CrossMark

On the Computation of Value Correspondences for Dynamic Games

Christopher Sleet · Şevin Yeltekin

Abstract Recursive game theory provides theoretic procedures for computing the equilibrium payoff or value sets of repeated games and the equilibrium payoff or value correspondences of dynamic games. In this paper, we propose and implement outer and inner approximation methods for equilibrium value correspondences that naturally occur in the analysis of dynamic games. The procedure utilizes set-valued step functions. We provide an application to a bilateral insurance game with storage.

1 Introduction

Abreu et al. [1,2] (henceforth, APS) provide set-valued dynamic programming techniques for solving *repeated games*. They give conditions such that the set of sequential equilibrium payoffs in a repeated game is a fixed point of an operator analogous to the Bellman operator in dynamic programming. This operator is monotone on a space of compact sets and is the basis of an iterative procedure for calculating the equilibrium payoff set. The approach of APS can be extended to cover a large class of *dynamic games* that arise naturally in industrial organization, macroeconomics, and public finance.[1] In the dynamic case, the object of interest is a value correspondence that maps physical state variables to sets of equilibrium payoffs. In this setting, numerical implementation of the APS approach requires approximation of (candidate value) correspondences. Our paper provides such a scheme that is applicable to

---

[1] See Atkeson [3], Phelan and Stacchetti [14], Sleet [15], and Fernandez-Villaverde and Tsyvinski [10] for applications to international lending and macroeconomic policy.

C. Sleet (✉) · Ş. Yeltekin
Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA 15213, USA
e-mail: csleet@andrew.cmu.edu

Ş. Yeltekin
e-mail: sevin@andrew.cmu.edu

Birkhäuser

games with uncountable endogenous states and integrates it into an APS-type algorithm for computing equilibrium value correspondences.

As a first step, following Cronshaw and Luenberger [8] and Judd et al. [12]'s analyses of repeated games, we convexify a dynamic game by introducing a public randomization device. We approximate candidate value correspondences with simple (convex valued) step correspondences. We describe how to construct sequences of such correspondences that contain or are contained by the approximated candidate value correspondence. Integration of these approximation procedures into an APS-type algorithm permits computation of bounding *inner* and *outer* approximations to the equilibrium payoff correspondence. Since the latter lies between these approximations, the difference between them provides us with an error bound for gauging the accuracy of the procedure.[2] This approach improves on both discretization and linear interpolation methods that have been used elsewhere.[3] The former is non-parsimonious and is vulnerable to the curse of dimensionality, the latter delivers neither an inner nor an outer approximation and lacks error bounds. To illustrate the value of our methods, we briefly describe an application of them to a bilateral insurance game with storage.

1.1 Literature

Cronshaw and Luenberger [8] provide an algorithm for calculating the payoffs of strongly symmetric equilibria in convexified repeated games. Judd et al. [12] give a numerical implementation of the APS iterative procedure for such games. Their procedure permits computation of all sub-game perfect equilibrium payoffs. Baldauf et al. [5] extend the approach of Judd, Yeltekin, and Conklin to dynamic games with a finite number of (endogenous) states. Our procedure can accommodate an uncountable number of endogenous states. Many alternative numerical methods exist for characterizing the equilibria of dynamic games. These alternatives provide direct characterization of specific sub-game perfect equilibria in certain settings. Examples include Herings and Peeters [11] who provide a homotopy method for computing stationary equilibria in dynamic games and Balbus et al. [4] who show how to calculate extremal Markov equilibria in stochastic games with complementarities. We mention also Feng et al. [9] who use a similar approach to us to characterize equilibria in economies with distortions and taxes.

2 Recursive Dynamic Games and the APS Algorithm

In this section, we lay out a class of recursive dynamic games and describe the theoretic application of the APS algorithm to this class.

2.1 The Game

A group of infinitely lived agents $i = 1, \ldots, I$ play a dynamic game. Agents enter successive periods $t = 0, 1, 2, \ldots$ with an inherited pair of state variables $(s_t, k_t) \in S \times K$, where S is

---

[2] The usefulness of inner and outer approximations in providing error bounds was first emphasized by Judd et al. [12] in the context of repeated games.

[3] See Chang [7] for an application of discrete methods to a repeated policy game. See Phelan and Stacchetti [14] for an application of linear interpolation methods to a dynamic policy game.

finite and $K \subset \mathbb{R}^N$ is compact.[4] The former is a shock and evolves according to a Markov chain with transition Π; the latter is an endogenous state whose evolution is defined below. On entering period t, the agents simultaneously select a profile of actions $a_t = \{a_t^i\}_{i=1}^I$, with each agent i choosing his or her action a_t^i from a compact set $A^i(s_t, k_t) \subset \mathbb{R}^D$. Let $A = \times_{i=1}^I A^i$ denote the feasibility correspondence for actions. The inherited state k_t and the action profile a_t induce a successor state $k_{t+1} = f(s_t, k_t, a_t)$, with each $f(s, \cdot, \cdot) :$ Graph $A(s) \to K$ a continuous function.[5]

Each agent has a per-period utility function $u^i :$ Graph $A \to \mathbb{R}$, with each $u^i(s, \cdot, \cdot) :$ Graph $A(s) \to \mathbb{R}$ a continuous function. Let $\overline{V} : S \times K \rightrightarrows \mathbb{R}^I$ denote the feasible payoff correspondence:

$$\overline{V}(s, k) = \{w \in \mathbb{R}^I : w^i = u^i(s, k, a), a \in A(s, k), i = 1, \cdots I\}.$$

The graph of $\overline{V}$ is clearly compact.

At $t = 0$, the history of the game h_0 is just the initial aggregate state: (s_0, k_0). For $t > 0$, a history, h_t, is a sequence $\{s_0, k_0, \{a_\tau, s_{\tau+1}, k_{\tau+1}\}_{\tau=0}^{t-1}\}$, with each $k_{\tau+1} = f(s_\tau, k_\tau, a_\tau)$ and $a_\tau \in A(s_\tau, k_\tau)$. Let H_t denote the set of feasible t period histories from a given (s_0, k_0). The ith player values processes for states and actions according to:

$$(1 - \beta) \sum_{\tau=0}^{\infty} \beta^\tau E[u_i(s_\tau, k_\tau, a_\tau)], \tag{1}$$

where $\beta \in (0, 1)$ is the common discount factor among agents. A pure strategy for agent i is a sequence of functions $\{\sigma_t^i\}_{t=0}^\infty$ with $\sigma_t^i : H_t \to \mathbb{R}^D$ and $\sigma_t^i(h_{t-1}, a_{t-1}, s_t, k_t) \in A^i(s_t, k_t)$.

Let $V^* : S \times K \rightrightarrows \mathbb{R}^I$ denote the sub-game perfect equilibrium payoff or *value* correspondence for this game. We assume that V^* is non-empty-valued. Let $\mathscr{V}$ denote the set of all non-empty-valued correspondences $V : S \times K \rightrightarrows \mathbb{R}^I$ such that i) Graph V is compact and ii) it is contained within Graph $\overline{V}$. By an elaboration of arguments in APS, V^* is an element of $\mathscr{V}$. Since each element $V \in \mathscr{V}$ can be identified with its graph, $\mathscr{V}$ can be identified with the collection of non-empty, compact subsets of Graph $\overline{V} \subset S \times K \times \mathbb{R}^I$. We furnish $\mathscr{V}$ with the set inclusion order ($V_1 \geq$ (resp. $\leq$)V_2 if Graph $V_1 \subseteq$ (resp. $\supseteq$) Graph V_2) and the metric:

$$d_V(\text{Graph } V_1, \text{Graph } V_2) = \max_S d_H(\text{Graph } V_1(s), \text{Graph } V_2(s))$$

where d_H is the Hausdorff metric

$$d_H(\text{Graph } V_1(s), \text{Graph } V_2(s)) =$$

$$\max \left\{ \max_{x \in \text{Graph } V_1(s)} \min_{y \in \text{Graph } V_2(s)} d(x, y), \max_{y \in \text{Graph } V_2(s)} \min_{x \in \text{Graph} V_1(s)} d(x, y) \right\},$$

and $d((k, v), (k', v'))$ is the Euclidean metric on $\mathbb{R}^{N+I}$. $(\mathscr{V}, \geq, d_V)$ is a complete, partially ordered metric space.

2.2 Correspondence-valued Dynamic Programming

The recursive approach of APS utilizes a multi-player version of the principle of optimality: each sub-game perfect equilibrium payoff vector is supported by a profile of actions consistent

---

[4] Thus, we allow for an uncountable number of endogenous states. The restriction to a finite set S avoids measure-theoretic complications.

[5] For a correspondence $C : S \times K \rightrightarrows X$, Graph $C(s) = \{(k, x) : x \in C(s, k)\}$ and Graph $C = \{(s, k, x) : x \in C(s, k)\}$.

with Nash play in the current period and a vector of continuation payoffs that are themselves sub-game perfect. In analogy with the Bellman operator of dynamic programming, this observation motivates the introduction of the operator $B : \mathscr{V} \to \mathscr{V}$:

$$B(V)(s, k)$$
$$= \left\{ v \left| \begin{array}{l} v'(s') \in V(s', f(s, k, a)), a \in A(s, k), \text{ and for each } i \in 1, \cdots, I, \\ v^i = (1 - \beta)u^i(s, k, a) + \beta \sum_S v^{i\prime}(s')\Pi(s'|s) \\ v^i \geq \max\limits_{A^i(s,k)} \min\limits_{\{V(s', f(s,k,a^{-i},a'))\}} (1 - \beta)u^i(s, k, a^{-i}, a') + \beta \sum_S \tilde{v}^i(s')\Pi(s'|s) \end{array} \right. \right\}.$$

A value $v = \{v^i\}_{i=1}^I$ is in $B(V)(s, k)$ if two conditions are satisfied. First, there is an action profile $a \in A(s, k)$ and continuation payoffs $v'(s') = \{v^{i\prime}\}(s') \in V(s', f(s, k, a))$ such that each player i receives v^i when the current period action is a and the player's continuation payoff is $v^{i\prime}$:

$$v^i = (1 - \beta)u^i(s, k, a) + \beta \sum_S v^{i\prime}(s')\Pi(s'|s). \qquad (2)$$

Second, each player is better off adhering to a^i in the current period and receiving the continuation payoff $v^{i\prime}$ than deviating and receiving the worst feasible continuation payoff consistent with V:

$$v^i \geq \max\limits_{A^i(s,k)} \min\limits_{\{V(s', f(s,k,a^{-i},a'))\}} (1 - \beta)u_i(s, k, a^{-i}, a') + \beta \sum_S \tilde{v}^i(s')\Pi(s'|s). \qquad (3)$$

Inequality (3) will be referred to as the incentive compatibility constraint.

Remark 1 In some problems, it may be possible to find:

$$\underline{v}^i(s, k) := \max\limits_{A^i(s,k)} \min\limits_{\{V(s', f(s,k,a^{-i},a'))\}} (1 - \beta)u_i(s, k, a^{-i}, a') + \beta \sum_S \tilde{v}^i(s')\Pi(s'|s). \qquad (4)$$

directly. In these cases, this fact may be exploited by substituting $\underline{v}^i(s, k)$ directly into the incentive compatibility constraint:

$$v^i \geq \underline{v}^i(s, k). \qquad (5)$$

The modified incentive constraint (5) may then be incorporated into a (modified) B-operator. We do so in Sect. 4.

For a game with a finite number of states (and compact action sets), Baldauf et al. [5] extend arguments in APS to show that the equilibrium payoff correspondence V^* is a fixed point of B. In addition, they show that B is monotone (in the set inclusion order), i.e., $V_1 \geq V_2$ implies $B(V_1) \geq B(V_2)$ and a self-map on $\mathscr{V}$, i.e., $B(\mathscr{V}) \subset \mathscr{V}$. It follows that for $j = 1, 2, \ldots, \overline{V} \geq B(\overline{V}) \geq B^j(\overline{V}) \geq B^{j+1}(\overline{V}) \downarrow \cap_{j=1}^\infty B^j(\overline{V}) := B^\infty(\overline{V})$. Arguments in Baldauf et al. [5] also show that $B^\infty(\overline{V}) = V^*$, thus repeated application of B to $\overline{V}$ (and, by similar arguments, any set V in the interval $V^* \leq V \leq \overline{V}$) generates a sequence of sets that contain V^* and converge monotonically to it in the Hausdorff metric. Although, Baldauf et al. [5] prove these results for games with a finite number of states, they are readily extended to those with a finite number of exogenous shocks and a compact endogenous state space (in $\mathbb{R}^N$), the case considered in this paper. In addition, the monotonicity of B and the fact that V^* is a fixed point of B can also be exploited to show that if $V \leq B(V) \leq V^*$, then $V \leq B(V) \leq B^j(V) \leq B^{j+1}(V) \uparrow \cup_{j=1}^\infty B^j(V) := B^\infty(V) \leq V^*$.

2.3 A Recursive Game with Lotteries

To facilitate computation, we "convexify" the preceding game via the introduction of lotteries. Specifically, at the beginning of each period, a lottery is held over a set Ω. Strategies condition player choices on histories of lottery outcomes as well as histories of states and actions. The lottery probabilities are determined by a profile of functions that map histories into probability distributions on Ω. These and the lottery outcomes are common knowledge. It is readily shown that this formulation's equilibrium payoff or value correspondence V^{L*} is convex valued.[6] Let $\mathscr{V}^L$ denote the set of all correspondences $V : S \times K \rightrightarrows \mathbb{R}^I$ such that (1) Graph V is compact, (2) Graph V is convex valued, and (3) Graph V is contained within Graph $\overline{V}^L$, where for each $(s, k) \in S \times K$, $\overline{V}^L(s, k) := \text{co } \overline{V}^L(s, k)$ and co X denotes the convex hull of a set X. Define the operator $B^L : \mathscr{V}^L \rightarrow \mathscr{V}^L$ according to: $B^L(V)(s, k) = \text{co } B(V)(s, k)$. It is not hard to show that analogs of the results in the previous subsection go through with B^L replacing B and $\mathscr{V}^L$ replacing $\mathscr{V}$. Specifically, (1) V^{L*} is a fixed point of B^L, (2) repeatedly applying B^L to $\overline{V}^L$ induces a sequence of convex-valued correspondences that converge monotonically from above to V^{L*}, and (3) repeatedly applying B^L to $V \leq B^L(V) \leq V^{L*}$ induces a sequence of convex-valued correspondences that converge monotonically from below to $B^{L\infty}(V) \leq V^{L*}$.

3 Approximating Value Correspondences

For a *repeated game without state variables*, the introduction of lotteries implies that the equilibrium payoff set is convex. Furthermore, (the analog of) B^L can be restricted to a domain of compact and convex sets. Well-known "outer" and "inner" approximation procedures exist for such sets. Specifically, an outer approximation can be formed from the intersection of a finite number of closed half-spaces that contain the convex set; an inner approximation from the convex hull of a finite number of points contained by the set. The outer approximation provides an upper bound for the set; the inner approximation a lower bound. Together they provide an error bound. Judd et al. [12] integrate such approximations into the monotone iterative scheme described in the previous section and, hence, create practical algorithms for calculating the equilibrium payoff sets of repeated games.

In the setting of *dynamic games with state variables*, the situation is more complicated. Lotteries ensure that the equilibrium value correspondence V^{L*} is convex valued (which is useful), but not that the sets Graph $V^{L*}(s)$ are convex. Consequently, implementation of the theoretic algorithms described in the previous section typically requires the approximation of a sequence of non-convex sets (i.e., the graphs of value correspondences).

To fix language and notation, we introduce some general classes of approximation procedure. We then turn to the specific procedures that we use to represent correspondences.

3.1 Approximation Schemes

We introduce the following general definition of an approximation scheme.

Definition 1 (Approximation scheme). Let $(\mathscr{W}, \geq)$ be a collection of non-empty, compact, convex-valued correspondences $W : S \times K \rightrightarrows \mathbb{R}^I$ with common domain $S \times K$, $K \subset \mathbb{R}^N$ compact, and partially ordered by set inclusion. An *approximation rule* is a pair $(F, \hat{\mathscr{W}})$ with

---

6   However, each Graph $V^{L*}(s)$ is not generally convex.

$\hat{\mathcal{W}} \subset \mathcal{W}$ and $F : \mathcal{W} \to \hat{\mathcal{W}}$. An *approximation scheme* is a sequence of approximation rules $\{F_j, \hat{\mathcal{W}}_j\}_{j=0}^{\infty}$.

1. An approximation rule $(F, \hat{\mathcal{W}})$ is finitely parameterized if $\hat{\mathcal{W}} = g(T)$, for a set of parameters $T \subset \mathbb{R}^P$ and a one-to-one mapping $g : T \to \hat{\mathcal{W}}$ that associates each parameter value with a unique correspondence. An approximation scheme $\{F_j, \hat{\mathcal{W}}_j\}_{j=0}^{\infty}$ is finitely parameterized if each approximation rule $(F_j, \hat{\mathcal{W}}_j)$ is finitely parameterized.
2. An approximation rule $(F, \hat{\mathcal{W}})$ is monotone if for $W_1, W_2 \in \mathcal{W}$, $W_1 \geq W_2$ implies $F(W_1) \geq F(W_2)$. A monotone approximation scheme is an approximation scheme composed of monotone approximation rules.
3. An approximation rule $(F, \hat{\mathcal{W}})$ is outer (resp. inner) if for all $W \in \mathcal{W}$, $F(W) \geq W$ (resp. $W \geq F(W)$). An outer (inner) approximation scheme is an approximation scheme composed of outer (inner) approximation rules.
4. An approximation scheme $\{F_n, \hat{\mathcal{W}}_j\}_{j=0}^{\infty}$ is pointwise (resp. uniformly) convergent if, $\forall W \in \mathcal{W}$, the sequence $F_j(W)$ converges to W pointwise (resp. uniformly) in the metric d_V.

Computation of the convexified equilibrium value correspondence V^{L*} requires the integration of a monotone iterative scheme with a practical finitely parameterized approximation procedure. Outer and inner approximations that respect the monotonicity of the iterative scheme are especially useful. Identifying $\mathcal{W}$ with $\mathcal{V}^L$, let $F : \mathcal{V}^L \to \hat{\mathcal{V}}^L$ be an approximation rule (with range $\hat{\mathcal{V}}^L$). Then $\hat{B} = F \circ B^L$ maps $\hat{\mathcal{V}}^L$ into itself. If F is monotone, then $\hat{B}$ is also monotone and if $\hat{V} \geq \hat{B}(\hat{V})$, $\hat{V} \in \hat{\mathcal{V}}^L$, then $\{\hat{B}^j(\hat{V})\}_{j=1}^{\infty}$ is a decreasing and pointwise convergent sequence of correspondences in $\hat{\mathcal{V}}^L$. If, in addition, F is a monotone outer approximation rule and $\hat{V} \geq V^{L*}$, then for all j, $V^{L*} \leq F(V^{L*}) = \hat{B}(V^{L*}) \leq \hat{B}^j(\hat{V}) \leq \ldots \leq \hat{B}(\hat{V}) \leq \hat{V}$. Hence, the limit of an iteration of $\hat{B}$ from $\hat{V}$ is an outer approximation of V^{L*}. Similarly, if $\hat{V} \in \hat{\mathcal{V}}^L$, $\hat{V} \leq V^{L*}$, $\hat{V} \leq \hat{B}(\hat{V})$ and F is a monotone inner approximation, the limit of $\{\hat{B}^j(\hat{V})\}$ is an inner approximation to V^{L*}.

3.2 Block-Specific and Step Correspondences

As a precursor to developing approximation schemes for value correspondences, we introduce certain simple correspondences which we will use to construct approximations. Assume that $K = \times_{n=1}^{N}[\underline{k}_n, \overline{k}_n] \subseteq \mathbb{R}^N$ and that each interval $[\underline{k}_n, \overline{k}_n]$ is subdivided into M_n subintervals $\{[k_{n,m}, k_{n,m+1}]\}_{m=0}^{M_n}$, $\underline{k}_n = k_{n,0} < k_{n,1} < \cdots < k_{n,M_n} = \overline{k}_n$. The Cartesian product of these subintervals defines a partition $\mathcal{K}$ of K, with $\mathcal{K} = \{K_q\}_{q=1}^{Q}$, $Q = \prod_{n=1}^{N} M_n$ and $K_q = \prod_{n=1}^{N}[k_{n,m_q}, k_{n,m_q+1}]$. $K_q \in \mathcal{K}$ is called a *block*.

Definition 2 (Block-specific and step correspondences). Let $\mathcal{K}$ be a partition of $K = \times_{n=1}^{N}[\underline{k}_n, \overline{k}_n]$. A correspondence $W : K \rightrightarrows \mathbb{R}^I$ is *block specific* if it vanishes (is empty-valued) or is equal to all of $\mathbb{R}^I$ outside of a specific block in $\mathcal{K}$. W is a *step correspondence* if it is constant on the interior of each block in $\mathcal{K}$.

The following approximation methods rely on correspondences constructed from finite numbers of block-specific or step correspondences.

3.3 An Outer Approximation Method

This method approximates the underlying correspondence with a collection of convex-valued step correspondences. The procedure is as follows. Let $V : S \times K \rightrightarrows \mathbb{R}^N$ be the underlying correspondence. Assume that each $V(s, \cdot)$ is upper hemi-continuous and non-empty, compact and convex valued. Fix a partition $\mathcal{K} = \{K_q\}$ of K and define the family of convex-valued step correspondences $\omega_{s,q}$ according to:

$$\omega_{s,q}(k) = \begin{cases} \text{co } \cup_{k' \in K_q} V(s, k') & \text{if } k \in K_q, \\ \emptyset & \text{otherwise.} \end{cases}$$

Each $\omega_{s,q}$ is completely described by the convex set co $\cup_{k' \in K_q} V(s, k')$ (and the partition element K_q). Consequently, it can be approximated by the outer approximation procedure of Judd et al. [12]. More precisely, let $\{v_\ell(s, q) | \ell = 1, \ldots, \bar{\ell}\}$ be a set of $\bar{\ell}$ points on the boundary of co $\cup_{k' \in K_q} V(s, k')$; let $\{z_1, \ldots, z_{\bar{\ell}}\} \subset \mathbb{R}^{I\bar{\ell}}$ be a corresponding set of normals oriented so that $(v_\ell(s, q) - v) \cdot z_\ell \geq 0$ for all $v \in$ co $\cup_{k' \in K_q} V(s, k')$. The outer approximation $\omega_{s,q}^{\mathbb{O}}(k)$ is given by:

$$\omega_{s,q}^{\mathbb{O}}(k) = \begin{cases} \cap_{\ell=1}^{\bar{\ell}} \{v \mid z_\ell \cdot v \leqq c_\ell(s, q)\}, & \text{if } k \in K_q, \\ \emptyset & \text{otherwise.} \end{cases}$$

where for each ℓ, $c_\ell(s, q) = z_\ell \cdot v_\ell(s, q)$.[7]

The outer approximation to V is then defined according to $V^{\mathbb{O}}(s, k) = \cup_q \omega_{s,q}^{\mathbb{O}}(k)$. Using this scheme, an approximating correspondence $V^{\mathbb{O}}$ is described by a collection $\{c_\ell(s, q)\}$ in $\mathbb{R}^{SQI\bar{\ell}}$, with each $c_\ell(s, q)$ determining the position of the ℓ-th bounding hyperplane on the (s, q)-th block.

Let $\{V_j^{\mathbb{O}}\}_{j=0}^{\infty}$ be a sequence of such outer approximations to V. Let $\mathcal{K}_j$ denote the partition associated with the j-th approximation and let $\omega_{s,q,j}^{\mathbb{O}}$ denote the (s, q)-th block-specific correspondence of the j-th approximation. Assume that $\mathcal{K}_j$ is constructed from 2^j uniformly distributed blocks and that for a fixed decreasing sequence $\{\varepsilon_j\}$ with $\lim_{j \to \infty} \varepsilon_j = 0$, the approximations $\omega_{s,q,j}^{\mathbb{O}}$ are constructed such that $d_H(\omega_{s,q,j}^{\mathbb{O}}, \omega_{s,q,j}) < \varepsilon_j$ for $q = 1, \cdots, 2^j$ (see Appendix 5 for details).

Theorem 1 *The sequence $\{V_j^{\mathbb{O}}\}_{j=0}^{\infty}$ is a pointwise convergent, monotone outer approximation scheme. Each $V_j^{\mathbb{O}}(s, \cdot)$ is upper hemi-continuous and non-empty and compact valued.*

Proof Upper hemi-continuity of each $\omega_{s,q,j}^{\mathbb{O}}$ is immediate. The upper hemi-continuity of $V_j^{\mathbb{O}}(s, \cdot)$ then follows since $V_j^{\mathbb{O}}(s, \cdot)$ is equal to the union of a finite number of upper hemi-continuity correspondences. Each $V_j^{\mathbb{O}}(s, k)$ is bounded by construction. For the closed valuedness, see Beer [6]. Non-emptiness follows from $\emptyset \neq V(s, k) \subset V_j^{\mathbb{O}}(s, k)$. Fix $\varepsilon > 0$ and a value $v \in \mathbb{R}^I$. Let $\mathcal{O}_\varepsilon(X)$ denote the union of all ε balls around points in $X \subset \mathbb{R}^I$. By the upper hemi-continuity of $V(s, \cdot)$, there is a $\lambda > 0$ such that: $\cup \{V(s, k') : \|k - k'\| < \lambda\} \subset \mathcal{O}_\varepsilon(V(s, k))$. Now, there is a j_1 such that for all $j > j_1$, the diameter of each block in the j-th partition is less than λ. Additionally, by construction, there is a j_2 such that for all $j > j_2$, $d_H(\omega_{s,q,j}^{\mathbb{O}}, \omega_{s,q,j}) < \varepsilon$, $\forall q = 1, \cdots, 2^j$. Fix

---

[7] Geometrically, for $k \in K_q$, $\omega_{s,q}^{\mathbb{O}}(k)$ equals the intersection of a finite number of half-spaces containing and with hyperplanes tangent to co $\cup_{k' \in K_q} V(s, k')$.

$j > \max(j_1, j_2)$ and let $Q(k)$ denote the set of block indices q such that $k \in K_q \in \mathcal{K}_j$. It follows that $\cup \{V(s, k') : k' \in \cup_{q \in Q(k)} K_q\} \subset \mathcal{O}_\varepsilon(V(s, k))$. Since V is convex valued, it follows that $\mathcal{O}_\varepsilon(V(s, k))$ is convex. Hence,

$$V(s, k) \subset \omega_{s,q(k),j}(k) := \text{co} \cup \{V(s, k') : k' \in \cup_{q \in Q(k)} K_q\} \subset \mathcal{O}_\varepsilon(V(s, k))$$

But since $j > j_2$, for all $q \in Q(k)$, $d_H(\omega_{s,q,j}^{\mathbb{O}}, \omega_{s,q,j}) < \varepsilon$. Hence, $V(s, k) \subset V_j^{\mathbb{O}}(s, k) = \cup_{q \in Q(k)} \omega_{s,q,j}^{\mathbb{O}}(k) \subset \mathcal{O}_{2\varepsilon}(V(s, k))$, and the approximation is convergent in d_V. Monotonicity of the approximation scheme is immediate from the definition. $\qquad\square$

In conjunction with the previous result, the following result of Beer [6] shows that a uniformly convergent approximation scheme can be obtained if the correspondence to be approximated is lower hemi-continuous and non-empty and compact valued.

Theorem 2 *(Beer [6]). Suppose that each $V(s, \cdot) : K \rightrightarrows \mathbb{R}^I$ is a lower hemi-continuous, compact and non-empty-valued correspondence. Let $\{V_j(s, \cdot)\}_{j=1}^\infty$ be a sequence of upper hemi-continuous, compact-valued correspondences such that for each $k \in K$ and j, $V(s, k) \subseteq V_{j+1}(s, k) \subseteq V_j(s, k)$. If $\cap_{j=1}^\infty V_j(s, k) = V(s, k)$ for each k, then $V_j(s, \cdot)$ converges to $V(s, \cdot)$ uniformly in the Hausdorff metric, d_H.*

Under the assumptions made in Sect. 2, the B^L-operator preserves continuity. Thus, the equilibrium value correspondences $V^{L*}(s, \cdot)$ can be obtained as the limit of a sequence of continuous correspondences. It is the elements of this sequence that must be approximated. Theorem 2 ensures that there are uniformly convergent approximation schemes available for this purpose.

3.4 The Inner Approximation Step Correspondence Method

This section describes a method for approximating a lower hemi-continuous correspondence with collections of step correspondences whose graphs are contained by the graph of the approximant. Let $V : S \times K \rightrightarrows \mathbb{R}^I$ be the underlying correspondence with each $V(s, \cdot)$ lower hemi-continuous and with each $V(s, k)$ convex body, i.e., a compact and convex set with a non-empty interior.[8] Fix $s \in S$ and a partition $\mathcal{K}$ of K and define the correspondence $\omega_{s,q}$ according to:

$$\omega_{s,q}(k) = \begin{cases} \cap_{k' \in K_q} V(s, k') & \text{if } k \in K_q, \\ \mathbb{R}^I & \text{otherwise.} \end{cases}$$

Thus, each $\omega_{s,q}$ is completely described by the convex set $\cap_{k' \in K_q} V(s, k')$ (and the partition element K_q). Consequently, it can be approximated by the inner approximation procedure of Judd et al. [12]. More precisely, let $\{v_\ell(s, q) | \ell = 1, \ldots, \bar{\ell}\}$ be a set of $\bar{\ell}$ points in $\cap_{k' \in K_q} V(s, k')$. The inner approximation $\omega_{s,q}^{\mathbb{I}}(k)$ is given by:

$$\omega_{s,q}^{\mathbb{I}}(k) = \begin{cases} \text{co}\, \{v_\ell(s, q) | \ell = 1, \ldots, \bar{\ell}\} & \text{if } k \in K_q, \\ \mathbb{R}^I & \text{otherwise.} \end{cases}$$

The inner approximation to V is then defined according to: $V^{\mathbb{I}}(s, k) = \cap_q \omega_{s,q}^{\mathbb{I}}(k)$. Using this scheme, an approximating correspondence $V^{\mathbb{I}}$ is described by a collection of points $\{v_\ell(s, q)\}$ in $\mathbb{R}^{SQI\bar{\ell}}$ from which the convex hulls are constructed.

---

[8] The stronger assumption of a non-empty interior is needed in this case.

Let $\{V_j^{\mathbb{I}}\}_{j=0}^{\infty}$ denote a sequence of inner approximations to V. Let $\mathcal{K}_j$ denote the partition associated with the j-th approximation and define $\omega_{s,q,j}^{\mathbb{I}}$ and $\omega_{s,q,j}$ analogously to $\omega_{s,q}^{\mathbb{I}}$ and $\omega_{s,q}$ above (e.g., $\omega_{q,j}^{\mathbb{I}}$ is the q-th step correspondence of the j-th approximation). Assume that $\mathcal{K}_j$ is constructed from 2^j uniformly distributed blocks. Additionally, let $\overline{\ell}_j$ denote the number of points used in the construction of the inner approximations $\{\omega_{s,q,j}^{\mathbb{I}}\}$. Assume that $\overline{\ell}_{j+1} > \overline{\ell}_j$ and that for a fixed decreasing sequence $\{\varepsilon_j\}$ with $\lim_{j \to \infty} \varepsilon_j = 0$, each $\overline{\ell}_j$ is chosen sufficiently large that $d_H(\omega_{s,q,j}^{\mathbb{I}}, \omega_{s,q,j}) < \varepsilon_j$ for $q = 1, \ldots, 2^j$.

Theorem 3 *The sequence $\{V_j^{\mathbb{I}}\}_{j=0}^{\infty}$ is a pointwise convergent, monotone inner approximation scheme. Each $V_j^{\mathbb{I}}(s, \cdot)$ is lower hemi-continuous and compact and convex valued.*

The proof is a small adaptation of arguments on p. 15 of Beer [6].

4 Application: A Bilateral Insurance Game

This section considers a problem in which two risk averse agents accumulate a joint fund to hedge endowment risk and neither can commit not to steal from the fund. The problem extends the physical environment of Kocherlakota [13] by permitting the agents to store. The details of the model follow.

An environment is inhabited by two infinitely lived agents with identical preferences. Each period a shock $s_t \in S$ is drawn from a finite set S according to a distribution Π. For simplicity, shock draws are i.i.d. A function $\gamma : S \to \mathbb{R}_+^2$, $\gamma(s) = \{\gamma^i(s)\}$, gives the endowment of each agent in each shock state. The history of endowments received by each agent is common knowledge. Agents begin the game in a mutual risk-sharing arrangement with $k_0 \geq 0$ in a joint fund. If the arrangement is still active at t, they enter period t with an amount $k_t \geq 0$ in the fund. On observing the current endowment realization $\gamma(s_t)$, each agent transfers $\tau_t^i \in \Gamma^i(k_t) = [-\delta k_t, \gamma^i(s_t)]$, $\delta \in [0, 1/2]$, to the fund. The agents then choose to remain in the risk-sharing arrangement or walk away from it. If they walk away, any resources left in the fund are destroyed and the agents live in autarky thereafter (though with the ability to accumulate their own fund). While the relationship is active, the fund evolves according to $k_{t+1} = f(k_t + \tau_t^1 + \tau_t^2)$, with f bounded above (and below by zero). If it is dissolved, each agent separately uses the same storage technology f to accumulate their own fund. The i-th, $i = 1, 2$, agent's preferences over current consumption are described by utility functions u^i, where each u^i is a continuous, strictly concave, strictly increasing function. Each agent discounts the future with discount factor β.

At any date t, the agent is free to keep her endowment $\gamma^i(s_t)$, withdraw the maximal amount δk_t from the joint fund and live in autarky thereafter. Denote the associated payoff:

$$\underline{v}^i(s_t, k_t) := \max_{b \in [0, \gamma^i(s_t) + \delta k_t]} (1 - \beta) u^i(\gamma^i(s_t) + \delta k_t - b) + \beta v_{aut}^i(f(b)),$$

where v_{aut} gives the expected payoff to an agent in autarky as a function of her savings. It is not hard to see that $\underline{v}^i(s_t, k_t)$ gives the worst possible sub-game perfect payoff to agent i if the relationship is active and in state (s_t, k_t). The payoffs $\underline{v}^i$ may be calculated directly and substituted into the B-operator as discussed in Remark 1. The B-operator so obtained calculates equilibrium payoffs to the agents while the risk-sharing arrangement is active (and uses states sufficient for that purpose). Specifically, it is given by:
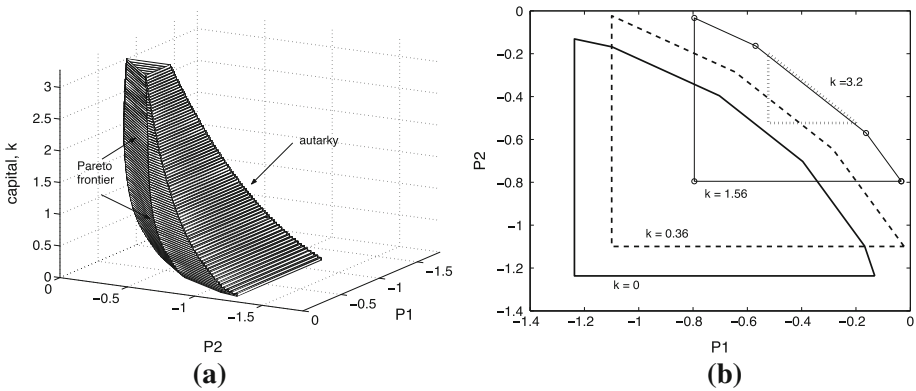
Fig. 1 Value correspondence from the insurance game. **a** Value correspondence graph **b** cross section

$$B(V)(s,k) = \left\{ (v^1, v^2) \left| \begin{array}{l} \text{for } i = 1, 2, \\ v^i = \{(1-\beta)u^i(\gamma^i(s) - \tau^i(s)) + \beta \sum_{s' \in \mathcal{S}} v^{i'}(s')\Pi(s'), \\ v' \in V(f(k + \sum_i \{s^i - \tau^i(s)\}), s) \\ (1-\beta)u^i(\gamma^i(s) - \tau^i(s)) + \beta \sum_{s' \in \mathcal{S}} v^{i'}(s')\Pi(s') \geq \underline{v}^i(s,k). \end{array} \right. \right\}$$

Lotteries may be introduced and the corresponding B^L operator obtained for the convexified game.

4.1 Calculations

The following parametric forms are assumed for our numerical example: $u^i(c) = \frac{c}{1-\sigma}^{1-\sigma}$ and $f(k) = \min(\bar{k}, k)$. The preference and technology parameters are set to: $\beta = 0.8, \sigma = 1$, $\bar{k} = 4, \gamma^i(S) = \{0.1, 0.5\}, \Pi = \{0.5, 0.5\}, \delta = 0.5$. Calculation using either the outer or inner approximation methods yields very similar results. Figure 1a below illustrates the outer-approximated value correspondence for the case of high endowment shocks for both agents. Figure 1b displays the cross section of the correspondence in Fig. 1a for a sample of k levels.

The worst equilibrium payoff function for each agent is concave in capital and is independent of the payoff received by the other agent. This functional form is, of course, inherited from the autarkic value function. As the capital level increases, autarky becomes more attractive and the set of equilibrium values decrease, but (symmetric) payoffs on the Pareto frontier increase. These changes can be clearly seen in Fig. 1b. Although incentives for leaving the insurance relationships are higher with high autarky values, high capital levels provide better insurance against aggregate shocks and higher consumption. However, in this example, for capital levels above 3.2, no risk-sharing is feasible. For $k \in (3.2, 4]$, the autarky values are too high, and thus the incentives to deviate are too strong to enable risk-sharing. There is, consequently, a limit to the amount of capital that can be accumulated in this risk-sharing arrangement. That limit depends on the parameters of the model, in particular on the amount of capital that can be grabbed by one agent.

5 Appendix

The following algorithms provide the details of the computational procedure for the outer and inner monotone approximations of $B(V)(s,k)$ using hyperplanes. The correspondence $V(s,k)$ is a candidate for the equilibrium value correspondence.

Outer Approximation Algorithm

1. Fixed Inputs:

   (a) Exogenous parameters.
   (b) Partition $\mathcal{K} = \{K_q\}$.
   (c) Search directions: $\mathcal{R} = \{r_1, \ldots, r_{\overline{\ell}_r}\}$, with each $\{r_\ell \in \mathbb{R}^I \ : \ ||r_\ell|| = 1\}$.
   (d) Convergence criterion $\varepsilon > 0$.

2. First Iteration Inputs:

   (a) Set of boundary points $\{v_\ell(s, q) | \ell = 1, \ldots, \overline{\ell}\}$ for each (s, q).
   (b) Set of corresponding normals: $\mathcal{Z} = \{z_1, \ldots, z_{\overline{\ell}}\}$.

   Together (a) and (b) define the initial outer approximation $V^{\mathbb{O}}(s, k) = \cup_q \omega_{s,q}^{\mathbb{O}}(k)$
   where

   $$\omega_{s,q}^{\mathbb{O}}(k) = \begin{cases} \cap_{\ell=1}^{\overline{\ell}} \{v \mid z_\ell \cdot v \leq z_\ell \cdot v_\ell(s, q) = c_\ell(s, q)\}, & \text{if } k \in K_q, \\ \emptyset & \text{otherwise.} \end{cases}$$

   for all $v \in \text{co} \cup_{k' \in K_q} V(s, k')$.

3. Optimization: Compute the new subgradient and boundary point sets that together represent an outer approximation to $B^L(V)$. For each $s \in S$, each partition element K_q and each $r_\ell \in \mathcal{R}$

   $$c_\ell^+(s, q) = \max_{a, v', k \in K_q} r_\ell \cdot [(1 - \beta)u(s, k, a) + \beta \sum_S v'(s')\Pi(s'|s)] \text{ s.t.} \tag{6}$$

   $(i) \ v'(s') \in V^{\mathbb{O}}(s', f(s, k, a)), \ a \in A(s, k),$

   (ii) for each $i \in \{1, \ldots, I\}$,

   $$(1 - \beta)u^i(s, k, a) + \beta \sum_S v^{i'}(s')\Pi(s'|s) \geq$$

   $$\max_{a_{-i} \in A^i(s,k)} \min_{\tilde{v}^i(s') \in V^{\mathbb{O}}(s', f(s,k,a_{-i},a'))} (1 - \beta)u^i(s, k, a_{-i}, a')$$

   $$+ \beta \sum_S \tilde{v}^i(s')\Pi(s'|s).$$

4. Output: $V^+(k, s) = \hat{B}(V)(s, k) = \cup_q \omega_{s,q}^{\mathbb{O}^+}(k)$, where

   $$\omega_{s,q}^{\mathbb{O}^+}(k) = \begin{cases} \cap_{\ell=1}^{\overline{\ell}_r} \{v \mid r_\ell \cdot v \leq c_\ell^+(s, q)\} & \text{if } k \in K_q, \\ \emptyset & \text{otherwise.} \end{cases}$$

5. Check for convergence: Stop if $d_V(V^+, V) < \varepsilon$. Else set $V^{\mathbb{O}} = V^+$, $\mathcal{Z} = R$, and $\overline{\ell} = \overline{\ell}_r$ and return to Step 3.

The main computational work is done in the optimization step. In this step, for each search direction r_ℓ and element of the partition $\mathcal{K}$, we find an action profile a, continuation values

$v'(s') \in V(s', f(s, k, a))$, and current state k in K_q which makes a incentive compatible and maximizes a weighted sum of player payoffs where the weights are given by r_ℓ.[9]

In the outer approximation case, the set of search directions, $\mathcal{R}$, and the set of hyperplane normals for the approximation of $V(s, k)$, $\mathcal{Z}$, coincide after the initial iteration. Correspondences are represented by sets of bounding hyperplanes, one set for each block in the partition and each state s. In this way, an approximated correspondence can be reduced to a collection $\{c_\ell(s, q)\}$ in $\mathbb{R}^{SQI\bar{\ell}}$, with each $c_\ell(s, q)$ determining the position of the ℓ-th hyperplane on the q-th block in state s.

For inner approximations, the set of hyperplane normals vary from iteration to iteration, depending on the number and position of the points on the convex hull of $V(s, k)$. The bounding hyperplanes are constructed from the line segments that connect the extreme points belonging to the convex hull of $V(s, k)$. In this case, an approximated correspondence cannot be reduced to a collection of $\{c_\ell(s, q)\}$ in $\mathbb{R}^{SQI\bar{\ell}}$ but can be reduced to a collection of $\{v_\ell(s, q)\}$ in $\cap_{k' \in K_q} V(s, k')$ for each (s, q). This requires a slight modification to the output and updating used in the outer approximation algorithm. Below is the step by step description of the algorithm for computing an inner approximation to $B(V)$.

_____

Inner Approximation Algorithm

1. Fixed Inputs:

   (a) Exogenous parameters
   (b) Partition $\mathcal{K} = \{K_q\}$.
   (c) Search directions: $\mathcal{R} = \{r_1, \ldots, r_{\bar{\ell}_r}\}$, with each $\{r_\ell \in \mathbb{R}^I \ : \ ||r_\ell|| = 1\}$.
   (d) Convergence criterion: $\varepsilon > 0$.

2. First Iteration Inputs:

   (a) Set of points $\{v_\ell(s, q)|\ell = 1, \ldots, \bar{\ell}\}$ in $\cap_{k' \in K_q} V(s, k')$, for each (s, q).
   These points define the initial inner approximation $V^{\mathbb{I}}(s, k) = \cap_q \omega^{\mathbb{I}}_{s,q}(k)$ where

   $$\omega^{\mathbb{I}}_{s,q}(k) = \begin{cases} \mathrm{co}\ \{v_\ell(s, q)|\ell = 1, \ldots, \bar{\ell}\} & \text{if } k \in K_q, \\ \mathbb{R}^I & \text{otherwise.} \end{cases}$$

3. Optimization: Compute the new sets of point from which the convex hulls that together represent an inner approximation to $B^L(V)$ are constructed. For each $s \in S$, each partition element K_q and each $r_\ell \in \mathcal{R}$

   $$\{a, v', k\}^+_\ell(s, q) = \underset{a, v', k \in K_q}{\mathrm{argmax}} \quad r_\ell \cdot [(1 - \beta)u(s, k, a) + \beta \sum_S v'(s')\Pi(s'|s)] \text{ s.t} \quad (7)$$

   $(i)\ v'(s') \in V^{\mathbb{I}}(s', f(s, k, a)),\ a \in A(s, k),$

   (ii) for each $i \in \{1, \ldots, I\}$,

   $$(1 - \beta)u^i(s, k, a) + \beta \sum_S v^{i'}(s')\Pi(s'|s) \geq$$

_____

$$\max_{a_{-i} \in A^i(s,k)} \min_{\tilde{v}^i(s') \in V^{\mathbb{I}}(s', f(s,k,a_{-i},a'))} (1 - \beta) u^i(s, k, a_{-i}, a')$$

$$+ \beta \sum_S \tilde{v}^i(s') \Pi(s'|s).$$

4. Output: $V^+(k, s) = \hat{B}(V)(s, k) = \cap_q \omega_{s,q}^{\mathbb{I}+}(k)$, with

$$\omega_{s,q}^{\mathbb{I}+}(k) = \begin{cases} \text{co } \{v_\ell^+(s, q) | \ell = 1, \ldots, \overline{\ell}\} & \text{if } k \in K_q, \\ \mathbb{R}^I & \text{otherwise.} \end{cases}$$

where $v_\ell^+(s, q)$ is equal to $[(1 - \beta)u(s, k, a) + \beta \sum_S v'(s')\Pi(s'|s)]$ evaluated at $\{a, v', k\}_\ell^+(s, q)$.

5. Check for convergence: Stop if $d_V(V^+, V) < \varepsilon$. Else set $V^{\mathbb{I}}(s, k) = V^+(s, k)$ and go back to Step 3.

See Judd et al. [12] and Yeltekin et al. [16] for the details of constructing convex hulls similar to those required in Step 4.

References

1. Abreu D, Pearce D, Stacchetti E (1986) Optimal cartel equilibria with imperfect monitoring. J Econ Theory 39(1): 251–269
2. Abreu D, Pearce D, Stacchetti E (1990) Toward a theory of discounted repeated games with imperfect monitoring. Econometrica 58(5):1041–1063
3. Atkeson A (1991) International lending with moral hazard and risk of repudiation. Econometrica 59:1069–1089
4. Balbus L, Reffett K, Wozny L (2014) A constructive study of markov equilibria in stochastic games with strategic complementarities. J Econ Theory 150:815–840
5. Baldauf M, Judd K, Mollner J, Yeltekin Ş (2013) Supergames with states. Stanford University, Mimeo
6. Beer G (1980) The approximation of upper semicontinuous multifunctions by step multifunctions. Pac J Math 87:11–19
7. Chang R (1998) Credible monetary policy in an infinite horizon model: recursive approaches. J Econ Theory 81:431–461
8. Cronshaw M, Luenberger D (1994) Strongly symmetric subgame perfect equilibrium in infinitely repeated games with perfect monitoring and discounting. Games Econ Behav 6:220–237
9. Feng Z, Miao J, Peralta-Alva A, Santos M (2014) Numerical simulation of nonoptimal dynamic equilibrium models. Int Econ Rev 55(1):83–110
10. Fernandez-Villaverde J, Tsyvinski A (2002) Optimal fiscal policy in a business cycle model without commitment. University of Pennsylvania, Mimeo
11. Herings PJ-J, Peeters R (2004) Stationary equilibria in stochastic games: structure, selection and computation. J Econ Theory 118:32–60
12. Judd K, Yeltekin Ş, Conklin J (2003) Computing supergame equilibria. Econometrica 71:1239–1255
13. Kocherlakota N (1996) Efficient bilateral risk sharing without commitment. Rev Econ Stud 63:595–609
14. Phelan C, Stacchetti E (2001) Subgame perfect equilibria in a ramsey taxes model. Econometrica 69:1491–1518
15. Sleet C (1996) New recursive methods for government policy problems. Stanford University, Mimeo
16. Yeltekin Ş, Judd K, Cai Y (2014) Computing equilibria of dynamic games. Working Paper. Carnegie Mellon University