The authors analyze two-player repeated games with imperfect monitoring. Each player $i = 1, 2$ has a finite action set A_i and a finite set of signals Σ_i. An action profile is an element of $A = A_1 \times A_2$. For each possible action profile $a \in A$, the monitoring distribution $m(\cdot \mid a)$ specifies a joint probability distribution over the set of signal profiles $\Sigma = \Sigma_1 \times \Sigma_2$. When action profile a is played and signal profile σ is realized, player i observes his corresponding signal σ_i. Stage payoffs can be expressed as a function of mixed action profile alone:

$$u_i(\alpha) = \sum_{a \in A} \sum_{\sigma_i \in \Sigma_i} \alpha(a) m_i(\sigma_i \mid a) \tilde{u}_i(a_i, \sigma_i),$$

where $\tilde{u}_i(a_i, \sigma_i)$ denotes the payoff to player i from action a_i and signal σ_i. Repeated game payoffs are evaluated using the discounted average criterion. The players share a common discount factor $\delta < 1$.

A t-length (private) history for player i is an element of $H_i^t = (A_i \times \Sigma_i)^t$. A pair of t-length histories, called a history, is denoted h^t. The set $H_i = \cup_t H_i^t$ denotes the set of private histories for i. A repeated-game strategy for player i is a mapping $s_i : H_i \to \Delta A_i$, where ΔA_i represents the set of mixed actions for player i. For history h_i^t, let $s \mid_{h_i^t}$ denote the continuation strategy derived from s following history h_i^t. Specifically, if $h_i \hat{h}_i$ denotes the concatenation of the two histories h_i and $\hat{h}_i$, then $s \mid_{h_i}$ is the strategy defined by $s \mid_{h_i} (\hat{h}_i) = s(h_i \hat{h}_i)$. Given a strategy profile s, for each t and $h_{-i}^t \in H_{-i}^t$ let $B_i(s \mid h_{-i}^t)$ denote the set of continuation strategies for i that are best replies to $s_{-i} \mid_{h_{-i}^t}$. A strategy profile s is belief-free if for every h^t, $s_i \mid_{h_i^t} \in B_i(s \mid h_{-i}^t)$ for $i = 1, 2$. It can be shown that belief-free profiles are sequential equilibria. Belief-free equilibria exist, since any history-independent sequence of static equilibrium action profiles is belief-free.

Suppose s is an equilibrium belief-free strategy profile. It is convenient to describe belief-free equilibria in terms of optimal actions in a given period t. Let

$$\mathcal{A}_i^t = \{a_i \in A_i : \exists z_i \in B_i^t(s),\ \exists h_i^t \text{ such that } z_i(h_i^t)[a_i] > 0\}.$$

The set $\mathcal{A}^t = \mathcal{A}_i^t \times \mathcal{A}_{-i}^t$ is referred to as the regime that prevails at date t. Every belief-free equilibrium gives rise to a sequence of non-empty regimes $\{\mathcal{A}^t\}$.

Characterizing belief-free equilibrium payoffs is considerably simplified when the players have access to a public randomization device. It is therefore assumed that in each period, all players observe a public signal y from a set of possible public signal realizations Y. A public signal is an i.i.d. draw from the same distribution $p \in \Delta Y$. A strategy now depends on private history as well as the sequence of realizations y^t.

In Section 3, the authors characterize the set of payoffs that are achieved by belief-free equilibria in the presence of a public randomization device. The public randomization device suggests the regime $\mathcal{A}^t$ and the strategies ensure that all actions belonging to $\mathcal{A}_i^t$ are best-replies for player i. Suppose that p is the distribution of the public randomization device so that $p(\mathcal{A})$ is the probability of regime $\mathcal{A}$. A payoff vector (v_1, v_2) is strongly p-generated by a set of vectors $W_1 \times W_2$ if for each possible regime $\mathcal{A}$ and for each player i there is a mixed action profile α_i from the stage game whose support is included in $\mathcal{A}_i$, and continuation values can be selected from W_i so that α_i is a best reply for player i and results in total p-expected payoff v_i. Let $B_p(W)$ be the set of all vectors strongly p-generated by W. A set W for which $W \subset B_p(W)$ is said to be strongly self-generating and it is shown that all members of a strongly self-generating set are belief-free equilibrium payoffs. Furthermore, the set of all belief-free equilibrium payoffs of a give game using public randomization p is itself the largest self-generating set.

In Section it is shown that for increasing patience the limiting set of equilibria can be characterized by a family of linear programs. For the case where players monitor another with increasing accuracy, a simple formula is found that can be computed by linear programming methods. By means of examples it is shown that belief-free equilibria can support a large set of payoffs, but in general the set is not large enough to establish a folk theorem.

In Section 5 is is argued that for discount factors close enough to 1, any equilibrium payoff obtained using an i.i.d. public randomization can also be achieved in a belief-free equilibrium with a deterministic sequence of regimes. Finally, the results are applied to provide a sufficient condition for equilibrium payoffs in arbitrary two-action games with independent monitoring.