), \quad \forall c \in C, \forall t \in T \setminus \{s\}$$

then we again get

$$\tau(c) = \sum_{t \in T} r(t) \sigma(c|t) = \sum_{t \in T} r(t) \sigma'(c|t), \quad \forall c \in C.$$

Thus we get the following theorem on invariance of equilibria under type-splitting.

Theorem 4. If the Poisson game Γ' is derived from the Poisson game Γ by splitting a type into two utility-equivalent types that have the same total probability, then the set of marginal distributions on the action set C that are generated by equilibria is the same in Γ and Γ' .

This theorem can be interpreted as saying that, as long as we stay within the class of Poisson games, there is essentially no loss of generality in assuming that all players who have the same utility function are expected to independently implement the same randomized strategy. Admitting that players have other payoff-irrelevant characteristics which might affect their behavior does not essentially increase the set of equilibria, as long as the beliefs about the numbers of players who have each of these characteristics still fit the Poisson model.

7. An example

In this section we consider an example of a large voting game, to illustrate the technical advantages of using a Poisson model of population uncertainty. Unfortunately, illustrating these advantages is a rather unpleasant task, because it requires that I first show how complicated the formulas get without the Poisson approach before I can show that they become simpler with a Poisson model. (For more applications of Poisson games, see Engelbrecht-Wiggans, 1987, and Myerson, 1994, 1997.)

We consider here a voting game that was studied by Palfrey and Rosenthal (1983,1985). The players in this game are the eligible voters in an election, and we expect to have millions of such eligible voters. Suppose that each voter has a type that is rightist or leftist. There are two candidates in this election, one candidate for the right and one candidate for the left. The winner will be the candidate who gets more votes, and in the case of a tie the winner will be determined by a fair coin toss that gives each candidate a $1/2$ probability of winning.

Let us suppose that each voter's utility payoff is the sum of two terms. The first term in the utility function equals +1 if the voter's preferred candidate wins the election, and equals 0 otherwise. The second term in the utility function equals -0.05 if the voter votes in the election, and equals 0 if the voter abstains. With this utility function, a voter wants to vote for his preferred candidate if the probability of his vote making a difference is more than .05, but the voter would prefer to abstain if the probability of his vote making a difference is less than .05. The probability that a voter's vote would change the outcome of the election is called the voter's pivot probability.

One may pose the intuitive conjecture that, in such a voting game where the act of voting

is costly for every voter, expected turnout in equilibrium cannot be large. Following Palfrey and Rosenthal (1983, 1985), let us analyze this example with the goal of evaluating this conjecture.

It would be very easy to construct an equilibrium in which everyone votes if there were exactly 1,000,000 rightists and exactly 1,000,000 leftists in the electorate. In such an equilibrium, every voter's pivot probability would be $1/2$ (the probability that the preferred candidate will win the tie that his vote creates), and so every voter would want to vote (because $1/2 > .05$). But Palfrey and Rosenthal (1983) showed that such counterexamples to the low-turnout conjecture can also be found in more general games, as long as we stay in the case where there is no population uncertainty and the types of all voters are common knowledge.

So let us consider a first version of this game in which there are 1,000,000 rightist voters and 2,000,000 leftist voters in the electorate. If everyone was expected to vote then no one would have any chance of changing the outcome of the election, and so no one would want to pay the cost of voting. But if no one was expected to vote then everyone's pivot probability would be $1/2$ (the probability that the preferred candidate would lose the coin toss in the case of a 0-0 tie), and so everyone would want to vote. So to find equilibria of this game, we must look for equilibria in which some players randomize, which they will only do if their pivot probabilities are exactly 0.05.

To construct an equilibrium that violates the low-turnout conjecture, we can begin by partitioning the 2,000,000 leftists into two subsets of 1,000,000 each. All leftists in the first set of 1,000,000 are expected to vote for the candidate on the left with probability one in this equilibrium. The other 1,000,000 leftists are expected to abstain for sure in this equilibrium. Each of the 1,000,000 rightists is expected to use a randomized strategy, voting for

the candidate on the right with probability 0.9999977, and abstaining with probability 0.0000023. In this equilibrium, the number of abstaining rightists is a random variable that has a binomial distribution with parameters $n=1,000,000$ and $p = 0.0000023$. Such a binomial distribution with large n and small p can also be closely approximated by a Poisson distribution with mean $np=2.3$.

To verify that this scenario is an equilibrium, notice first that each rightist voter thinks that his pivot probability is $1/2$ times the probability that no other rightist abstains (that is, the probability that his vote will make a tie which the right candidate will win), and this probability is

$$(1/2) \times .9999977^{999999} = .05$$

Thus the rightists are willing to randomize.

Each leftist who is expected to vote figures that he could be pivotal two ways: if all rightists vote and the left candidate would win a coin toss, and if all but one rightists vote and the left candidate would lose a coin toss. So any of the leftists who are expected to vote has a pivot probability

$$(1/2) \times .9999977^{1000000} + (1/2) \times 1000000 \times .9999977^{999999} \times .0000023 = .115.$$

and so each of these 1,000,000 leftists wants to vote (because $.115 > .05$).

Each leftist who is expected to abstain figures that he would be pivotal only if all rightists vote and the right candidate would win the coin toss, which has probability

$$(1/2) \times .9999977^{1000000} = .04999985$$

This pivot probability is slightly less than $.05$, and so the leftists who are supposed to not vote actually prefer to not vote.

There may be something intuitively unreasonable about this large-turnout equilibrium.

however. In particular, it may be surprising that the pivot probabilities of the leftists who are supposed to vote (.115) are so much larger than the pivot probabilities of the otherwise identical leftists who are not supposed to vote (.04999985). The difference in these pivot probabilities ultimately depends on the fact that each leftist who is supposed to vote perceives in his environment that there are 999,999 other leftists who are supposed to vote, whereas each leftist who is not supposed to vote perceives in his environment 1,000,000 other leftists who are supposed to vote. Thus, the surprising difference in the pivot probabilities relies crucially on the assumption that everyone knows exactly how many leftists there are and which of them are expected to vote. Adding uncertainty to this game will make it impossible to so cleanly divide the leftists into different groups with such different incentives, as Palfrey and Rosenthal (1985) have shown.

This game has other equilibria, including one small-turnout equilibrium which (we will show) is more robust to the addition of uncertainty. This equilibrium is symmetric, in the sense that each voter uses the same randomized strategy as all other voters of the same type. In this symmetric equilibrium, each of the 1,000,000 rightists votes with a probability $\rho \approx .000032$, and each of the 2,000,000 leftists votes with a probability $\lambda \approx .000016$. These probabilities are determined by the need to make each voter's pivot probability equal to .05.

The pivot probability for a rightist voter is the sum of two probabilities: the probability that without his vote there would be a tie that the left candidate would win, plus the probability that without his vote the left candidate would be one vote ahead but the right candidate would win a coin toss. Without this one rightist's vote, the number of other rightist votes has a binomial distribution with parameters $n=999,999$ and $p=\rho$, while the number of leftist votes has a binomial

distribution with parameters $n=2,000,000$ and $p=\lambda$. So a rightist's pivot probability in this symmetric scenario is

$$(7.1) \quad \sum_{k=0}^{999999} \frac{999999!}{k!(999999-k)!} \rho^k (1-\rho)^{999999-k} \frac{2000000!}{k!(2000000-k)!} \lambda^k (1-\lambda)^{2000000-k} \times \left(1 + \frac{(2000000-k)\lambda}{(k+1)(1-\lambda)} \right) \left(\frac{1}{2} \right)$$

Similarly, the pivot probability for a leftist voter is the sum of two probabilities: the probability that without his vote there would be a tie that the right candidate would win, plus the probability that without his vote the right candidate would be one vote ahead but the left candidate would win a coin toss. Without this leftist's vote, the number of other leftist votes has a binomial distribution with parameters $n=1,999,999$ and $p=\lambda$ while the number of rightist votes has a binomial distribution with parameters $n=1,000,000$ and $p=\rho$. So a leftist's pivot probability in this scenario is

$$(7.2) \quad \sum_{k=0}^{1000000} \frac{1999999!}{k!(1999999-k)!} \lambda^k (1-\lambda)^{1999999-k} \frac{1000000!}{k!(1000000-k)!} \rho^k (1-\rho)^{1000000-k} \times \left(1 + \frac{(1000000-k)\rho}{(k+1)(1-\rho)} \right) \left(\frac{1}{2} \right)$$

Finding for the parameters ρ and λ that make these two pivot probabilities equal to .05 might seem to be a difficult task, because of the complexity of the formulas. But these formulas can be simplified by using the fact that a binomial distribution with large n and small p parameters can be closely approximated by a Poisson distribution with mean np . So in this scenario, the numbers of votes for the two candidates are independent and are approximately Poisson random variables, with means 1000000ρ for the right candidate and 2000000λ for the

left candidate. Then the probability that an additional rightist vote would change the outcome (by making a tie where the right candidate would win the coin toss or by breaking a tie where the right candidate would lose the coin toss) is well approximated by the simpler formula

$$(7.3) \quad \sum_{k=0}^{\infty} \left(\frac{e^{-1000000\rho} (1000000\rho)^k}{k!} \right) \left(\frac{e^{-2000000\lambda} (2000000\lambda)^k}{k!} \right) \left(1 + \frac{2000000\lambda}{(k+1)} \right) (1/2)$$

Similarly, the probability that an additional leftist vote would change the outcome (by making a tie where the left candidate would win the coin toss or by breaking a tie where the left candidate would lose the coin toss) well approximated by the simpler formula

$$(7.4) \quad \sum_{k=0}^{\infty} \left(\frac{e^{-1000000\rho} (1000000\rho)^k}{k!} \right) \left(\frac{e^{-2000000\lambda} (2000000\lambda)^k}{k!} \right) \left(1 + \frac{1000000\rho}{(k+1)} \right) (1/2)$$

When the means 1000000ρ and 2000000λ are relatively small (say less than 50), these formulas (7.3) and (7.4) can be easily estimated numerically (summing the first hundred terms is enough) to show that these pivot probabilities equal .05 when the means are

$$1000000\rho = 2000000\lambda \approx 31.7,$$

which gives us the voting probabilities $\rho \approx .000032$ and $\lambda \approx .000016$ as described above.

Palfrey and Rosenthal (1985) showed that the unrealistic large-turnout equilibrium of this game disappears when we change the model to one in which there is uncertainty about the number of leftists and rightists, and only the low-turnout equilibrium remains. As in Ledyard (1984), however, Palfrey and Rosenthal's model of uncertainty retained the assumption that the total number of players was known. So to follow Palfrey and Rosenthal (1985), we would revise our example by continuing to assume that there were exactly 3,000,000 voters in the electorate, but now we would assume instead that each player's type was an independent random variable, with probability 1/3 of being rightist, and probability 2/3 of being leftist. In this model, if each

rightist votes with probability ρ and each leftist votes with probability λ then, in any voter's environment, the number of rightist votes, the number of leftist votes, and the number of abstentions, together have a joint multinomial distribution with parameters $n = 2999999$, $p_1 = \rho/3$, $p_2 = 2\lambda/3$, and $p_3 = (1 - \rho/3 - 2\lambda/3)$. So the pivot probability for a rightist voter (half the probability of a tie plus half the probability of the left candidate being one vote ahead) is

$$(7.5) \quad \sum_{k=0}^{1499999} \left(\frac{2999999!}{k! k! (2999999 - 2k)!} \right) \left(\frac{\rho}{3} \right)^k \left(\frac{2\lambda}{3} \right)^k \left(1 - \frac{\rho}{3} - \frac{2\lambda}{3} \right)^{2999999 - 2k} \\ \times \left(1 + \frac{(2999999 - 2k) 2\lambda/3}{(k+1)(1 - \rho/3 - 2\lambda/3)} \right) \left(\frac{1}{2} \right).$$

Similarly, the pivot probability of a leftist voter in this Palfrey-Rosenthal (1985) model is

$$(7.6) \quad \sum_{k=0}^{1499999} \left(\frac{2999999!}{k! k! (2999999 - 2k)!} \right) \left(\frac{\rho}{3} \right)^k \left(\frac{2\lambda}{3} \right)^k \left(1 - \frac{\rho}{3} - \frac{2\lambda}{3} \right)^{2999999 - 2k} \\ \times \left(1 + \frac{(2999999 - 2k) \rho/3}{(k+1)(1 - \rho/3 - 2\lambda/3)} \right) \left(\frac{1}{2} \right)$$

When ρ and λ are small, then these formulas (7.5) and (7.6) can also be closely approximated by the simpler Poisson formulas (7.3) and (7.4) respectively, and so we find again an equilibrium in which the expected turnout for each candidate is approximately 32.

In our Poisson version of this game, however, we drop the assumption that the number of voters is known, and instead we assume that the number of rightist voters and the number of leftist voters are independent Poisson random variables with means 1,000,000 and 2,000,000 respectively. (To get a sense of the degree of uncertainty that this assumption entails, notice that the standard deviations of these random variables are only 1000 and 1414, because the standard deviation of a Poisson random variable is the square root of its mean.) In this Poisson model, when each rightist votes with probability ρ and each leftist votes with probability λ , the pivot

probabilities for rightist and leftist voters can be computed exactly from the formulas (7.3) and (7.4) respectively. That is, the Poisson formulas which were only simplifying approximations in the two previous models are now exact formulas. Thus, the Poisson version of this game has a small-turnout equilibrium like the previous models, in which each rightist votes with probability $\rho \approx .000032$, each leftist votes with probability $\lambda \approx .000016$, and expected turnout for each candidate is close to 32.

As we search for other equilibria, we quickly find another major simplification arising from the population-uncertainty approach, because it directly eliminates the nonsymmetric large-turnout equilibria of the complete-information game. Our population-uncertainty models do not allow any such nonsymmetric equilibria in which different players of the same type are expected to behave differently. Thus, every rightist must be expected to vote with the same probability, which we have denoted by ρ , and every leftist must be expected to vote with some other common probability, which we have denoted by λ . Furthermore, by the type-splitting theorem in Section 6, making some payoff-irrelevant subdivision of the leftists and rightists into subtypes that could be perceived differently would not increase the set of equilibria, as long as there was still Poisson uncertainty about the numbers of these subtypes.

It is easy to verify that there is no equilibrium in which the leftist or rights all use a pure strategy (either all voting for sure or all abstaining for sure). Thus, the only question is whether the Poisson pivot probabilities (7.3) and (7.4) can be both set equal to .05 for some other values of ρ and λ , besides the pair that we have already found. The answer is that this small-turnout solution is unique. To prove uniqueness, we can rule out the possibility of large-turnout equilibria by using an approximation formula for the Poisson pivot probabilities that has been

derived in a paper on large Poisson games (Myerson, 1997). This result asserts that, if the expected vote total for candidate 1 and 2 are independent Poisson random variables with means α and β respectively, then the pivot probability of an additional vote for candidate 1 satisfies the approximate equality:

$$(7.7) \quad \sum_{k=0}^{\infty} \left(\frac{e^{-\alpha} \alpha^k}{k!} \right) \left(\frac{e^{-\beta} \beta^k}{k!} \right) \left(1 + \frac{\beta}{(k+1)} \right) (1/2) \approx \frac{e^{-(\alpha+\beta-2\sqrt{\alpha\beta})}}{4\sqrt{\pi\sqrt{\alpha\beta}}} \left(\frac{\sqrt{\alpha} + \sqrt{\beta}}{\sqrt{\alpha}} \right)$$

The ratio of the two sides of this approximate equality is close to one, as long as α and β are not small. In fact, errors in this approximation (7.7) are less than 1% when α and β are larger than 7.

In our example, the expected turnouts for the two candidates are 1000000ρ and 2000000λ . So by (7.7), a large-turnout equilibrium would have to satisfy the two equations

$$\frac{e^{-1000000(2\lambda-\rho-2\sqrt{2\lambda\rho})}}{4\sqrt{\pi 1000000\sqrt{2\lambda\rho}}} \left(\frac{\sqrt{2\lambda} + \sqrt{\rho}}{\sqrt{\rho}} \right) = .05$$

$$\frac{e^{-1000000(2\lambda-\rho-2\sqrt{2\lambda\rho})}}{4\sqrt{\pi 1000000\sqrt{2\lambda\rho}}} \left(\frac{\sqrt{2\lambda} + \sqrt{\rho}}{\sqrt{2\lambda}} \right) = .05$$

These two equations can be satisfied only if $\rho = 2\lambda$. So we are left with just one unknown, and it is then easy to verify numerically that these pivot formulas cannot be both close to .05 unless ρ and 2λ are approximately .000032. That is, the only equilibrium of this Poisson game with costly voting is the small-turnout equilibrium in which each candidate's expected vote total is close to 32. Thus, the perverse large-turnout equilibrium of the complete-information model has been straightforwardly eliminated in the Poisson model, and the analysis of the Poisson model supports our original conjecture that expected turnout cannot be large if the act of voting is costly for all voters.

REFERENCES

- L. Daboni, "A property of the Poisson distribution," Bolletino della Unione Matematica Italiana 14 (1959), 318-320.
- R. Engelbrecht-Wiggans, "On optimal reservation prices in auctions," Management Science 33 (1987), 763-770.
- J. Harsanyi, "Games with incomplete information played by 'Bayesian' players" Management Science 14 (1967-8), 159-182, 320-334, 486-502.
- F. A. Haight, Handbook of the Poisson Distribution, New York: Wiley (1967).
- N. L. Johnson and S. Kotz, Discrete Distributions, New York: Wiley (1969).
- J. Ledyard, "The pure theory of large two-candidate elections," Public Choice 44 (1984), 7-41.
- I. Milchtaich, "Random-player games." Northwestern University discussion paper (1997).
- R. Myerson, "Extended Poisson games and the Condorcet jury theorem." Northwestern University discussion paper (1994). To appear in Games and Economic Behavior.
- R. Myerson, "Large Poisson games." Northwestern University discussion paper (1997).
- T. R. Palfrey and H. Rosenthal, "A strategic calculus of voting," Public Choice 41 (1983).
- T. R. Palfrey and H. Rosenthal, "Voter participation and strategic uncertainty," American Political Science Review 79 (1985).