On a Dynamic Public Insurance Game with Heterogeneous Agents

TIMOTHY KAM AND RONALD STAUBER

ABSTRACT. We develop a class of dynamic taxation and public insurance games between a government whose aim is to maximize a measure of discounted expected total social welfare, and a continuum of private agents. Private agents are endogenously distributed on a finite set of individual states. We show that the sequential equilibrium payoff correspondence exists and can be found recursively. If the government can commit, there exists an optimal fixed tax-and-insure policy for any initial distribution of private agent states. We also prove that a socially optimal steady state exists in such a case.

KEYWORDS: Heterogeneity; limited commitment; public insurance; dynamic games.

J.E.L. CODE: C61; C73; D86; E24; H21

1. INTRODUCTION

WE DEVELOP A CLASS of dynamic taxation and public insurance games between a government whose aim is to maximize a measure of discounted expected total social welfare, and a heterogeneous society of private agents that is continuously distributed over a finite set of individual states. Each private agent can influence the stochastic evolution of his personal state through the choice of an effort variable. The cross-sectional distribution over private-agent personal states evolve according to a Markov operator that depends directly on private actions in a sequential equilibrium, and indirectly on equilibrium government policy via private agents' actions.

Potential applications of our framework include public unemployment and public health insurance design. In the case of public unemployment insurance design, private agents' personal states may represent their duration of employment or unemployment, their actions may represent their job-search intensity while unemployed,

This version: April 1, 2008

and work effort while employed. The vector of government policy actions will be relevant state-specific unemployment benefits or employment taxes. In the latter case of public health insurance design, a private agent's action may represent an index of his lifestyle choices such as diet, exercise, and social habits. With suitable assumptions on payoffs and the Markov transition law on private-agent states, the agents' personal states can correspond to their levels of current health and thus determine their claim on, or contribution to, the public health insurance fund. In this case, government policy contains a list of health insurance levies for the healthy and productive, and health insurance payouts for the infirmed.

We provide three existence results. First, although the public insurance allocation is history dependent, we show that for any given initial distribution over private-agent states and any initial government subsidy to the insurance scheme, the sequential equilibrium payoff correspondence exists and can be found recursively. Second, we prove that if the government can commit, there exists an optimal fixed tax-and-insure policy for any initial distribution of private agent states and government subsidy. Third, we also prove the existence of a socially optimal steady state which maximizes the long-run social welfare.

Our model extends the seminal work of Phelan and Stacchetti [12] to a game comprising a government that is faced with a non-degenerate distribution of anonymous players. In Phelan and Stacchetti [12] there is a continuum of private agents possessing identical individual state variables at each point in time. Therefore, the dynamic game characterized in Phelan and Stacchetti [12] reduces to one between a government and a representative private agent.[1]

Previous literature on public insurance focuses more on principal-agent problems (see e.g. Spear and Srivastava [16], Phelan and Townsend [13]). In providing a public tax-and-insurance mechanism, the government in our model is constrained by an initial subsidy to the insurance scheme – a realistic but conveniently missing assumption in the existing papers on public insurance design. For example, see the literature on unemployment insurance (e.g. Shavell and Weiss [14], Hopenhayn and

[1]This is also the effect in Atkeson [3] who has a single borrower and a sequence of finitely-lived lenders, where the public state variable is a one-dimensional capital stock held by the borrower.

Nicolini [6], Wang and Williamson [18], Zhao [20], Werning [19], Pavoni [11]). This realistic feature also allows us to characterize the equilibrium payoff correspondence without needing to exogenously fix promised utilities.

The remainder of this note is organized as follows. In section 2, we outline the model and its assumptions. In section 3, we define sequential equilibrium for our game and characterize the equilibrium payoff set as a function of the initial distribution over agent states and government subsidy. In section 4 we describe steady states, and discuss the case where the government can commit to a fixed policy.

2. The Model

The players of the insurance game are given by a continuum of heterogeneous private agents distributed on the unit interval $[0, 1]$, and a single insurance fund planner, or in short, a "government".

2.1 Transitions between private agent states

At each time period, every private agent is characterized by his *personal state*, which is drawn from the finite set $\mathcal{Z} := \{-N, \ldots, -1, 1, \ldots, M\}$, where $M, N \in \mathbb{Z}_+$. To clarify the intuition behind the model, we use the interpretation as an unemployment insurance design problem. In this environment, positive states $j > 0$ denote an agent who has been working for j periods, and negative states $j < 0$ denote an agent who has been unemployed for j periods. An agent in state $j > 0$ can only move to state $j+1$ or state -1, with the transition probabilities depending on an unobservable effort exerted by the agent, and on his current state. Similarly, an agent in state $j < 0$ can only move to state $j - 1$ or state 1, with transition probabilities depending on an unobservable effort and on the current state of the agent. By assuming that the transition probabilities are constant after M periods for $j > 0$, or N periods for $j < 0$, we can use state M to characterize agents who have been employed for more than M periods, and state $-N$ to characterize agents who have been unemployed for more than N periods.

Let Y_t denote the state of an individual agent at time t, and let $a_t(j) \in \mathbb{R}_+$ denote effort exerted by an agent who is in state j in period t. Thus, in our example, effort $a_t(j)$ represents job-search intensity if $j < 0$, and work intensity if $j > 0$. The transition probabilities for private-agent states are then characterized by a controlled Markov chain, where effort is the control variable. The transition probabilities of the corresponding Markov chains are defined for each state $j \in \mathcal{Z}$ by functions $p^j : \mathbb{R}_+ \to [0, 1]$, where

$$p^j(a) = \begin{cases} \mathbb{P}(Y_{t+1} = 1 \,|\, Y_t = j, a_t(j) = a), & \text{for } j < 0 \\ \mathbb{P}(Y_{t+1} = j+1 \,|\, Y_t = j, a_t(j) = a), & \text{for } 0 < j < M \\ \mathbb{P}(Y_{t+1} = j \,|\, Y_t = j, a_t(j) = a), & \text{for } j = M \end{cases} \cdot$$

It follows that

$$1 - p^j(a) = \begin{cases} \mathbb{P}(Y_{t+1} = j \,|\, Y_t = j, a_t(j) = a), & \text{for } j = -N \\ \mathbb{P}(Y_{t+1} = j-1 \,|\, Y_t = j, a_t(j) = a), & \text{for } -N < j < 0 \\ \mathbb{P}(Y_{t+1} = -1 \,|\, Y_t = j, a_t(j) = a), & \text{for } j > 0 \end{cases} \cdot$$

We assume that for all $j \in \mathcal{Z}$, p^j is strictly increasing, strictly concave, and continuously differentiable, $p^j(0) = 0$ for $j < 0$ and $p^j(0) \geq 0$ for $j > 0$, and for every $a \in \mathbb{R}_+$: (i) $p^j(a) < p^{j+1}(a)$, (ii) $dp^j(a)/da$ is increasing (decreasing) in j for $j < 0$ ($j > 0$). These assumptions provide a natural model for empirical conditional hazard rates of unemployment duration, for example.

For any function $w : \mathcal{Z} \to \mathbb{R}$, define $\mathbb{E}_{p^j(a)}[w(i)] := \mathbb{E}\left[w(Y_{t+1}) \middle| Y_t = j, a_t(j) = a\right]$. Denote the distribution over states in period t by λ_t, and let $\Delta(\mathcal{Z})$ denote the set of distributions over $\mathcal{Z}$. We will consider symmetric equilibria where all agents in a particular state exert the same amount of effort. We refer to the effort levels as actions, and denote an action vector by $a_t \in \mathbb{R}^{\mathcal{Z}}$.[2] It follows that any fixed sequence of symmetric effort levels $\{a_t\}_{t=0}^{\infty}$ defines the evolution of the distribution of states, which will be deterministic in the aggregate, assuming that an appropriate law of

[2]We use $\mathbb{R}^{\mathcal{Z}}$ to denote the set of real-valued functions with domain $\mathcal{Z}$. If the cardinality of $\mathcal{Z}$ is finite, then $\mathbb{R}^{\mathcal{Z}}$ is the $|\mathcal{Z}|$-dimensional Euclidean space. If $a \in \mathbb{R}^{\mathcal{Z}}$, we use $a(j)$ to denote the j-th coordinate of a, i.e., the action taken by an agent in state j.

large numbers holds (see Judd [8]). Given any fixed period t action vector a_t, denote the induced matrix of transition probabilities by $P(a_t)$. The period $t+1$ distribution over personal states induced by a period t distribution λ_t and symmetric action vector a_t can then be computed as $\lambda_{t+1} = \lambda_t P(a_t)$.

2.2 The government and agent consumption insurance

For simplicity, suppose that all agents receive a constant per-period wage $m > 0$ during periods when they are in a state $j > 0$, and no wage in states $j < 0$. Agents have no opportunity to save or borrow, and thus consume their entire income in each period. Agents in this world can only insure their consumption risk by participating in the public insurance scheme.

The government cannot observe effort levels, and can only observe agents' current personal states. Thus, the government is restricted to state-contingent transfers, where taxes or insurance payouts are conditioned on the personal state of each agent. Note that agents are assumed to be anonymous, in the sense that each agent's individual history of personal states is not observed by the government.[3] The government can however observe the aggregate distribution over personal states at the beginning of each period t, and has a record of the evolution of these distributions $\lambda^t := (\lambda_0, \dots, \lambda_t)$. Thus, the government can chose its time t transfer policy as a function of the history λ^t. Since agents are distributed on a continuum, individual deviations have no effect on the evolution of the aggregate distribution over states, and therefore no effect on the future government policies or other agents' actions.

We denote by $b_t(j)$ the period t net transfer that an agent in state j receives. A period t policy vector for the government is then given by a vector $b_t \in \mathbb{R}^{\mathcal{Z}}$. A policy vector must satisfy $b_t(j) \geq -m$ for all $j > 0$, and $b_t(j) \geq 0$ for all $j < 0$. In addition, we assume that there is an exogenously given upper bound on the amount of benefits that the government can pay, i.e., there is a constant $\overline{m}$ so that $b_t(j) \leq \overline{m}$

[3] The model can easily be extended to a more general framework, where transfers can be conditioned on any finite history of agent employment states. This can be achieved by redefining personal states to include a finite history of employment states. Doing this yields a model where the insurance scheme provides a finer incentive structure, that does not only depend on the current employment state, but also on a finite number of previous employment states.

for all j and t. Letting $c_t^{b_t}(j)$ denote consumption in period t of an agent in state j, we get

$$c_t^{b_t}(j) = \begin{cases} m + b_t(j), & \text{if } j > 0 \\ b_t(j), & \text{if } j < 0 \end{cases}.$$

We assume that the agents' common utility function is separable in consumption and effort levels, so that the period t utility of an agent who is in state j can be expressed by $u(c_t^{b_t}(j)) - \phi(a_t(j))$, where u is non-negative, strictly increasing and concave and ϕ is non-negative, strictly increasing and strictly convex.

At the initial time $t = 0$, the government has an amount $s_0 \geq 0$ at its disposal, which it can use to subsidize the insurance scheme. We will refer to s_0 as government savings at time $t = 0$, and use s_t to denote government savings at time t. The government's period t net surplus from the insurance program, $-\sum_{j \in \mathcal{Z}} \lambda_t(j) b_t(j)$, is added to the start-of-period savings s_t, and the resulting amount earns an exogenous interest rate $r = (1 - \delta)/\delta$ to yield the savings at the start of the following period. The amount of the subsidy s_0 is exogenous, and no additional outside funds can be added to the savings at later time periods. Moreover, the government cannot borrow funds, so any increase in savings must come from a surplus.

We assume that the government cannot commit to a sequence of policy vectors, and that its objective is to maximize the normalized expected discounted average utilities of the agents

$$\mathbb{E}_{\lambda_0} \left\{ (1 - \delta) \sum_{t=0}^{\infty} \delta^t \sum_{j \in \mathcal{Z}} \lambda_t(j) \left[u(c_t^{b_t}(j)) - \phi(a_t(j)) \right] \right\},$$

given the initial subsidy s_0 and initial distribution over states λ_0.

Depending on the value of the initial distribution λ_0, we can interpret λ_0 as a situation where there is a recession, when unemployment is high, or a boom, when unemployment is low. By varying the initial subsidy s_0, our framework allows us to analyze how additional subsidies to the insurance scheme can yield higher welfare, or more favorable distributions over employment states in the long run. In addition, since the objective of the optimal public insurance scheme lies in maximizing the welfare over a population of agents, given initial values for s_0 and λ_0, we can determine the optimal welfare of any individual agent as a function of his initial state endogenously. This is in contrast to principal-agent models where only a single agent is considered, and promised utilities are given exogenously, as in Hopenhayn and Nicolini [6].

2.3 Game states, histories, and strategies

The model described above defines a dynamic game, where the continuation game at the beginning of each period t is characterized by the values of s_t and λ_t. We define a *game state* to be a pair (s_t, λ_t). Game states change over the course of play as a function of the players' actions. We restrict agents to use symmetric actions, and assume that an appropriate law of large numbers holds (as in Judd [8] or Uhlig [17]), in which case the game state evolves deterministically as a function of the government policy vector b_t and agents' action vector a_t. To simplify the characterization of the equilibrium correspondence, we assume that there is an upper bound $\overline{s}$ on the amount of government savings, in the sense that if the government accumulates more than $\overline{s}$, the amount exceeding $\overline{s}$ is forfeited without accruing any benefit. The transition function for s_t is then defined by the following function:

$$F(s_t, \lambda_t, b_t) := \min\left\{\frac{1}{\delta}(s_t - \lambda_t \cdot b_t), \overline{s}\right\},$$

where $\lambda_t \cdot b_t$ denotes the inner product $\sum_j \lambda_t(j) b_t(j)$. The transition function for λ_t is simply $\lambda_{t+1} = \lambda_t P(a_t)$. Since the government's spending is constrained by the available savings, the set of feasible government policies in period t is a function of the game state (s_t, λ_t), as defined by the following compact-valued correspondence:

$$\mathcal{B}(s_t, \lambda_t) := \left\{b \in \mathbb{R}^{\mathcal{Z}} \mid s_t - \lambda_t \cdot b \geq 0,\ -m \leq b(j) \leq \overline{m} \text{ for } j > 0,\ 0 \leq b(j) \leq \overline{m} \text{ for } j < 0\right\}.$$

At the beginning of each period, all players can observe the realization of a public random variable $X_t \sim$ i.i.d.$U[0,1]$, and can condition their strategies on the history of past realizations of this variable, $x^t = (x_0, \ldots, x_t)$. The use of public correlation is a standard tool that convexifies the set of equilibrium payoffs (see e.g. Mailath and

Samuelson [10], Judd et al. [7], Phelan and Stacchetti [12]). Denote a public history at the beginning of period t by $h^t = (s^t, \lambda^t, x^t, b^{t-1})$, where $s^t = (s_0, \ldots, s_t)$ is the realized history of government savings, $\lambda^t = (\lambda_0, \ldots, \lambda_t)$ the history of distributions over private agents' states, $x^t = (x_0, \ldots, x_t)$ the realized history of the correlation variables, and $b^{t-1} = (b_0, \ldots, b_{t-1})$ the policy history. A strategy for the government is a sequence of functions β_t, where each β_t maps a history h^t to a policy vector $b_t = \beta_t(h^t) \in \mathbb{R}^{\mathcal{Z}}$. A symmetric strategy for the agents is a sequence of functions α_t, where each α_t maps a history (h^t, b_t) to an action vector $a_t = \alpha_t(h^t, b_t) \in \mathbb{R}^{\mathcal{Z}}$. A strategy profile is defined by a pair $\sigma = (\beta, \alpha) := (\{\beta_t\}_t, \{\alpha_t\}_t)$.

Figure 1 summarizes the timing of information and actions in each period.

FIGURE 1. Information and timing of the actions

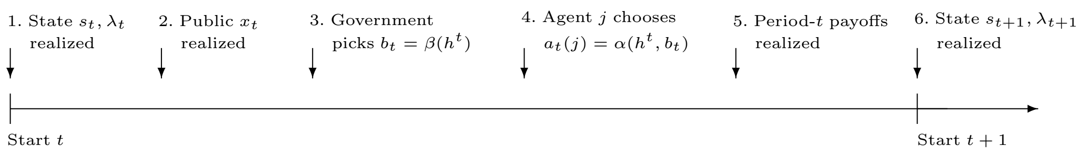

The j'th component of $\alpha_t(h^t, b_t)$ specifies the action prescribed by the strategy α_t for an agent whose personal state Y_t is equal to j. Since agents are "small" players, individual unilateral deviations have no effect on the evolution of λ_t, and thus no effect on future play. Therefore, the private history of an agent's past actions and states will not affect his optimal action. We can thus ignore that agents could also condition their actions on their private histories.

Any initial game state state (s_0, λ_0) and strategy profile σ recursively generate the following payoffs, where the subscript G denotes the government payoffs, i.e., the expected social welfare, and a subscript j denotes the expected utility of an agent who starts out in state j:

$$V_G(s_0, \lambda_0, \sigma) = \sum_{j \in \mathcal{Z}} \lambda_0(j) V_j(s_0, \lambda_0, \sigma),$$

where

$$V_j(s_0, \lambda_0, \sigma) = \mathbb{E}_{\lambda_0,\sigma}\left\{(1-\delta)\sum_{t=0}^{\infty}\delta^t\left[u(c_t^{b_t}(Y_t)) - \phi(a_t(Y_t))\right]\middle| Y_0 = j\right\}.$$

We use $V(s_0, \lambda_0, \sigma)$ to denote the vector containing the individual personal state-contingent expected payoffs and the government payoffs, so $V(s_0, \lambda_0, \sigma) \in \mathbb{R}^{\mathcal{Z}\cup\{G\}}$ for every (s_0, λ_0, σ). Let $\overline{\mathcal{Z}} := \mathcal{Z} \cup \{G\}$. For any vector $v \in \mathbb{R}^{\overline{\mathcal{Z}}}$, we use $v(j)$ to denote the payoff of an agent in state j, and $v(G)$ to denote the government payoff.

3. Sequential Equilibria

We consider symmetric sequential equilibria defined as follows:

Definition 1. *A strategy profile σ is a* symmetric sequential equilibrium (SSE) *for (s_0, λ_0), if for all t, and all h^t accessible from (s_0, λ_0),*

(i) $V_G(s_t, \lambda_t, \sigma|_{h^t}) \geq V_G(s_t, \lambda_t, \gamma, \alpha|_{h^t})$, *where γ denotes any alternative government continuation strategy;*

(ii) *For all j and b_t, $\alpha_t(h^t, b_t)(j)$ is an optimal action for a state j agent if the subsequent policies and agent actions are generated by the strategy $\sigma|_{(h^t,b_t)}$.*

Since any individual agent is small, he cannot affect λ_{t+1} by changing his personal effort level. Thus, as long as α_t prescribes an optimal action for every private-agent state, we can ignore individual deviations by the agents. Thus, we only need to determine what happens after a deviation by the government. We follow Abreu [1] and Phelan and Stacchetti [12] and consider only extreme punishments that yield the lowest continuation equilibrium payoffs for the government. Since agents are small players, such punishments must generate optimal actions for the agents. Thus, in response to a deviation by the government, agents play an action profile $\tilde{a}$ that generates a distribution over states $\lambda_{t+1} = \lambda_t P(\tilde{a})$, and an equilibrium continuation strategy for the government as a function of λ_{t+1} that yields the lowest continuation equilibrium payoff for the government, and is such that $\tilde{a}$ is optimal given the continuation equilibrium.

We define the equilibrium value correspondence $\mathcal{V} : [0, \overline{s}] \times \Delta(\mathcal{Z}) \rightrightarrows \mathbb{R}^{\overline{\mathcal{Z}}}$ to be the set of payoff vectors that can be achieved in some SSE, as a function of the initial game state, i.e.,

$$\mathcal{V}(s_0, \lambda_0) = \{V(s_0, \lambda_0, \sigma) \,|\, \sigma \text{ is a SSE for } (s_0, \lambda_0)\}.$$

We show next that the correspondence $\mathcal{V}$ can be characterized as the unique fixed point of a monotone set-valued operator, by adapting the approach of Phelan and Stacchetti [12], and Atkeson [3], who was the first to introduce a public state variable. These, in turn extend the techniques of self-generation and factorization pioneered by Abreu et al. [2] for repeated games. The proofs for our results follow standard methods. The consequence of these results is that in principle, the equilibrium correspondence can be computed recursively by applying the operator to a suitably defined initial correspondence.

Definition 2. *Let* $\mathcal{W} : [0, \overline{s}] \times \Delta(\mathcal{Z}) \rightrightarrows \mathbb{R}^{\overline{\mathcal{Z}}}$ *be a compact- and convex-valued correspondence having the property that* $w(G) = \sum_{j \in \mathcal{Z}} \lambda(j) w(j)$ *for all* $(s, \lambda, w) \in \operatorname{graph}(\mathcal{W})$*. A vector* $(b, a, s', \lambda', w) \in \mathbb{R}^{\mathcal{Z}} \times \mathbb{R}^{\mathcal{Z}} \times [0, \overline{s}] \times \Delta(\mathcal{Z}) \times \mathbb{R}^{\overline{\mathcal{Z}}}$ *is* consistent with respect to $\mathcal{W}$ at (s, λ) *if*

(i) $b \in \mathcal{B}(s, \lambda)$
(ii) $s' = F(s, \lambda, b)$;
(iii) $\lambda' = \lambda P(a)$;
(iv) $w \in \mathcal{W}(s', \lambda')$;
(v) *For all* $j \in \mathcal{Z}$,

$$a(j) \in \operatorname{argmax}_{a'} \left\{ (1 - \delta) \left[u(c^b(j)) - \phi(a') \right] + \delta \mathbb{E}_{p^j(a')}[w(i)] \right\}.$$

Definition 3. *For* $(s, \lambda, b) \in [0, \overline{s}] \times \Delta(\mathcal{Z}) \times \mathbb{R}^{\mathcal{Z}}$ *such that* $b \in \mathcal{B}(s, \lambda)$*, let*

$$\pi(s, \lambda, b) := \min_{(a', s'', \lambda'', w')} \left[(1 - \delta) \sum_{j \in \mathcal{Z}} \lambda(j) [u(c^b(j)) - \phi(a'(j))] + \delta \sum_{j \in \mathcal{Z}} \lambda''(j) w'(j) \right],$$

subject to $(b, a', s'', \lambda'', w')$ *is consistent with respect to* $\mathcal{W}$ *at* (s, λ)*. Let* $(\tilde{a}(s, \lambda, b), \tilde{s}'(s, \lambda, b), \tilde{\lambda}'(s, \lambda, b), \tilde{w}(s, \lambda, b))$ *denote the solutions to the corresponding minimization problem. A vector* $(b, a, s', \lambda', w) \in \mathbb{R}^{\mathcal{Z}} \times \mathbb{R}^{\mathcal{Z}} \times [0, \overline{s}] \times \Delta(\mathcal{Z}) \times \mathbb{R}^{\overline{\mathcal{Z}}}$ *is said to be* admissible with respect to $\mathcal{W}$ at (s, λ) *if*

(i) (b, a, s', λ', w) *is consistent with respect to* $\mathcal{W}$ *at* (s, λ);

(ii) $(1-\delta)\sum_{j\in\mathcal{Z}}\lambda(j)[u(c^b(j))-\phi(a(j))]+\delta\sum_{j\in\mathcal{Z}}\lambda'(j)w(j)\geq\max_{b'\in\mathcal{B}(s,\lambda)}\pi(s,\lambda,b')$.

The payoff vector defined by an admissible vector (b,a,s',λ',w) at (s,λ) is given by

$$E_G(b,a,s',\lambda',w)(s,\lambda)=(1-\delta)\sum_{j\in\mathcal{Z}}\lambda(j)[u(c^b(j))-\phi(a(j))]+\delta\sum_{j\in\mathcal{Z}}\lambda'(j)w(j),\ \text{and}$$

$$E_j(b,a,s',\lambda',w)(s,\lambda)=(1-\delta)\left[u(c^b(j))-\phi(a(j))\right]+\delta\mathbb{E}_{p^j(a(j))}[w(i)].$$

Note that $E_G(b,a,s',\lambda',w)(s,\lambda)=\sum_{j\in\mathcal{Z}}\lambda(j)E_j(b,a,s',\lambda',w)(s,\lambda)$. Let

$$\mathbf{B}(\mathcal{W})(s,\lambda):=\text{co}\{E(b,a,s',\lambda',w)(s,\lambda)\,|\,(b,a,s',\lambda',w)\text{ is admissible with respect to }\mathcal{W}\text{ at }(s,\lambda)\},$$

where co denotes the convex hull of a set.

Lemma 1. *$\mathcal{V}(s_0,\lambda_0)$ is a bounded subset of $\mathbb{R}^{\overline{\mathcal{Z}}}$ for every (s_0,λ_0). Furthermore, graph($\mathcal{V}$) $\subset[0,\overline{s}]\times\Delta(\mathcal{Z})\times\mathbb{R}^{\overline{\mathcal{Z}}}$ is bounded.*

Proof: Since the vectors b_t are assumed to be bounded, each agent's payoff is bounded above by $u(m+\overline{m})$, and therefore, so is the government's payoff. For the lower bound, note that the greatest incentive to exert effort a is given when having a $j>0$ guarantees the upper bound on utility $u(m+\overline{m})$, and having a $j<0$ yields a utility of zero. Thus, an upper bound on exerted effort is given by the maximum over $j\in\{-1,1\}$, of the effort level a^* that solves $-(1-\delta)\phi'(a^*)+\delta(p^j)'(a^*)u(m+\overline{m})=0$. $-\phi(a^*)$ yields a lower bound on the payoffs of the agents and of the government. Since $\mathcal{V}$ has compact domain, it follows that graph($\mathcal{V}$) is a bounded subset of $[0,\overline{s}]\times\Delta(\mathcal{Z})\times\mathbb{R}^{\overline{\mathcal{Z}}}$. □

Lemma 2. *$\mathbf{B}(\mathcal{W})\subset\mathbf{B}(\mathcal{W}')$ if $\mathcal{W}\subset\mathcal{W}'$.*

Proof: Given any (s,λ), every vector which is consistent w.r.t. $\mathcal{W}$ must also be consistent w.r.t. $\mathcal{W}'$. Thus, $\pi(s,\lambda,b)$ calculated w.r.t. $\mathcal{W}'$ cannot exceed $\pi(s,\lambda,b)$ calculated w.r.t. $\mathcal{W}$, for every $b\in\mathcal{B}(s,\lambda)$. It follows that vectors that are admissible w.r.t. $\mathcal{W}$ are also admissible w.r.t. $\mathcal{W}'$, which implies the result. □

Lemma 3. *If $\mathcal{W}$ has compact graph, then $\mathbf{B}(\mathcal{W})$ has compact graph.*

Proof: If $\mathcal{W}$ is compact, then $\mathbf{B}(\mathcal{W})$ is bounded since one-period utilities and continuation payoffs, which are drawn from $\mathcal{W}$, are bounded. $\mathbf{B}(\mathcal{W})$ is closed, since the limit of any sequence of admissible vectors must also be an admissible vector. □

PROPOSITION 1 (Self-generation). *If* $\mathcal{W}(s,\lambda) \subset \mathbf{B}(\mathcal{W})(s,\lambda)$ *for all* $(s,\lambda) \in [0,\overline{s}] \times \Delta(\mathcal{Z})$*, then* $\mathbf{B}(\mathcal{W}) \subset \mathcal{V}$.

Proof: The proof proceeds by constructing, for every $\tau_0 \equiv (s_0, \lambda_0, v_0) \in \text{graph}(\mathbf{B}(\mathcal{W}))$, an equilibrium strategy σ^{τ_0} with $V(s_0, \lambda_0, \sigma^{\tau_0}) = v_0$. For every $\tau \equiv (s, \lambda, v) \in \text{graph}(\mathbf{B}(\mathcal{W}))$ and $x \in [0,1]$, let $A(x,\tau) \in \mathbb{R}^{\mathcal{Z}} \times \mathbb{R}^{\mathcal{Z}} \times [0,\overline{s}] \times \Delta(\mathcal{Z}) \times \mathbb{R}^{\overline{\mathcal{Z}}}$ denote a vector that is admissible with respect to $\mathcal{W}$ at (s,λ) and has the property that $\int_0^1 E(A(x,\tau))(s,\lambda)dx = v$, and use superscripts b, a, s', λ' and w for A to denote the corresponding components of the vector $A(x,\tau)$. By definition of $\mathbf{B}(\mathcal{W})$, such a function A exists for every $\tau \in \text{graph}(\mathbf{B}(\mathcal{W}))$ and $x \in [0,1]$, and can be assumed to be measurable.

We define a strategy σ^{τ_0} recursively: At $t = 0$, let

$$\beta_0^{\tau_0}(s_0, \lambda_0, x_0) = A^b(x_0, \tau_0),$$

and

$$\alpha_0^{\tau_0}(s_0, \lambda_0, x_0, b_0) = \begin{cases} A^a(x_0, \tau_0), & \text{if } b_0 = \beta_0^{\tau_0}(s_0, \lambda_0, x_0) \\ \tilde{a}(s_0, \lambda_0, b_0), & \text{otherwise} \end{cases}$$

where the function $\tilde{a}$ is defined in Definition 2. At $t = 1$, let

$$\beta_1^{\tau_0}(s^1, \lambda^1, x^1, b_0) = \begin{cases} A^b(x_1, A^{s'}(x_0, \tau_0), A^{\lambda'}(x_0, \tau_0), A^w(x_0, \tau_0)), & \text{if } b_0 = \beta_0^{\tau_0}(s_0, \lambda_0, x_0) \\ A^b(x_1, \tilde{s}'(s_0, \lambda_0, b_0), \tilde{\lambda}'(s_0, \lambda_0, b_0), \tilde{w}(s_0, \lambda_0, b_0)), & \text{otherwise} \end{cases},$$

and

$$\alpha_1^{\tau_0}(s^1, \lambda^1, x^1, b^1) = \begin{cases} A^a(x_1, A^{s'}(x_0, \tau_0), A^{\lambda'}(x_0, \tau_0), A^w(x_0, \tau_0)), & \text{if } b_0 = \beta_0^{\tau_0}(s_0, \lambda_0, x_0) \text{ and} \\ & b_1 = \beta_1^{\tau_0}(s^1, \lambda^1, x^1, b_0) \\ \tilde{a}(s_1, \lambda_1, b_1), & \text{otherwise} \end{cases}.$$

Let $W_{t+1}(s_0, x^t, \lambda^t, b^t)$ denote the continuation payoffs after period t. Then

$$W_1(s_0, \lambda_0, x_0, b_0) = \begin{cases} A^w(x_0, \tau_0), & \text{if } b_0 = \beta_0^{\tau_0}(s_0, \lambda_0, x_0) \\ \tilde{w}(s_0, \lambda_0, b_0), & \text{otherwise} \end{cases},$$

and

$$W_2(s^1,\lambda^1,x^1,b^1) =$$
$$\begin{cases} A^w(x_1, A^{s'}(x_0,\tau_0), A^{\lambda'}(x_0,\tau_0), A^w(x_0,\tau_0)), & \text{if } b_0 = \beta_0^{\tau_0}(s_0,\lambda_0,x_0) \text{ and} \\ & b_1 = \beta_1^{\tau_0}(s^1,\lambda^1,x^1,b_0) \\ \tilde{w}(s_1,\lambda_1,b_1), & \text{otherwise} \end{cases}.$$

We can continue in this way to recursively define continuation payoffs and corresponding strategies $\sigma^{\tau_0} = (\beta^{\tau_0}, \alpha^{\tau_0})$, to get

$$W_{t+1}(s^t,\lambda^t,x^t,b^t) = \begin{cases} A^w(x_t,s_t,\lambda_t,W_t(s^{t-1},\lambda^{t-1},x^{t-1},b^{t-1})), & \text{if the government has} \\ & \text{never deviated} \\ \tilde{w}(s_t,\lambda_t,b_t), & \text{otherwise} \end{cases},$$

$$\beta_t^{\tau_0}(s^t,\lambda^t,x^t,b^{t-1}) = A^b(x_t,s_t,\lambda_t,W_t(s^{t-1},\lambda^{t-1},x^{t-1},b^{t-1})), \text{ and}$$

$$\alpha_t^{\tau_0}(s^t,\lambda^t,x^t,b^t) = \begin{cases} A^a(x_t,s_t,\lambda_t,W_t(s^{t-1},\lambda^{t-1},x^{t-1},b^{t-1})), & \text{if the government has} \\ & \text{never deviated} \\ \tilde{a}(s_t,\lambda_t,b_t), & \text{otherwise} \end{cases}.$$

For any $\tau \equiv (s,\lambda,v) \in \text{graph}(\mathbf{B}(\mathcal{W}))$, we now show that $v = V(s,\lambda,\sigma^\tau)$. Since $v(G) = \sum_{j\in\mathcal{Z}} \lambda(j)v(j)$, and $V_G(s,\lambda,\sigma^\tau) = \sum_{j\in\mathcal{Z}}\lambda(j)V_j(s,\lambda,\sigma^\tau)$, it suffices to show that $v(j) = V_j(s,\lambda,\sigma^\tau)$ for all $j \in \mathcal{Z}$.

For every $\tau \in \text{graph}(\mathbf{B}(\mathcal{W}))$, let $b_x := A^b(x,\tau)$, $a_x := A^a(x,\tau)$, $s'_x := A^{s'}(x,\tau)$, $\lambda'_x := A^{\lambda'}(x,\tau)$ and $w_x := A^w(x,\tau)$. Denote continuation strategies following (s,λ,x,b_x) as $\sigma^\tau|_{s,\lambda,x,b_x} = \sigma^{\tau'_x}$. Then

$$v(j) = \int_0^1 \left\{(1-\delta)\left[u(c_x^{b_x}(j)) - \phi(a_x(j))\right] + \delta\mathbb{E}_{p^j(a_x(j))}[w_x(i)]\right\} dx,$$

and

$$V_j(s,\lambda,\sigma^\tau) = \int_0^1 \left\{(1-\delta)\left[u(c_x^{b_x}(j)) - \phi(a_x(j))\right] + \delta\mathbb{E}_{p^j(a_x(j))}[V_i(s'_x,\lambda'_x,\sigma^{\tau'_x})]\right\} dx.$$

Subtracting the corresponding equations for any j, yields

$$|v(j) - V_j(s,\lambda,\sigma^\tau)| \le \delta \sup_{(i,s',\lambda',w(i))\in\text{graph}(\mathbf{B}(\mathcal{W}))} \left| w(i) - V_i(s',\lambda',\sigma^{s',\lambda',w})\right|.$$

Since this equation must hold for all $(j, s, \lambda, v(j)) \in \text{graph}(\mathbf{B}(\mathcal{W}))$, and all payoffs are bounded by Lemma 1, we get

$$\sup_{(j,s,\lambda,v(j))\in \text{graph}(\mathbf{B}(\mathcal{W}))} |v(j) - V_j(s, \lambda, \sigma^\tau)| \leq \delta \sup_{(i,s',\lambda',w(i))\in \text{graph}(\mathbf{B}(\mathcal{W}))} \left| w(i) - V_i(s', \lambda', \sigma^{s',\lambda',w}) \right|.$$

Therefore, $v(j) = V_j(s, \lambda, \sigma^\tau)$ for all j and $\tau \equiv (s, \lambda, v) \in \text{graph}(\mathbf{B}(\mathcal{W}))$.

It follows that σ^{τ_0} defines a sequential equilibrium, once we show that there does not exist a profitable multi-round deviation for the government. But this is a consequence of a standard "one-shot deviation principle", which can be derived along the lines of Mailath and Samuelson [10]. □

Together with the next proposition, self-generation implies that $\mathcal{V}$ is the unique fixed point of $\mathbf{B}(\mathcal{V})$.

Proposition 2 (Factorization). *$\mathcal{V}(s, \lambda) \subset \mathbf{B}(\mathcal{V})(s, \lambda)$ for all (s, λ), and $\mathcal{V}$ has compact graph.*

Proof: The closure of $\text{graph}(\mathcal{V})$, $\overline{\text{graph}(\mathcal{V})}$ defines a compact correspondence denoted by $\overline{\mathcal{V}}$. $\mathbf{B}(\overline{\mathcal{V}})$ is compact by Lemma 3. If $v \in \mathcal{V}(s, \lambda)$ for some (s, λ), and $\sigma = (\beta, \alpha)$ is a corresponding equilibrium strategy, define $b(x) = \beta_0(s, \lambda, x)$, $a(x) = \alpha_0(s, \lambda, x, b(x))$, $s'(x) = F(s, \lambda, b(x))$, $\lambda'(x) = \lambda P(a(x))$, and $w(x) = V(s'(x), \lambda'(x), \sigma|_{(s,\lambda,x,b(x))})$. Then for every $x \in [0, 1]$, $(b(x), a(x), s'(x), \lambda'(x), w(x))$ is admissible with respect to $\mathcal{V}$ at (s, λ), and $v = \int_0^1 E(b(x), a(x), s'(x), \lambda'(x), w(x))dx$, which implies that $\mathcal{V} \subset \mathbf{B}(\overline{\mathcal{V}})$. Since $\mathbf{B}(\overline{\mathcal{V}})$ is compact, it follows that $\overline{\mathcal{V}} \subset \mathbf{B}(\overline{\mathcal{V}})$. Hence, $\overline{\mathcal{V}}$ is self-generating, so $\overline{\mathcal{V}} \subset \mathcal{V}$ and thus $\mathcal{V} = \mathbf{B}(\mathcal{V})$. □

We call an equilibrium strategy corresponding to an initial state (s_0, λ_0) an *optimal equilibrium* if it maximizes the government's payoff among all equilibria for (s_0, λ_0). Since $\mathcal{V}$ is compact-valued by the previous theorem, an optimal equilibrium exists for all (s_0, λ_0). Let $\mathcal{V}_G(s, \lambda)$ denote the projection of $\mathcal{V}(s, \lambda)$ onto the G coordinate, i.e., the government's component of the equilibrium payoff correspondence, and let $\overline{v}_G(s, \lambda) := \max_{v\in\mathcal{V}_G(s,\lambda)} v$. Then a strategy profile σ constitutes an optimal equilibrium if $V_G(s_0, \lambda_0, \sigma) = \overline{v}_G(s_0, \lambda_0)$.[4]

[4]Phelan and Stacchetti [12] define *upper* and *lower boundaries* for equilibrium value correspondences, and define *best* and *worst equilibria* to be equilibria for which payoffs lie on the upper and respectively, lower boundary of the equilibrium value correspondence. Defining upper and lower boundaries in our model would involve maximizing, respectively minimizing government payoffs given an initial distribution and given a initial vector of agents' payoffs. Since the government payoff is just the expected value of the agents' payoffs using the initial distribution, it is uniquely defined by the specification of a payoff vector and initial distribution, and thus, the upper and lower boundaries would coincide.

Note that if σ is an optimal equilibrium for (s_0, λ_0), then the continuation strategies $\sigma|_{(s_0,\lambda_0,x_0,\beta_0(s_0,\lambda_0,x_0),\alpha_0(s_0,\lambda_0,x_0))}$ need not be optimal equilibrium strategies for the initial distribution $\lambda_1 = \lambda_0 P(\alpha_0(s_0, \lambda_0, x_0))$, for any realization of x_0. This is because the agents' continuation payoffs corresponding to an optimal equilibrium for $\lambda_0 P(\alpha_0(s_0, \lambda_0, x_0))$ may prevent $\alpha_0(s_0, \lambda_0, x_0)$ from being optimal in the initial period, and a different effort vector may yield a transition to a second-period distribution that is different from $\lambda_0 P(\alpha_0(s_0, \lambda_0, x_0))$.

4. Steady States and Fixed Policies

We now define and characterize steady states of our model, and analyze the evolution of the system when the government can commit to a fixed policy.

Definition 4. *A vector* $(b, a, s, \lambda, v) \in \mathbb{R}^{\mathcal{Z}} \times \mathbb{R}^{\mathcal{Z}} \times \mathbb{R} \times \Delta(\mathcal{Z}) \times \mathbb{R}^{\overline{\mathcal{Z}}}$ *is a* steady state vector *if:*

(i) $\lambda = \lambda P(a)$,

(ii) *For all* $j \in \mathcal{Z}$, $a(j) \in \operatorname{argmax}_{a'} \left\{(1-\delta)\left[u(c^b(j)) - \phi(a')\right] + \delta \mathbb{E}_{p^j(a')}[v(i)]\right\}$, *and*

$v(j) = (1-\delta)\left[u(c^b(j)) - \phi(a(j))\right] + \delta \mathbb{E}_{p^j(a(j))}[v(i)]$,

(iii) $v(G) = \sum_{j\in\mathcal{Z}} \lambda(j) v(j)$,

(iv) $s - \frac{\lambda \cdot b}{(1-\delta)} = 0$.

A steady state vector describes a path along which all variables describing our system are constant. In a steady state, agents are required to choose optimal effort levels, but no maximization of social welfare by the government is imposed. Each steady state does however define a unique corresponding level of social welfare, $v(G)$. The one-period steady state government budget surplus, $-\lambda \cdot b$, yields a present value of $-\frac{\lambda \cdot b}{(1-\delta)}$ for the total surplus. If the government wants to sustain a steady state with $\lambda \cdot b > 0$, a subsidy $\frac{\lambda \cdot b}{(1-\delta)}$ is required, which yields a constant stream of interest payments to finance the per-period deficit $\lambda \cdot b$. If $-\lambda \cdot b$ is non-negative, i.e., if $\lambda \cdot b \leq 0$, the government does not need to subsidize the steady state. In that case, we can interpret s as the value of outstanding government debt whose interest can be serviced using the surplus from the insurance scheme. Thus, the variable s in

Definition 4, which can also be negative, is the cost to the government of sustaining the corresponding steady state.

We assumed that any vector of government transfers is bounded, such that $-m \leq b(j) \leq \overline{m}$ for $j > 0$, and $0 \leq b(j) \leq \overline{m}$ for $j < 0$, where $\overline{m}$ is an exogenous constant. A transfer vector is *unconstrained feasible* if it satisfies these bounds. The following proposition shows that any unconstrained feasible transfer vector b defines a unique corresponding steady state vector.[5]

PROPOSITION 3. *Given any unconstrained feasible transfer vector b, there exists a unique corresponding steady state vector $(b, a^b, s^b, \lambda^b, v^b)$.*

Proof: Fix any feasible vector b. Given the assumptions we made regarding the transition probabilities $p^j(a)$ and the cost function $\phi(a)$, standard results from dynamic programming imply the existence of a unique value function $v^b : \mathcal{Z} \to \mathbb{R}$, and a unique corresponding vector of optimal actions a^b, that together satisfy condition (ii) of Definition 4. We show that the Markov chain defined by a^b must always have a unique ergodic set, which will also be regular.[6] If $a^b(-N) = 0$, $p^{-N}(a^b(-N)) = 0$, in which case $-N$ is absorbing and all other states are transient. This is a consequence of the assumption that $p^j(a^b(j)) < 1$ for all states j. Thus, $\{-N\}$ is a unique ergodic set. If $a^b(-N) > 0$, define $\bar{j}$ to be the smallest positive state such that $p^{\bar{j}}(a^b(\bar{j})) = 0$, if such a state exists, and set $\bar{j} = M$ otherwise. Then $\{-N, \dots, -1, 1, \dots, \bar{j}\}$ is the unique ergodic set of the Markov chain. Since $p^{-N}(a^b(-N)) < 1$, it is also a regular ergodic set. Therefore, the Markov chain will have a unique invariant distribution λ^b, with $\lambda^b(j) > 0$ for all ergodic states j. We can then set $v^b(G) = \sum_{j \in \mathcal{Z}} \lambda^b(j) v^b(j)$, and $s^b = \lambda^b \cdot b/(1-\delta)$. □

The elements the steady state vectors corresponding to the transfer vectors b are continuous as a function of b:

LEMMA 4. *v^b, a^b, and λ^b are continuous in b.*

Proof: Let $T^b : \mathbb{R}^{\mathcal{Z}} \to \mathbb{R}^{\mathcal{Z}}$ denote the operator that defines the Bellman equation of the dynamic programming problem defined by the policy b. Thus, for every $v \in \mathbb{R}^{\mathcal{Z}}$,

$$T^b(v)(j) := \sup_{a'} \left\{ (1-\delta) \left[u(c^b(j)) - \phi(a') \right] + \delta \mathbb{E}_{p^j(a')}[v(i)] \right\}.$$

[5] A special case where this result applies is the case of *autarky*, where the government does not run a public insurance program, i.e., where b is equal to zero.

[6] See Kemeny and Snell [9] for the terminology and results on finite Markov chains which are used in the remainder of the proof.

Standard results from dynamic programming imply that every T^b is a contraction mapping, and is therefore continuous and has a unique fixed point v^b. By Theorem 2 in Fort [5], a unique fixed point must be an *essential fixed point*.[7] To show that v^b is continuous in b, it remains to show that for every $\varepsilon > 0$ there exists a $\zeta > 0$ such that $d^{sup}(T^b, T^{b'}) < \varepsilon$ whenever $d(b, b') < \zeta$, where d denotes the sup-metric on $\mathbb{R}^{\mathcal{Z}}$. But this follows from the continuity of the utility function u, and from the fact that given any vector $v \in \mathbb{R}^{\mathcal{Z}}$, we have $T^b(v)(j) - T^{b'}(v)(j) = (1-\delta)[u(c^b(j)) - u(c^{b'}(j))]$.

Since v^b is continuous in b and the objective function in the corresponding dynamic programming problem is strictly concave in a, the Maximum Theorem implies that a^b is continuous in b.

λ^b is the unique fixed point of the linear map on $\Delta(\mathcal{Z})$ defined by $P(a^b)$, and is therefore an essential fixed point for every b. Since the functions p^j are continuous in actions, continuity of λ^b follows by a similar argument to the one used to prove the continuity of v^b. □

If we fix a transfer vector b, and consider any initial distribution over states λ_0, we can characterize the evolution of the system induced by b and λ_0. Let $(b, a^b, s^b, \lambda^b, v^b)$ be the unique steady state vector corresponding to b, as defined in Proposition 3. Then the distribution over states at time t is $\lambda_t = \lambda_0[P(a^b)]^t$, and the subsidy required by the government to implement the fixed policy b is given by

$$s_0^b(\lambda_0) := \sum_{t=0}^{\infty} \delta^t \left[\lambda_0[P(a^b)]^t\right] \cdot b.$$

Note that this formulation allows the government to also borrow funds, as long as it can repay them using future proceeds from the insurance scheme. This constraint is well defined, as all future surpluses and deficits are deterministic.

Since the value of the policy b to an agent in state j is given by $v^b(j)$, the social welfare induced by b and the initial distribution λ_0 can be calculated as $v_G^b(\lambda_0) := \sum_{j\in\mathcal{Z}} \lambda_0(j) v^b(j)$. Therefore, if the government must choose among all unconstrained feasible fixed policies, it will choose the policy b that solves $\max_{b'} v_G^{b'}(\lambda_0)$ subject to $s_0 \geq s_0^{b'}(\lambda_0)$, where s_0 denotes the actual subsidy to the insurance scheme. Call a policy that solves this maximization problem an *optimal fixed policy for* (s_0, λ_0).

PROPOSITION 4. *There exists an optimal fixed policy for every* (s_0, λ_0).

[7] A fixed point y^* of a continuous function $f : Y \to Y$ defined on a compact metric space Y is an *essential fixed point*, if for every $\xi > 0$ there exists a $\varepsilon > 0$, such that every continuous function $g : Y \to Y$ with $d^{sup}(f, g) < \varepsilon$ has a fixed point y^{**} such that $d(y^*, y^{**}) < \xi$.

Proof: Lemma 4 implies that the problem of finding an optimal fixed policy involves maximizing a continuous function over a compact set. Existence is then a consequence of Weierstrass' Theorem. □

A policy b which is optimal for (s_0, λ_0), may not be optimal one period later, so the choice of an optimal fixed policy b will in general be time inconsistent, unless $\lambda_0 = \lambda^b$ and $s_0 = s^b$. In discussing fixed policies, we have assumed that commitment to such a policy is possible for the government. Even when such a commitment is possible, the government may want to reconsider its commitment unless $\lambda_0 = \lambda^b$ and $s_0 = s^b$. Notice however that for every λ_0, the fact that each b defines a transition matrix $P(a^b)$ with a unique regular ergodic set, implies that the distribution over states will converge to the invariant distribution. Thus, $\lambda_0[P(a^b)]^t \to \lambda^b$ for every λ_0 and b. If the government is restricted to a fixed policy, it may want to choose as its objective to maximize only the long-run social welfare, and ignore the transition towards the invariant distribution. This can be achieved by choosing a policy b^* that solves

$$\max_{b'} v^{b'}(G) \equiv \max_{b'} \sum_{j \in \mathcal{Z}} \lambda^{b'}(j) v^{b'}(j), \text{ subject to } s_0 \geq s_0^{b'}(\lambda_0).$$

We define an *optimal steady state for* s_0 to be a steady state vector $(b^*, a^{b^*}, s^{b^*}, \lambda^{b^*}, v^{b^*})$, such that $v^{b^*}(G) = \max_{b'} v^{b'}(G)$ subject to $s_0 \geq s_0^{b'}(\lambda_0)$.

PROPOSITION 5. *An optimal steady state exists for every* s_0.

Proof: Existence follows again by Weierstrass' Theorem, after applying Lemma 4. □

5. Further Remarks

In the proposed form, our dynamic public insurance game can be directly interpreted as a game between a public unemployment insurance provider and a population of workers with heterogeneous durations of employment or unemployment. With straightforward modifications to per-period payoff functions (e.g. allowing utilities to also depend on private-agent states) and gradual upward transitions from negative health states and vice-versa, one can easily port the framework to describe and analyze a public health insurance game.

To investigate further the properties of equilibrium strategies, we will have to resort to numerical computations. We can recursively approximate the correspondence-valued monotone operator defined by the SSE using convex polytopes (e.g. Sleet and Yeltekin [15], Judd et al. [7], Cronshaw [4]). Given numerical approximations of these value correspondences, we can construct the supporting equilibrium strategies. One computational drawback of the model is the exponential relationship between the payoff-space dimension and the "volume" of the model's natural state space. We defer the computation of equilibria to future generations of cheap computing power.

TIMOTHY KAM, *School of Economics & C.A.M.A., H.W. Arndt Building 25a, The Australian National University, ACT 0200, Australia. Email:* timothy.kam@anu.edu.au.

RONALD STAUBER, *School of Economics, H.W. Arndt Building 25a, The Australian National University, ACT 0200, Australia. Email:* ronald.stauber@anu.edu.au.

REFERENCES

[1] D. Abreu. On the theory of infinitely repeated games with discounting. *Econometrica*, 56(2):383–396, 1988.

[2] D. Abreu, D. Pearce, and E. Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):1041–1063, 1990.

[3] A. Atkeson. International lending with moral hazard and risk of repudiation. *Econometrica*, 59(4):1069–1089, 1991.

[4] M. Cronshaw. Algorithms for finding repeated game equilibria. *Computational Economics*, 10:139–168, 1997.

[5] M. K. J. Fort. Essential and non essential fixed points. *American Journal of Mathematics*, 72(2):315–322, 1950.

[6] H. A. Hopenhayn and J. P. Nicolini. Optimal unemployment insurance. *Journal of Political Economy*, 105(2):412–438, 1997.

[7] K. Judd, Ş. Yeltekin, and J. Conklin. Computing supergame equilibria. *Econometrica*, 71(4):1239–1254, 2003.

[8] K. L. Judd. The law of large numbers with a continuum of IID random variables. *Journal of Economic Theory*, 35(1):19–25, 1985.

[9] J. G. Kemeny and J. L. Snell. *Finite Markov Chains*. Springer, 1983. originally published by Van Nostrand, 1960.

[10] G. Mailath and L. Samuelson. *Repeated Games and Reputations: Long Run Relationships*. Oxford University Press, New York, 2006.

[11] N. Pavoni. Optimal unemployment insurance, with human capital depreciation, and duration dependence. University College London, 2004.

[12] C. Phelan and E. Stacchetti. Sequential equilibria in a Ramsey tax model. *Econometrica*, 69(6):1491–1518, 2001.

[13] C. Phelan and R. M. Townsend. Computing multi-period, information-constrained optima. *Review of Economic Studies*, 58(5):853–881, 1991.

[14] S. Shavell and L. Weiss. The optimal payment of unemployment insurance benefits over time. *Journal of Political Economy*, 87:1347–1362, 1979.

[15] C. Sleet and Ş. Yeltekin. On the computation of value correspondences. KGMS-MEDS, Northwestern University, 2000.

[16] S. E. Spear and S. Srivastava. On repeated moral hazard with discounting. *Review of Economic Studies*, 54(4):599–617, 1987.

[17] H. Uhlig. A law of large numbers for large economies. *Economic Theory*, 8(1): 41–50, 1996.

[18] C. Wang and S. D. Williamson. Uneployment insurance with moral hazard in a dynamic economy. *Carnegie-Rochester Conference Series on Public Policy*, 44:1–41, 1996.

[19] I. Werning. Optimal unemployment insurance with unobservable savings. MIT, 2003.

[20] R. Zhao. The optimal unemployment insurance contract: Why a replacement ratio? University of Illinois at Urbana-Champaign, 2001.