Algorithms for Stochastic Games Online Appendix

Dilip Abreu Benjamin Brooks Yuliy Sannikov

January 30, 2020

A A repeated game with infinitely many extreme equilibrium payoffs

In this Online Appendix, we give additional details on the three-player example from Section 5.1, depicted in Figure 4, that has infinitely many extreme equilibrium payoffs. First we construct a self-generating set that turns out to be V. We will then argue that this is in fact the equilibrium payoff set.

A.1 The equilibrium payoff set

Recall that only four action profiles can be played in equilibrium, which induce payoffs (4, 4, 4) and permutations of (8, 8, 0). Note that (4, 4, 4) is one of the equilibrium payoffs.

We will generate two sequences of payoffs $\{u^l\}_{l=0}^{\infty}$ and $\{v^l\}_{l=0}^{\infty}$. The payoff u^0 corresponds to u in the right panel of Figure 5, and the subsequent sequence is the sequence of extreme payoffs that move counter-clockwise around the frontier. The payoff v^0 corresponds to v in the right panel of Figure 5, and the sequence of extreme points moves clockwise around the frontier. Aside from (4, 4, 4), the extreme equilibrium payoffs are permutations of points in these sequences.

Every v^l is generated the same way, by randomizing between u^l and (4, 4, 4), to make the incentive constraint for player 1 bind, i.e.,

$$v^{l} = \left(6 - \frac{1}{u_{2}^{l} - 4}, 6 + \frac{u_{1}^{l} - 4}{u_{2}^{l} - 4}, 3\right) = \frac{1}{2}(8, 8, 0) + \frac{1}{2}\left[\beta^{l}(3, u_{1}^{l}, u_{2}^{l}) + (1 - \beta^{l})(4, 4, 4)\right], \quad (1)$$

where

$$\beta^{l} = \frac{2}{u_{2}^{l} - 4}.$$
(2)

The payoffs u^l are generated in three different ways. First, the permutations of u^0 , i.e., the extreme points on the efficient frontier comprise a self-generating set and are generated according to

$$u^{0} = \left(\frac{11}{2}, \frac{15}{2}, 3\right) = \frac{1}{2}(8, 8, 0) + \frac{1}{2}\left[\frac{1}{4}\left(3, \frac{11}{2}, \frac{15}{2}\right) + \frac{3}{4}\left(3, \frac{15}{2}, \frac{11}{2}\right)\right],$$

i.e., by playing (B, B, C) for one period, followed by randomizing over two other efficient extreme payoffs to make the incentive constraint (player 1's in this case) bind.

Given u^0 , we can generate v^0 according to (1) and (2), which turns out to be $v^0 = (40/7, 45/7, 3)$, with $\beta^0 = 4/7$. The payoff u^1 is then generated by playing (B, B, C) for one period, followed by randomization between two permutations of v^0 :

$$u^{1} = \left(\frac{11}{2}, \frac{99}{14}, 3\right) = \frac{1}{2}(8, 8, 0) + \frac{1}{2}\left[\alpha^{1}(3, v_{2}^{0}, v_{1}^{0}) + (1 - \alpha^{1})(3, v_{1}^{0}, v_{2}^{0})\right].$$

where $\alpha^1 = 3/5$ is again chosen to make player 1's incentive constraint bind. Finally, the rest of the u^l sequence for $l \ge 2$ is generated according to

$$u^{l} = \left(6 - \frac{1}{v_{2}^{l-2} - 4}, 6 + \frac{v_{1}^{l-2} - 4}{v_{2}^{l-2} - 4}, 3\right) = \frac{1}{2}(8, 8, 0) + \frac{1}{2}\left[\alpha^{l}(3, v_{1}^{l-2}, v_{2}^{l-2}) + (1 - \alpha^{l})(4, 4, 4)\right],$$

where

$$\alpha^{l} = \frac{2}{v_{2}^{l-2} - 4}$$

is again chosen to make player 1's constraint bind.

Finally, these sequences converge to the accumulation points in Figure 5, which are permutations of $((9+\sqrt{5})/2, (11+\sqrt{5})/2, 3)$. These payoffs, together with (4, 4, 4), comprise another self-generating set, where

$$\left(\frac{9+\sqrt{5}}{2},\frac{11+\sqrt{5}}{2},3\right) = \frac{1}{2}(8,8,0) + \frac{1}{2}\left[\alpha^*\left(3,\frac{9+\sqrt{5}}{2},\frac{11+\sqrt{5}}{2}\right) + (1-\alpha^*)(4,4,4)\right],$$

where $\alpha^* = 3 - \sqrt{5}$.



Figure 6: Three different bounds on the equilibrium payoff set. Left: The convex hull of the flow payoffs. Center: The left set less the payoffs that are below the threat point. Right: Additional payoffs removed to create the set \widetilde{W}_0 . Faces that coincide with incentive constraints are colored tan.

A.2 Feasible set

We next argue that the equilibrium payoff set is the convex hull of the points constructed heretofore. The analysis consists of several steps. First, since only these three action profiles can possibly be played in equilibrium, we know that the equilibrium payoff set must be contained in the triangular pyramid with peak at (4, 4, 4) and base corners which are permutations of (0, 8, 8). Thus, the pyramid "points" in the direction (-1, -1, -1). In the sequel, we refer to this as the "feasible set."

A.3 Equilibrium threats

Clearly, the equilibrium threat point \underline{v} must be less than 4 (since the Nash equilibrium is certainly an equilibrium payoff). Thus, from the definition of \tilde{B} , the only way that player 3 can obtain a lower payoff is if (B, B, C) is played in the first period, with a flow payoff of (8, 8, 0). Moreover, any payoff we generate with this action must be weakly above (8, 8, 0) in the direction (0, 0, -1), and therefore it must be generated with a binding incentive constraint. But players 1 and 2 are playing myopic best responses at (B, B, C), so the only relevant incentive constraint is player 3's. Plugging in the specified payoffs and discount factor, we conclude that

$$\underline{v} = \frac{1}{2}3 + \frac{1}{2}\underline{v}$$
$$\iff \underline{v} = 3.$$

A.4 The efficient frontier

In addition, we claim that no equilibrium payoff can lie above the plane that contains (4, 4, 4), (11/2, 15/2, 3), and (3, 15/2, 11/2), i.e., with level x and direction λ such that

$$x = \lambda \cdot (4, 4, 4) = \lambda \cdot \left(\frac{11}{2}, \frac{15}{2}, 3\right) = \lambda \cdot \left(3, \frac{15}{2}, \frac{11}{2}\right)$$

$$\iff x = -52, \quad \lambda = (-7, 1, -7).$$
(3)

(The permutations of this statement also apply when we give the low payoff of 3 to player 1 or player 2). The reason is as follows. Consider maximizing payoffs in this direction. The optimal level must be at least -52, which is that of (4, 4, 4), the Nash equilibrium. But the flow payoff (8, 0, 8) has level -112, which is strictly below the Nash level, and hence cannot generate the optimal payoff. So, we may ask, what is the highest level that can be generated by (0, 8, 8) or (8, 8, 0)? We will consider the former, and the case for the latter is symmetric. In this direction, the flow payoffs (0, 8, 8) are maximal among all payoffs in the feasible pyramid, so that the minimal regime must be APS. To satisfy incentive compatibility, the continuation value of player 1 must be at least 6. Player 3's continuation value must be at least 3 from incentive compatibility. Finally, the sum of the payoffs is at most 16 (from feasibility). It follows that the highest level that can be attained in this direction is

$$\lambda \cdot \left(\frac{1}{2}(0,8,8) + \frac{1}{2}(6,7,3)\right) = -52.$$

Moreover, the permutations of $u^0 = (11/2, 15/2, 3)$ are a self-generating set. In particular,

$$\left(\frac{11}{2}, \frac{15}{2}, 3\right) = \frac{1}{2}(8, 8, 0) + \frac{1}{2}\left[\frac{1}{4}\left(3, \frac{11}{2}, \frac{15}{2}, \right) + \frac{3}{4}\left(3, \frac{15}{2}, \frac{11}{2}\right)\right].$$

We conclude that these are all extreme equilibrium payoffs (being at the corners of the hyperplanes in (3), the minimum payoff constraints, and the efficient frontier. Moreover, the convex hull of these points is the set of Pareto efficient payoffs.

The equilibrium payoff set must lie inside the polyhedron defined by the hyperplanes in (3), the constraints $v_i \ge 3$ for all *i*, and the constraint $\sum_i v_i \le 16$. We denote this set by \widehat{W} .

A.5 Structure of minimal regimes

Note that since (A, A, A) is a Nash equilibrium, no matter what feasible set W we consider, as long as $V \subseteq W$, the recursive regime will be minimal for (A, A, A), i.e., $x(\lambda, (A, A, A), W) = \lambda \cdot (4, 4, 4)$. In addition, we claim that whenever (B, B, C) is maximal, the minimal regime must be APS. For we already know that the payoffs (4, 4, 4), and permutations of u^0 can be generated. This pins down the optimal level exactly in all directions except those which are in the interior of $\widehat{\Lambda}_1 = \operatorname{co}\{(-7, -7, 1), (-7, 1, -7), (-1, 0, 0)\}$, or permutations thereof. (Outside of these sets of directions, an optimal payoff must be one of the aforementioned extreme points). $\widehat{\Lambda}_i$ denotes the permutations of these directions, where we give the weight -1 to a different player. For directions in $\widehat{\Lambda}_1$, it is easy to argue that (0, 8, 8) is higher than all other payoffs in the feasible triangle, so that necessarily the minimal regime for (C, B, B)is *APS*. In addition, either (8, 0, 8) or (8, 8, 0) is minimal among all feasible payoffs, so that the corresponding minimal regimes are all recursive, and hence these action profiles cannot be maximal in directions in $\widehat{\Lambda}_1$.

This means that for directions in $\widehat{\Lambda}_1$, the optimal level is simply given by $\widehat{x}^{APS}((C, B, B), \lambda)$, and we can reduce the computation of \widetilde{B} to simply computing the sets C(a) (where we drop the argument **W** for notational simplicity) for each $a \neq (A, A, A)$. Specifically, for all Wcontained within \widehat{W} ,

$$B(W) = \operatorname{co}\left(\{(4,4,4)\} \cup_{a \in \{(C,B,B),(B,C,B),(B,B,C)\}} C(a)\right).$$

We also note for future reference that if $v \in C(a)$ and $v' \in C(a')$, then there is no direction in which both v and v' are both maximal. This comes from the fact that the sets of directions $\hat{\lambda}_i$ are disjoint.

A.6 Two more bounds on binding payoffs

The focus of the analysis now shifts to the sets C(a) where $a_i = C$ and $a_{-i} = (B, B)$ for some $i \in \{1, 2, 3\}$. Ultimately we will construct a sequence of iterates using the \tilde{B} operator that converge to V and demonstrate that the limit set has infinitely many extreme points. Before doing so, we will slightly refine the approximation so that the sequence converges in an orderly manner.

It is straightforward that any $v \in C(B, B, C)$ must satisfy $v_j \ge 11/2$ for j = 1, 2. This follows from the fact that the flow payoff is 8 and the minimal equilibrium payoff is 3.

In addition, consider the direction (-7, -7, -29). We claim that no payoff in C(a) can be above the level -172. This level is attained by the Nash payoff (4, 4, 4) and also by (8, 8, 0) with maximal continuation payoffs w such that $w_3 \ge 6$ and $w \cdot (1, -7, -7) \le -52$. In particular, the solution is attained by payoffs

$$v^{0} = \left(\frac{45}{7}, \frac{40}{7}, 3\right) = \frac{1}{2}(8, 8, 0) + \frac{1}{2}\left[\frac{4}{7}\left(\frac{11}{2}, 3, \frac{15}{2}\right) + \frac{3}{7}(4, 4, 4)\right].$$

Note that since the permutations of $u^0 = (15/2, 11/2, 3)$ are already known to be part of a self-generating set, we know that v^0 are also equilibrium payoffs, and hence the plane with level -172 in direction (-7, -7, -29) is a supporting hyperplane of V. In fact, it intersects V in a face that contains (45/7, 40/7, 3), (40/7, 45/7, 3), and (4, 4, 4).

We thus conclude that C(B, B, C) is contained within the trapezoid of payoffs v defined by $v_3 = 3$, $\sum_i v_i \leq 16$, $v_1 \geq 11/2$, $v_2 \geq 11/2$, and $-29v_3 - 7(v_1 + v_2) \leq -172$. This trapezoid is denoted by \widetilde{C}_3^0 (\widetilde{C}_i^k will later denote a sequence of minimal payoff sets for player i). We note for future reference that \widetilde{C}_3^0 is the convex hull of the payoffs (15/2, 11/2, 3) and

$$w^0 = (11/2, 93/14, 3)$$

and the permutations obtained by interchanging the payoffs of players 1 and 2. Note that the payoff w^0 is at the intersection of the bounds $v_1 \ge 11/2$, $v_3 = 3$, and $-29v_3 - 7(v_1 + v_2) \le$ -172. We correspondingly define the sets \widetilde{C}_1^0 and \widetilde{C}_2^0 by permuting players' payoffs. We let

$$\widetilde{W}^{0} = \operatorname{co}\left(\{(4,4,4)\} \cup_{i=1,2,3} \widetilde{C}_{i}^{0}\right),\$$

which will serve as the initial set for the sequence we generate in the next and final subsection.

A.7 The sequence $\{\widetilde{W}^k\}$

We now analyze the sequence of sets produced by iterative application of \widetilde{B} to \widetilde{W}^0 . The critical issue is to determine the shape of the sets \widetilde{C}_i^{k+1} . In the following discussion, we take the perspective of minimum payoffs for player i = 3, but the case is symmetric for the other players.

We will argue that at iteration $k \ge 0$, the set \widetilde{C}_3^k is the convex hull of the points $\{u^l\}_{l=0}^k$, $\{v^l\}_{l=0}^{k-1}$, the payoff w^k defined as above for k = 0 and by

$$w^{k} = \frac{1}{2} \left(8, 8, 0\right) + \frac{1}{2} \left[\frac{2}{w_{2}^{k-1} - 4} \left(3, w_{1}^{k-1}, w_{2}^{k-1}\right) + \left(1 - \frac{2}{w_{2}^{k-1}}\right) \left(4, 4, 4\right)\right]$$

for $k \ge 1$, and the permutations thereof obtained by interchanging the payoffs of players 1 and 2. The base case has already been given for k = 0 in the previous subsection.

Let us take as an inductive hypothesis that the set \widetilde{W}^k is comprised of the following



Figure 7: The sets \tilde{C}_i^k for $k \in \{0, 1, 2, 3\}$. At every iteration, four new faces are added to each \tilde{C}_i^k . The right two panels show the left-hand corner of the set at different levels of magnification.

edges: First, there are edges between the payoffs in \widetilde{C}_i^k . Second, there are edges between permutations of u^0 that are in different \widetilde{C}_i^k sets. Finally, there are edges that connect all of the payoffs in \widetilde{C}_i^0 with the Nash payoff (4, 4, 4).

Given this inductive hypothesis for k - 1, we can easily compute the set \widetilde{C}_3^k . First, we compute the intersection of \widetilde{W}^{k-1} with the plane $w_3 = 6$ to find the extreme binding continuation values for player 3. We then average these payoff with the flow payoff (8, 8, 0) to obtain \widehat{C}_3^k . The intersections with the $w_3 = 6$ plane must lie on edges of \widetilde{W}^{k-1} that have one point with v_3 higher than 6 and another point with v_3 less than 6. There are three kinds of such edges that have intersections with the $w_3 = 6$ plane: The edges between permutations of u^0 , e.g., (3, 11/2, 15/2) and (3, 15/2, 11/2), which will generate the point u^0 ; The edges between permutations of w^0 (when k = 1) or between permutations of v^0 , which generate u^1 (when k > 1); and the edges between one of the payoffs whose permutation is in $\{u^l\}_{l=0}^{k-1} \cup \{v^l\}_{l=0}^{k-2} \cup \{w^{k-1}\}$, and the Nash payoff (4, 4, 4), which generate a payoff v^l , u^{l+2} , or w^k , respectively. From the inductive hypothesis, all of these intersections must result in new extreme payoffs of \widetilde{C}_3^k .

As an example, when k = 1, the payoffs generated will be u^0 , u^1 , v^0 and w^1 , as well as their permutations when we swap the payoffs of players 1 and 2. The first five elements of the \widetilde{C}_3^k sequence are depicted in Figure 7.

Finally, it remains to argue that the inductive hypothesis will be true for k. The new payoffs generated in \tilde{C}_3^k can be divided into those where player 1's payoff is at least 6, and those where player 1's payoff is less than 6. Focus for now on the former. These payoffs are maximal for directions that are convex combinations of (-7, -7, -29), (1, -7, -7), and (0, 0, -1). Note that we have already characterized supporting hyperplanes of V in these



Figure 8: Two views of \widetilde{W}^5 . Flow payoffs are marked with red crosses. Efficient extreme payoffs are red dots, inefficient extreme payoffs are blue dots. The minimum incentive compatible continuation value for player 3 (whose payoff is on the z axis) is a blue plane. The intersection of this set with the payoff set, contracted towards the payoff (8, 8, 0), generates the bottom flat of V.

three directions, which are also necessarily supporting hyperplanes of \widetilde{W}^k . For directions other than (0, 0, -1) and (1, -7, -7), the only other optimal payoff is (4, 4, 4), so that edges on supporting hyperplanes in these directions will be composed of either two payoffs in \widehat{C}_3^k with $v_1 \ge 3$, or one of the payoffs in \widetilde{C}_3^k with $v_1 \ge 6$ and (4, 4, 4). For the direction (-7, -7, -29), when k = 0, the permutations of w^0 and the Nash equilibrium are all optimal, so there is one additional edge, between the permutations of w^0 . For $k \ge 1$, it is the permutations of v^0 and (4, 4, 4) that are optimal in this direction. Finally, for the direction (0, 0, -1), all of the payoffs in \widehat{C}_3^k are optimal, so edges here will be between points in \widehat{C}_3^k . A similar analysis applies to other extreme points, so that the inductive hypothesis is true for k.

Note that at the kth round, we drop the permutations of w^{k-1} , but add the permutations of w^k , v^{k-1} , and u^k . Thus, the number of extreme points increases by 12 on every iteration. Moreover, the points u^k and v^{k-1} , once added, are never dropped, so the set of extreme points increases without bound. In the limit, the sequence w^k converges to the accumulation point w^* , which is generated according to

$$w^* = \frac{1}{2}(8,8,0) + \frac{1}{2} \left[\alpha^* \left(3, w_1^*, w_2^*\right) + (1-\alpha^*)(4,4,4) \right].$$

The weight α^* must solve

$$0 = -\frac{1}{8}\alpha^3 + \frac{1}{2}\alpha^2 + \alpha - 1.$$

This equation has three real roots, only one of which is between 0 and 1, which is $\alpha^* = 3 - \sqrt{5}$. The resulting payoff is

$$w^* = \left(\frac{9+\sqrt{5}}{2}, \frac{11+\sqrt{5}}{2}, 3\right).$$

The set \widetilde{W}^5 is depicted in Figure 8. At this resolution, this set is indistinguishable from V.

As a final note, while the analysis of this game is involved, in many ways it is the simplest example possible. Four is the minimum number of equilibrium action profiles such that the equilibrium payoff set is full dimension, which is necessary for the number of extreme points to be unbounded. The incentive constraints are also quite simple: One action profile is a Nash equilibrium, and for each other equilibrium action profile, only a single incentive constraint binds, that of the player whose payoff is being minimized.

B Additional examples

B.1 Two-player two-state Prisoners' Dilemma

This example illustrates the utility of the test directions in iteratively computing optimal levels. There are two states, L and R, and the stage game in each state is a Prisoners' Dilemma with the payoffs in Figure 9. The probability of staying in the same state is 1/3 if the players take the same action, and it is 1/2 if the players take different actions. The discount factor is $\delta = 2/3$.

We computed the sequence $\widetilde{\mathbf{W}}^k$ until the Hausdorff distance between successive iterations was less than 10^{-8} . The computation took 0.37 seconds. The sequence of payoff correspondences is depicted in Figure 10. The final payoff set for the left state has six extreme points, and the right state has four.

It turns out that the equilibrium threat point is generated by a policy that plays (D, D)in both states in the recursive regime. The resulting threat point is

$$(\underline{v}_i(L), \underline{v}_i(R)) = \left(\frac{8}{11}, \frac{14}{11}\right)$$

The utilitarian efficient payoffs that are optimal in the direction (1, 1) are generated by playing (C, C) in both states in the recursive regime. The resulting symmetric payoffs are 19/11 in the left state and 25/11 in the right state.

We may ask, how will the optimal policy change as the direction rotates clockwise from (1,1)? A natural conjecture, which turns out to be correct, is that the optimal policy will change by switching from (C, C) to (D, C) in some state. But should this switch occur first in the left state or the right state? Both switches would move flow payoffs in the same direction of (1, -3). But switching from (C, C) to (D, C) in state L would increase the probability of staying at s = L where payoffs are lower, whereas switching in s = R leads to a higher probability s = R, where payoffs are higher. Thus, less surplus is burnt by switching when s = R, and indeed this is the correct substitution.

$$\begin{array}{c|ccccc} s = L & s = R \\ a_1/a_2 & D & C & D & C \\ C & (-1,2) & (1,1) & (1,4) & (3,3) \\ D & (0,0) & (2,-1) & (2,2) & (4,1) \end{array}$$

Figure 9: A two-state Prisoners' Dilemma. The probability of remaining in the same state is 1/3 if $a_i = a_j$, and otherwise it is 1/2.

Our algorithm resolves this question mechanically using the test directions, which are



Figure 10: The sequence of correspondences generated by the max-min-max operator.



Figure 11: Test directions for (D, C), relative to the symmetric efficient payoffs.

depicted in Figure 11. The flow payoffs from (D, C) are depicted with black stars, payoff sets in blue, expected continuation payoff sets in red, binding incentive constraints in magenta. The test directions are black arrows. The shallowest legitimate test directions point along the frontier, and are generated by (D, C) in the right state. Note that there is a tie between the recursive and APS substitutions: both move payoffs along the frontier,¹ although only the recursive substitution is "lexicographically legitimate" in the sense described in Section 4.2.4.

¹In fact, there is a three-way tie, since there is a binding substitution for (C, C) in the right state that moves payoffs in the same direction. This test direction is not depicted in Figure 11.

B.2 A three-player contribution game

We implemented the stochastic algorithm for three players as part of the aforementioned SGSolve package.² Let us illustrate the algorithm with two examples. The first example is a simple contribution game: N = 3, $S = \{1, 2\}$, $\mathbf{A}_i(s) = \{0, 1\}$, and

$$u_i(a,s) = 2\sum_{j=1}^N a_j - 3a_i + 20s.$$

The transition probabilities are $\pi(s|a) = 1/2$ for every s and a, and $\delta = 2/3$. The stage game in each state is effectively a three-player Prisoners' Dilemma.

This example illustrates how our algorithm can solve for the equilibrium payoff exactly. We initialized the algorithm with 214 directions that are distributed around the unit sphere. We used the convergence criterion that no directions were added or dropped between iterations, and the Hausdorff distance between consecutive iterations was less than 10^{-8} . Due to the stochastic nature of the algorithm, its performance varies on each run. On one series of five runs, the algorithm finished with 9 directions three times, and 10 directions the other two. Over the course of one of the runs that terminated with 9 face directions, the algorithm added 72 endogenous directions and dropped 277. In all cases, the algorithm converged in 45 iterations and took between 2.85 and 3.11 seconds.

One can analytically verify that the equilibrium payoff correspondence for this game has exactly 9 face directions. Thus, in the runs where the algorithm terminated with 9 directions, it correctly identified the structure of equilibrium payoffs, which are depicted in Figure 12. All sixteen action profiles can be sustained. The efficient points are generated by always playing a = (1, 1, 1) in both states. There is also an inefficient point which corresponds to the Nash equilibrium a = (0, 0, 0). The remaining points are generated by playing permutations of a = (1, 1, 0) and (1, 0, 0). The face in which a player's payoff is minimized is attained by $a_i = 1$ and $a_{-i} = (0, 0)$.

For comparison, we solved this game using our implementation of the JYC algorithm with the same set of 214 initial directions. The same tolerance was achieved in 49 iterations and 3 minutes and 38.45 seconds. So, the JYC code is between one and two orders of magnitude slower. All of our previous caveats still apply, but we find this suggestive that the stochastic max-min-max algorithm is significantly more efficient.

A natural question is, which features of the max-min-max operator explain the difference in performance? We also ran a version of our algorithm with the same 214 initial directions,

 $^{^{2}}$ We note that the graphical interface currently only works for two-player games, but the three-player routines are part of the callable library.



Figure 12: Equilibrium payoffs for the contribution game. The equilibrium payoff correspondence is in red, and the expected equilibrium payoffs are in gray.

but where we set $\widehat{\Lambda}^k = \widehat{\Lambda}^0$ for all k, i.e., the set of directions is held fixed. In this case, the algorithm converged in 44 iterations and 3.53 seconds. This suggests that most of the efficiency gain comes from using the max-min-max level rather than APS. The endogeneity of directions, however, leads to a tight limit set.

B.3 Three-player risk sharing and partial formal insurance

We solved a three-player risk-sharing game, as in Section 4. Each player now has an endowment e_i and their actions are vectors that specify how much they transfer to each other player. For this particular simulation, we used $u(c) = \sqrt{c}$, the endowment grid is $E = \{0, 0.5, 1\}$, and endowment distribution is independent across periods and uniform over endowment profiles that sum to 1. The discount factor is $\delta = 0.6$. For this simulation, we capped the algorithm at 300 directions and iterated until a convergence threshold of 10^{-8} . The algorithm converged in 68 seconds and 33 iterations. Over the course of the computation, 492 endogenous directions were added and 551 redundant directions were dropped. The computed expected equilibrium payoff set is depicted in the left-hand panel of Figure 13.

As a simple application, we used our algorithms to investigate the following question: What happens to equilibrium risk sharing and social welfare if the players can write formal insurance contracts? If all of the players can write a formal full insurance contract, so that they equally share their collective resources, then the welfare implications are obvious: The sum of the agents' surpluses must weakly increase. If only two of the three players can write such a contract, however, the implications are ambiguous. Suppose that players 2 and 3 write such a contract, so that $c_2 = c_3$ and each consumes half of their total endowment net of transfers to player 1, and their transfers to player 1 are chosen to maximize the joint



Figure 13: Risk sharing with three players. Left: Expected equilibrium payoffs. Right: Achievable $(v_1, v_2 + v_3)$ pairs. Non-cooperative play is in blue, and cooperation between players 2 and 3 is in red.

surplus $v_2 + v_3$. On the one hand, players 2 and 3 should be better off, because they are always guaranteed a minimal level of insurance, so that their autarkic payoffs are higher with such a contract than without. On the other hand, the higher autarkic payoffs tighten incentive constraints and may reduce risk sharing with player 1.

We used the two-player algorithm of Section 4 to investigate what would happen if players 2 and 3 behave cooperatively to maximize their joint surplus. Expected equilibrium payoffs are plotted in red in the right-hand panel of Figure 13. For comparison, the blue curve represents the possible $(v_1, v_2 + v_3)$ pairs in the game where players 2 and 3 behave non-cooperatively. The threat payoff for players 2 and 3 is clearly higher with the contract: their minimum joint surplus is approximately 0.805 in the non-cooperative case, and approximately 1.04 when they cooperate. A striking result is that the tightening of incentive constraints appears to overwhelm the benefits of additional risk sharing, and the Pareto frontier when players 2 and 3 cooperate is strictly below the Pareto frontier when they behave non-cooperatively. Thus, the example illustrates how partial insurance contracts may lead to lower social welfare.

B.4 Lower bounds on payoffs

Recall the two-state risk sharing example of Section 4 with $\delta = 0.7$. Figure 14 compares the upper and lower bounds on equilibrium payoffs, which are blue and red respectively. The lower bound was computed with $\epsilon = 0.005$. The red dotted line corresponds to the expansion of the lower bound by ϵ in every direction, which is contained in the correspondence that would be produced by the APS operator. The payoffs that induce the upper and lower



Figure 14: Lower bounds on V for $\epsilon = 0.005$. Dots represent the actual payoffs used to generate the bounds.

bounds are represented as dots. Note that the distance from the payoffs to the lower bound set varies depending on the direction of the bound. This distance is greater when more states have a minimal regime that is recursive. When both states are binding, such as the payoffs that approximate the threat point, the penalty is ϵ in both states. When one state is binding, such as when we maximize one player's payoff, the penalty in the binding state is still ϵ , but the penalty in the recursive state is $\epsilon/(1 - \delta/2)$. When both states are recursive, which is when the direction is close to maximizing the sum of payoffs, the penalties in both states are $\epsilon/(1 - \delta)$. At directions where the minimal regimes change, the penalties (and hence the level of the optimal payoffs) change discontinuously.

The computation depicted in Figure 14 used a relatively large value for ϵ for visual effect. When ϵ is small, the distance between the outer and inner bounds shrinks as well and appears to go to zero. For example, we computed upper bounds on **V** and **V**^{ϵ} to a tolerance of 10⁻⁷ when $\epsilon = 10^{-6}$. The extreme points of the upper bound on **V** are all within 10⁻⁶ of the bounds for **V**^{ϵ}, so that the lower and upper bounds are indistinguishable (up to the tolerance for computing the extreme points of the upper bound).