Operations Research

Computing Equilibria of Dynamic Games

Şevin Yeltekin, Yongyang Cai, Kenneth L. Judd

To cite this article:
Şevin Yeltekin, Yongyang Cai, Kenneth L. Judd (2017) Computing Equilibria of Dynamic Games. Operations Research 65(2):337-356. https://doi.org/10.1287/opre.2016.1572

OPERATIONS RESEARCH
Vol. 65, No. 2, March–April 2017, pp. 337–356
ISSN 0030-364X (print), ISSN 1526-5463 (online)
http://pubsonline.informs.org/journal/opre/

Computing Equilibria of Dynamic Games

Şevin Yeltekin,[a] Yongyang Cai,[b,c,d] Kenneth L. Judd[c,e]

[a] Tepper School of Business, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213; [b] Becker Friedman Institute, University of Chicago, Chicago, Illinois 60637; [c] Hoover Institution, Stanford University, Stanford, California 94305; [d] Department of Agricultural, Environment and Development Economics, The Ohio State University, Columbus, Ohio 43210; [e] National Bureau of Economic Research, Cambridge, Massachusetts, 02138

Contact: sevin@andrew.cmu.edu (ŞY); cai.619@osu.edu (YC); kennethjudd@mac.co (KLJ)

Received: March 24, 2015
Revised: June 4, 2016
Accepted: September 23, 2016
Published Online in Articles in Advance: February 2, 2017

Subject Classifications: games/group decisions: cooperative; analysis of algorithms; facilities/equipment planning: capacity expansion
Area of Review: Computational Economics

https://doi.org/10.1287/opre.2016.1572

Abstract. We develop a numerical method for computing *all* pure strategy subgame-perfect equilibrium values of dynamic strategic games with discrete states and actions. We define a monotone mapping that eliminates dominated strategies, and when applied iteratively, delivers an accurate approximation to the true equilibrium payoffs of the underlying game. Our algorithm has three parts. The first provides an outer approximation to equilibrium values, constructed so that any value outside of this approximation is *not* an equilibrium value. The second provides an inner approximation; any value contained within this approximation *is* an equilibrium value. Together, the two approximations deliver a practical check of approximation accuracy. The third part of our algorithm delivers sample equilibrium paths. To illustrate our method, we apply it to a dynamic oligopoly competition with endogenous production capacity.

Funding: This research is a part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation (Awards OCI-0725070 and ACI-1238993) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications. Yongyang Cai gratefully acknowledges support from the National Science Foundation [Grants SES-0951576 and SES-1463644].

Keywords: dynamic games • multiple equilibria • dynamic oligopoly • computation

1. Introduction

Analysis of strategic decision making is one of the fundamental areas of focus in economics. Static, repeated, and dynamic games are used widely in industrial organization, international trade, behavioral economics, macroeconomic policy making, and political economy. One of the issues that complicate their analysis is multiplicity of equilibria. This problem arises frequently in repeated games, and is more severe in dynamic games where the environment in which the strategic interaction takes place changes over time. To reduce the complexity of the analysis, economists have traditionally restricted strategies or chosen equilibria with certain features such as symmetry, stationarity, or Pareto optimality. In this paper, we take an alternative approach. Rather than focusing on reducing the multiplicity problem, we introduce a numerical method for computing *all* subgame-perfect equilibria of dynamic games with discrete states. We then apply this method to an oligopoly competition with endogenous production capacity and illustrate how it can be used in the analysis of such a dynamic game.

Specifically, we consider games with an infinite horizon, discrete states and actions, and with perfect monitoring where all players observe the entire history of the game, including actions of all other players. Additionally, players' actions affect the state of the world the game is played in. As a first step, we formulate such games in payoff space, rather than in action space, and show that the subgame-perfect equilibrium (SPE) payoffs are supported by player actions consistent with the Nash equilibrium in the current period and continuation payoffs, that are themselves payoffs in some SPE in the next period. These continuation equilibrium payoffs are drawn from equilibrium values consistent with the evolution of the discrete state.

Our methodology builds on the influential papers of Abreu et al. (1986, 1990) (APS) who developed techniques for characterizing repeated games of incomplete information and Cronshaw and Luenberger (1990, 1994) (CL) for repeated games of complete information. Computational methods for approximating the set of equilibria of repeated games with complete information were developed by Conklin and Judd (1996) and Judd et al. (2003) (JYC). In this paper, we extend the methods of APS, CL, and JYC to dynamic games with discrete states and show that the collection of SPE payoffs of such games can be obtained by repeated application of a monotone mapping that eliminates dominated strategies.

The practical implementation of our iterative method requires an efficient approximation scheme for equilibrium payoffs. In the first part of this paper, we provide such a scheme for the class of dynamic games we study. We follow a long tradition in game theory, starting with Fudenberg and Maskin (1986), and convexify the dynamic game by introducing a public randomization device. We then show that polygons provide an efficient and consistent means of approximating the equilibrium payoff set(s) from such convexified games. Moreover, they enable us to obtain two approximations to equilibrium values. The first, called an outer approximation, is constructed so that any value outside of this approximation is *not* an equilibrium value. The second is an inner approximation; any value contained within this approximation *is* an equilibrium value. Since the true equilibrium values lies between these approximations, the difference between them provides us with an error bound for gauging the accuracy of the procedure at each finite iteration.[1] Additionally, we provide an algorithm to compute sample equilibrium paths that support equilibrium values.

Our approach has several desirable features. First, it is applicable to a large class of dynamic games that arise naturally in industrial organization. Amongst them are oligopoly settings with research and development, advertising, learning by doing, capacity expansion, market entry and exit, inventory management, and technology adoption.[2] Second, it allows us to do comparative statics in value sets, for example, compare equilibrium value sets across different discount factors, technology, laws of motion for states and different forms of asymmetry between players. Third, it allows us to verify whether payoffs of interest, such as payoffs associated with collusive behavior or outcomes of interest, such as emergence of a monopoly are supported in equilibrium. Fourth, by providing a range of possible outcomes, it can shed light on policy interventions that affect the strategic environment.

To illustrate our method, we apply it to a specific dynamic game in industrial organization: oligopolistic competition with endogenous productive capacity. We choose this particular application because a large literature in game theory and industrial organization investigates how firms use capacity strategically, for example, to deter competitors from entering into the market or to ease tacit collusion. Due to the complex nature of this dynamic game and the multiplicity of equilibria, it is common to employ simplifying assumptions to make the problem more tractable. These assumptions include choosing capacity once and for all in the first period and/or restricting attention to symmetric Markov equilibria, or focusing on the best (Pareto-optimal) equilibria.[3] Although these assumptions help simplify the analysis, they may also eliminate equilibria of interest, while not substantially reducing the multiplicity problem. We take an alternative approach and compute all SPE of this dynamic game without such restrictions. We demonstrate how equilibrium values change as discount factors, maximum attainable capacity, reversibility of investment, and cost of production are modified. Since our method delivers inner and outer approximations that together "sandwich" the true equilibrium values, we show how it can be used to rule in or rule out equilibrium outcomes featuring tacit collusion, monopoly power, overinvestment or overproduction.

Section 2 describes the infinite-horizon game with finite states and actions. Section 3 introduces our approximation method and provides details of our algorithms. Section 4 applies our method to a dynamic capacity game with endogenous capacity investment and provides a selection of results from different parameterizations of this game. Section 5 concludes.

2. Supergames with State Variables

We start our analysis by describing the dynamic game with finite actions and finite states.

2.1. Histories and Strategies

N infinitely lived agents play a dynamic game. Let the finite set X_i denote agent i's set of states, and let $X = \bigtimes_{i=1}^{N} X_i$ denote the set of aggregate states of the game. The game unfolds with simultaneous moves at each stage where each player i chooses a perfectly observable action a_i from a finite set A_i. Elements of the set $A = \bigtimes_{i=1}^{N} A_i$ represent all possible combinations of player actions and are called *action profiles*. An action profile will be denoted by a. Additionally, we use the standard notation a_{-i} to refer to an action profile that excludes player i. The state evolves deterministically according to $g\colon A\times X \to X$, $x_{t+1} = g(a_t, x_t)$.[4] Let $\Pi_i\colon A\times X \to \Re$ be the current period payoff of player i.

We assume that for each player, the minimal and maximal period payoffs in each state are bounded by the scalars:

$$\overline{\Pi}_i = \max_{(a,x)\in A\times X} \Pi_i(a,x),$$
$$\underline{\Pi}_i = \min_{(a,x)\in A\times X} \Pi_i(a,x).$$

Note that these are well defined since the action space is assumed to be finite.

The action space for the dynamic game is A^∞. Agent i's average discounted payoffs from a specific sequence of states and action profiles are

$$U_i(a^\infty, x^\infty) = (1-\delta)\sum_{t=0}^{\infty} \delta^t \Pi_i(a_t, x_t),$$

where $\delta \in (0,1)$ is the common discount factor across agents.[5] A t-period history, h^t, is a pair of sequences $(\{a_s\}_{s=0}^{t-1}, \{x_s\}_{s=0}^{t})$. Therefore agents know the entire history of actions taken by each player prior to the current period and the aggregate states that realized prior to and including the current period. Let H^t denote the set of t-period histories. A *pure strategy* for player i is a sequence of functions $\{\sigma_{i,t}\}_{t=0}^{\infty}$ that map histories to actions with $\sigma_{i,t}\colon H^t \to A_i$. A strategy profile is a sequence of functions $\{\sigma_t\}_{t=0}^{\infty}$, where σ_t maps from H^t to A. We use $\sigma \mid h^t$ to denote the continuation strategy profile for the remainder of the game that follows the history h^t.

In each state x, the set of equilibrium payoffs is contained within the compact set $\Omega = \times_{i=1}^{N}[\underline{\Pi}_i, \overline{\Pi}_i]$. We define $\mathcal{P}^*$ as the set of all correspondences that map each point in $x \in X$ to some closed subset of Ω:

$$\mathcal{P}^* = \{W\colon X \rightrightarrows \Omega\}.$$

We now define a partial order on P^* that we will use throughout the paper. Consider two correspondences, $U, W \in \mathcal{P}^*$. Let U_x denote the state x component of correspondence U, and W_x the state x component of W. Using set notation, we say $U \subseteq W$ if $\forall\, x \in X$, $U_x \subseteq W_x$. Note that since X is finite, this partial order is well defined.[6]

2.2. Equilibrium and Its Characterization

The equilibrium concept we employ for our dynamic game is SPE. In a dynamic game, at any history, the "remaining game," called the subgame, can be regarded as a game of its own. In dynamic games, the Nash equilibrium is too permissive because it imposes no optimality conditions in these subgames, opening the door to violations of sequential rationality. Subgame perfection strengthens the Nash equilibrium by imposing the sequential rationality requirement that behavior be optimal in all circumstances (i.e., subgames), those that arise in equilibrium (as required by the Nash equilibrium), and those that arise out of equilibrium.

Definition 1. A strategy profile σ is a SPE if, for any history $h^t \in H^t$ ending in state x, the continuation strategy $\sigma \mid h^t$ is a Nash equilibrium of the continuation game.

Now, we can formally define the SPE payoff correspondence of our dynamic game.

Definition 2. Let V^* denote the correspondence that maps the current state into the set of average discounted payoffs that can be sustained in pure SPE.

In our formulation of this dynamic game, each SPE payoff vector $v \in V^*$ is supported by a profile of actions a consistent with Nash play in the current period and a vector of continuation payoffs w that are themselves payoffs in some SPE. The key to finding V^* involves defining an operator that maps future SPE payoffs into current SPE payoffs.

We use the one-stage deviation principle for infinite-horizon games, which provides a useful characterization of SPE. This principle applies to games where overall payoffs are a discounted sum of uniformly bounded stage payoffs, as is the case in our setting.

Theorem 1. *In a dynamic game with finite states and observed actions, profile a is subgame perfect if and only if there is no player i and strategy $\tilde{a}_i$ that agrees with a_i except at a single t and h^t, and such that $\tilde{a}_i$ is a better response to a_{-i} than a_i, conditional on history h^t being reached.*

Proof. See Fudenberg and Tirole (1991), Theorem 4.2.

An immediate consequence of the one-stage deviation principle is that V^* is equivalent to the correspondence that maps the current state into the set of average discounted payoffs that can be sustained by strategy profiles in which no player has a profitable one-stage deviation. This constraint, known as the incentive compatibility constraint, is central to our analysis.

Incentive compatibility. In any Nash equilibrium, player i must prefer the equilibrium action to any alternative, given the equilibrium actions of the other players. Suppose that W represents the set of possible continuation values of the dynamic game with W_x the state component of W. Also, suppose that in state x, player i is supposed to play a_i. Recall that a_{-i} is the action profile of all agents except for player i. The state in the next period will be $g(a_i, a_{-i}, x)$ and player i's continuation utility $w_{i,g(a,x)}$ has to be taken from the set of possible values in state $g(a_i, a_{-i}, x)$, $W_{g(a_i,a_{-i},x)}$. If instead player i chooses $\tilde{a}_i$ while the other players continue to play a_{-i}, he earns a current payoff of $\Pi_i(\tilde{a}_i, a_{-i}, x)$ and the next period's state will be $g(\tilde{a}_i, a_{-i}, x)$. We assume that if a player deviates this period, he will receive the smallest value given W, at next period's state $g(\tilde{a}_i, a_{-i}, x)$. We denote the worst payoff in the next period's state by $\mu_{i,g(\tilde{a}_i,a_{-i},x),W}$ with $\mu_{i,g(\tilde{a}_i,a_{-i},x),W} \in W_{g(\tilde{a}_i,a_{-i},x)}$. We express the gain to player i from playing a_i and receiving a continuation value of $w_{i,g(a,x)}$ in state x instead of playing $\tilde{a}_i$ and receiving continuation value $\mu_{i,g(\tilde{a}_i,a_{-i},x),W}$ as

$$\begin{aligned} IC_{x,W}&(i, a, \tilde{a}_i, w) \\ &\equiv (1-\delta)\Pi_i(a_i, a_{-i}, x) + \delta w_{i,g(a,x)} \\ &\quad - ((1-\delta)\Pi_i(\tilde{a}_i, a_{-i}, x) + \delta \mu_{i,g(\tilde{a}_i,a_{-i},x),W}). \end{aligned}$$

Given this definition, the incentive compatibility constraint for each agent i and each action profile a in state x given W can be written as

$$IC_{x,W}(i, a, \tilde{a}_i, w) \geqslant 0.$$

The incentive compatibility constraint, IC is therefore the temptation to deviate by playing $\tilde{a}_i$ and must be nonnegative for all players i and for all possible deviations $\tilde{a}_i$ for the action profile a to be incentive compatible.

B^* operator. We now turn to the formal description of our operator $B^*\colon \mathcal{P}^* \to \mathcal{P}^*$. Let $W \in \mathcal{P}^*$. We define $B^*(W)_x$ to be the set of possible payoffs consistent with a Nash equilibrium profile a in state x today and continuation payoffs drawn from the set $W_{g(a,x)}$.[7] That is,

$$B^*(W)_x = \bigcup_{(a, w_{g(a,x)})} \{(1-\delta)\Pi(a,x) + \delta w_{g(a,x)}\}$$

subject to

$$w_{g(a,x)} \in W_{g(a,x)}$$

and for each $1 \leqslant i \leqslant N$ and each $\tilde{a}_i \in A_i$, the incentive compatibility condition holds:

$$IC_{x,W}(i, a, \tilde{a}_i, w) \geqslant 0.$$

Given the previous definitions, a value v_x is in $B^*(W)_x$ if there is a continuation payoff profile $w_{g(a,x)} \in W_{g(a,x)}$ such that $v_x = (1-\delta)\Pi(a,x) + \delta w_{g(a,x)}$ is the value of playing a today, and for each i, player i will choose to play a_i because he believes that to do otherwise would yield him the worst continuation payoff in next period's state. Therefore the B^* operator maps a continuation value correspondence to a current value correspondence, analogous to the functional Bellman operator in dynamic programming, which maps a continuation value function to a current value function.

Self-generation. With the definition of B^* complete, we can now express the concept of self-generation central to our analysis.

A correspondence W is self-generating if $W \subseteq B^*(W)$. Using an extension of the arguments of Cronshaw and Luenberger (1994), in a related paper, Baldauf et al. (2014) show that any self-generating correspondence is contained within the equilibrium payoff correspondence, V^*, and that the equilibrium payoff correspondence is itself self-generating. The same paper also establishes that a unique maximal fixed point of B^* exists and that the equilibrium payoff correspondence is the maximal fixed point of the operator B^*.[8]

The operator B^* has important properties that can be exploited for computing the equilibrium correspondence. Specifically, B^* is monotone in the set inclusion ordering, so that if $U \subseteq W$, $U, W \in \mathscr{P}^*$, then $B^*(U) \subseteq B^*(W)$. Additionally, it preserves compactness. It is then possible to show that the equilibrium value correspondence V^*, the maximal fixed point of the mapping B^*, may be obtained by repeatedly applying this operator to a correspondence that is known to contain V^*. The details, as well as the proofs of these statements can be found in Baldauf et al. (2014) and Kitti (2016).

Applying the B^* operator numerically requires that the candidate value correspondences be efficiently represented on a computer and that the monotonicity property of the B^* operator is preserved. In the next section, we introduce approximation schemes specifically designed to have such features.

3. Approximating Equilibrium Value Correspondences

To approximate the equilibrium value correspondence of a dynamic game, we proceed in two steps. First, we convexify the underlying game and its equilibrium value correspondence via public randomization. Then, we develop methods for approximating convex-valued correspondences.

Before we proceed with the details of public randomization and approximation algorithms, we provide the following characterization of a convex hull.

Definition 3. If Y is a finite set of L points in $\Re^N$: $Y = \{y_i \mid i = 1, \ldots, L\}$, then the *convex hull of* Y is defined as

$$\text{co}(Y) = \left\{ b \in \Re^N \mid \exists\, \lambda \geqslant 0, \sum_{i=1}^{L} \lambda_i = 1, b = \sum_{i=1}^{L} \lambda_i y_i \right\}.$$

For the remainder of the paper, the notation $co(\cdot)$ will refer to the definition above.

3.1. Public Randomization

So far, we have not imposed any structure on the dynamic game to ensure that the equilibrium value correspondence has features amenable to approximation. As a first step, we follow the long tradition in the repeated game theory literature, starting with Fudenberg and Maskin (1986), and convexify the feasible payoffs in each state with a public randomization device at the start of each period.[9]

More precisely, we assume that in each stage of the game, there is a public lottery observable to all players, independent of all previous choices and realizations of the lottery. Histories now include previous action profiles as well as lottery outcomes. In each period, players make their simultaneous choices based on this augmented history. We also assume that the support of the lottery is contained in $W_{g(a,x)}$.

Now, W is defined as the possible ex ante continuation value correspondence at time t of the dynamic game with public randomization. Similarly, $co(B^*(W))$, the convex hull of $B^*(W)$, is defined as the set of ex ante continuation values available at $t-1$ for the dynamic game with public randomization.

It is important to note that the players still use pure strategies; player i does not mix between different actions $a_i \in A_i$. The public randomization delivers an "expected" continuation value profile $w_{g(a,x)}$ once the action profile a is chosen and next period's state $g(a,x)$ is determined. This makes that continuation value set convex. In the next period, the lottery outcome determines which equilibrium that is in the support of $w_{g(a,x)}$ will be played.

We let V denote the ex ante continuation value correspondence which can occur in equilibrium of the dynamic game with public randomization. Therefore V is convex valued and bounded. With minor modifications of the arguments presented in Baldauf et al. (2014), it can be shown that the repeated application of the operator with public randomization produces a sequence of convex-valued correspondences that converge to the equilibrium value correspondence V.

We now define $\mathcal{P}$ to be the set of all correspondences that map each point in $x \in X$ to some convex subset of Ω:

$$\mathcal{P} = \{W: X \rightrightarrows \Omega \mid \forall\, x \in X,\ W_x \text{ convex}\}.$$

We define $B: \mathcal{P} \to \mathcal{P}$ as

$$B(W) = \text{co}(B^*(W)), \quad W \in \mathcal{P}.$$

It then follows that B is monotone (i.e., for $U, W \in \mathcal{P}$, if $U \subseteq W$, then $B(U) \subseteq B(W)$), V is the largest fixed point of B and if $W^0 = \{W_x = \Omega, \forall\, x \in X\}$, and $W^{j+1} = B(W^j)$, then $V = \bigcap_j W^j$.

Sorin (1986) and Fudenberg and Maskin (1991) have shown that public randomization is dispensable if the discount the factor is sufficiently high. Our goal is to compute V for any discount factor. Public randomization is a useful device because it delivers equilibrium value sets that are convex for small discount factors.

Although public randomization convexifies equilibrium value sets, the challenge of representing an arbitrary convex set or correspondence on the computer remains. We represent a convex set computationally with a pair of approximations, called inner and outer approximations, the former a subset and the latter a superset of the approximated set. We proceed with the precise definitions of inner and outer approximations of convex sets, and then provide approximations for convex-valued correspondences.

3.2. Inner and Outer Approximations of Convex Sets

This section introduces the definitions of inner and outer approximations of a convex set.

Definition 4. If $Z \subset W$ and Z is a finite set of m points in $\Re^N$: $Z = \{z_i \mid i = 1, \ldots, m\}$, then the inner approximation W^I generated by Z is $W^I = \text{co}(Z)$.

Recall that with Definition 3, $w \in \text{co}(Z)$ will represent the linear equations:

$$\sum_{j=1}^{m} \lambda_j = 1, \quad \lambda \geqslant 0 \quad \text{and} \quad w = \sum_{j=1}^{m} \lambda_j z^j, \quad z_j \in Z.$$

Figure 1 illustrates a generic convex set W, and an inner approximation to it using the set of points $\{z^1, \ldots, z^8\}$. The piecewise linear boundary of the inner approximation allows us to represent W as a set of linear inequalities.[10] Similarly, our outer approximation of a convex set involves a set of linear inequalities, precisely defined below.

Figure 1. (Color online) Inner approximation

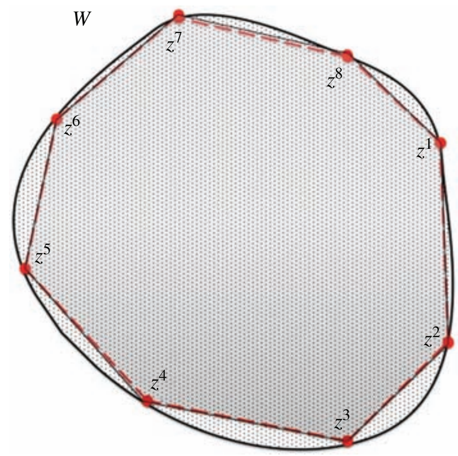

Definition 5. Assume $Z = \{z^1, \ldots, z^m\}$ is a set of m points on the boundary of W, and $R = \{r^1, \ldots, r^m\} \subset \Re^N$ is a set of m corresponding subgradients (normals) oriented such that $(z^l - w) \cdot r^l > 0$ for every $w \in W$. Then, the outer approximation W^O generated by Z and R is $W^O = \bigcap_{l=1}^{m} \{z \in \Re^N \mid r^l \cdot z \leqslant r^l \cdot z^l\}$.

Figure 2 illustrates the set of normals $\{r^l\}$ and the set of points $\{z^l\}$ used to construct an outer approximation to the two-dimensional set W. An outer approximation is constructed by the intersection of half spaces defined by the hyperplanes shown in Figure 2. Figure 3 plots the inner and the outer approximations generated by the set of points $\{z^l\}$ and the corresponding normals. The inner approximation lies within the outer approximation and the boundary of each set can be represented by a set of linear functions.

3.3. Inner and Outer Approximations of Convex-Valued Correspondences

We now introduce the general concept of inner and outer approximation of a convex-valued correspondence.

Figure 2. Outer approximation

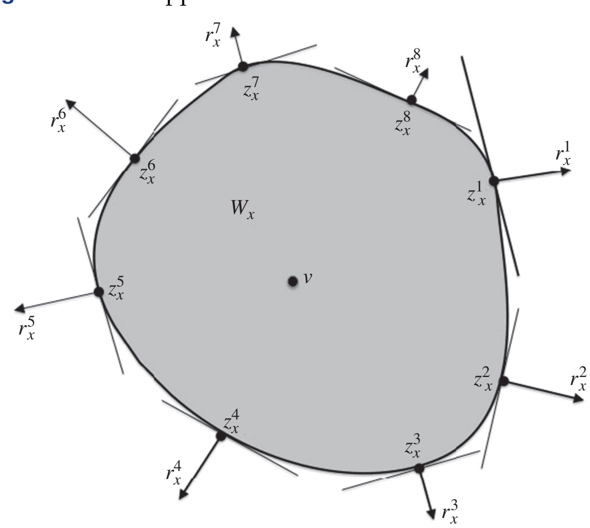

Figure 3. (Color online) Inner vs. outer approximations

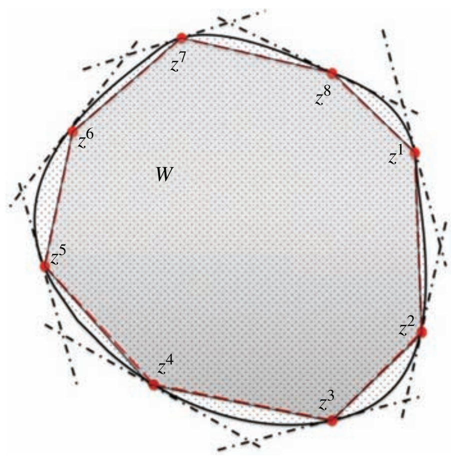

Definition 6. Suppose $W \in \mathcal{P}$ and Z is a correspondence such that the cardinality of Z_x is finite for each x. If for each x, $Z_x \subset W_x$, then the inner approximation W^I generated by Z is $W_x^I = \text{co}(Z_x)$ for each x, where $\text{co}(Z_x)$ denotes the convex hull of Z_x.

Definition 7. Suppose $W \in \mathcal{P}$ and Z is a correspondence such that the cardinality of Z_x is finite for each x. Also, suppose that for each x, $Z_x = \{z_x^1, \ldots, z_x^m\}$ is a set of m points on the boundary of W_x, and $R_x \subset \Re^N$ a set of m corresponding subgradients (normals) oriented such that $(z_x^l - w_x) \cdot r_x^l > 0$. Then, the outer approximation W^O generated by Z is $W_x^O = \bigcap_{l=1}^m \{z_x \in \Re^N \mid r_x^l \cdot z_x \leqslant r_x^l \cdot z_x^l\}$ for each x.

3.4. Outer Approximation of B

Having defined the inner and outer approximations for a correspondence W, we now turn to the approximation of the B operator. The critical property of B is that it maps convex-valued correspondences into convex-valued correspondences and that it is monotone. In particular, $B(W)$ maps $\mathcal{P}$ into itself. We first define an outer monotone approximation of the operator B that preserve these critical properties and detail the outer approximation algorithm. We then provide the same for the inner approximation.

It is important to note that the B operator can be discontinuous. A small reduction in the discount factor can cause the set of equilibria to collapse from substantial cooperation to one with no cooperation. The inherent sensitivity of equilibria to parameter values implies that there will also be sensitivity to numerical approximation errors. Our computational procedure is designed to handle this possibility.

Definition 8. A mapping $B^O: \mathcal{P} \to \mathcal{P}$ is an outer monotone approximation of B if
1. $\forall W \in \mathcal{P}$, $B^O(W) \supseteq B(W)$, and
2. $\forall U, W \in \mathcal{P}$ if $U \subseteq W$, then $B^O(U) \subseteq B^O(W)$.

The definition of an outer monotone approximation directly implies the following proposition, which relates the maximal fixed point of B^O to the maximal fixed point of B and also provides a sufficient condition for the outer monotone approximation scheme to converge.

Proposition 1. *Suppose $B^O(\cdot)$ is an outer monotone approximation of $B(\cdot)$. Then, the maximal fixed point of B^O contains V. More precisely, for $W \in \mathcal{P}$ if $W \supseteq B^O(W) \supseteq V$, then $B^O(W) \supseteq B^O(B^O(W)) \supseteq \cdots \supseteq V$.*

Proof. Follows from the self-generation of V and the definition of an outer monotone approximation.

Proposition 1 establishes that starting the the outer monotone approximation of the B operator with a correspondence that contains the true equilibrium value correspondence is sufficient for the approximation scheme to converge. The following Lemma asserts that the special correspondence $\mathcal{W} = \{W_x = \Omega, \forall x \in X\}$ is a good initial guess for the outer monotone approximation.

Lemma 1. $\mathcal{W} \supseteq B^O(\mathcal{W}) \supseteq V$.

Proof. Follows from the definition of $\mathcal{W}$ and self-generation.

Problems without state variables only require the approximation of a single value set at each iteration. In contrast, the dynamic problem we are interested in requires that a collection of approximated value sets, one for each element of the (finite-) state space, be found at each iteration. Additionally, the latter problem requires that the implications of each action profile (and each deviation from an action profile) for the future state be found, and that future continuation values be obtained appropriately from the relevant continuation value set. Next, we provide the details of our outer monotone approximation algorithm for dynamic games with discrete states and full information.

Outer Monotone Approximation Algorithm for Supergames with State Variables

1. Inputs:
   (a) Set of normals:

$$R = \{r^1, \ldots, r^m\}.$$

   (b) Set of boundary points for each state:

$$Z_x = \{z_x^1, \ldots, z_x^m\}.$$

Inputs (a) and (b) define the correspondence W, where for each $x \in X$

$$W_x = \bigcap_{j=1}^{m} \{b \in \Re^N \mid r^j \cdot (b - z_x^j) \leqslant 0\}.$$

2. Choose search subgradients:

$$S = \{s^1, \ldots, s^l\}.$$

3. Compute the new subgradient and boundary point sets, R^+ and Z^+, that together represent an outer approximation to $B(W)$. For each $x \in X$ and each $s \in S$, do

(a) for each action profile $a \in A$, solve the linear programming (LP) problem

$$c_x^a(s) = \max_{w_{g(a,x)}} s \cdot [(1-\delta)\Pi(a,x) + \delta w_{g(a,x)}],$$
$$\text{s.t.} \quad w_{g(a,x)} \in W_{g(a,x)},$$
$$IC_x(i, a, \tilde{a}_i, w) \geqslant 0, \quad \forall \tilde{a}_i \in A_i, \; i = 1, \ldots, N.$$

Let $w^*_{g(a,x)}(s)$ be an arg max of the previous problem.

Let $v_x^a(s) = (1-\delta)\Pi(a,x) + \delta w^*_{g(a,x)}(s)$ be the corresponding vector of player payoffs.

If the above problem is not feasible, then $c_x^a(s) = -\infty$, $w^*_{g(a,x)}(s) = \varnothing$, and $v_x^a(s)$ is a vector of $-\infty$.

(b) Choose the action profile that maximizes the weighted value

$$a_x^*(s) \in \arg\max_a c_x^a(s).$$

Let $v_x^*(s) = v_x^{a^*}(s)$ be the corresponding vector of payoffs and $c_x^*(s) = c_x^{a^*}(s)$ the corresponding vector of weighted payoffs.

4. Update R and Z:

(a) The new set of normals is

$$R^+ = S.$$

(b) The new set of boundary points is

$$Z_x^+ = \{v_x^*(s) \mid s \in S\}.$$

The sets R^+ and S^+ together define the outer approximation to $B(W)$, W^+. For each x,

$$W_x^+ = \bigcap_{j=1}^{l} \{b \in \mathfrak{R}^N \mid r^j \cdot b \leqslant r^j \cdot z_x^j\}, \quad r \in R^+, \; z_x \in Z_x^+.$$

(c) Check for convergence:

Stop if the Hausdorff distance between W_x^+ and W_x is less than $\epsilon > 0$ for all $x \in X$; else set $W_x = W_x^+$, $R = R^+$, $Z_x = Z_x^+$, and go back to Step 2.

The key step in our outer approximation is the collection of optimization problems we solve in Step 3. For a fixed directional search s and action profile a, each optimization problem is transformed into a LP problem in the continuation utility profile $w_{g(a,x)}$. In the objective function, the current period payoff becomes a scalar once a particular action profile is set, therefore, it is the weighted average of continuation utilities that are maximized. The constraint

$$w_{g(a,x)} \in W_{g(a,x)}$$

is replaced by a set of linear inequality constraints in $w_{g(a,x)}$ and these linear constraints define the intersection of the half spaces that represent an outer approximation to the set $W_{g(a,x)}$.

The incentive compatibility constraints in Step 3a can also be expressed as linear constraints on continuation utilities since the action space is finite. For a given search direction s, we run through the full set of action profiles, and then choose the action profile that maximizes the objective $c_x^a(s) = s \cdot [(1-\delta)\Pi(a,x) + \delta w_{g(a,x)}]$. The new maximized weighted values $\{c_x^*(s)\}$ are then used to construct a new outer approximation for each state x:

$$W_x^+ = \bigcap_{j=1}^{l} \{b \in \mathfrak{R}^N \mid s^j \cdot b \leqslant c_x^*(s^j)\}.$$

Therefore our outer approximation algorithm transforms a mapping problem in correspondences, to a series of LP problems. The inner monotone approximation uses a similar insight, the difference lies in how the approximation to the sets $\{W_x\}$ are constructed in each iteration, and how the initial correspondence is chosen.

Figure 4 illustrates features of a LP problem from Step 3 of the algorithm, having chosen a search direction s and an action profile a, for two players. Assume it's the case that the action profile involves staying in the same state $x = 1$. Therefore the continuation value profile w_1 will be chosen from the set W_1. The incentive compatibility conditions are linear constraints on these continuation values, as shown in the figure. Searching in the northeast direction, conditional on the chosen action profile, a continuation value pair w_1 that maximizes the linear objective is at the intersection of the two hyperplanes that help define the outer approximation to W_1. The incentive compatibility constraints shown are not binding. Additionally, the figure shows the continuation utility agent 1 would receive, if he or she were to deviate from the set action profile and

Figure 4. (Color online) Continuation value search

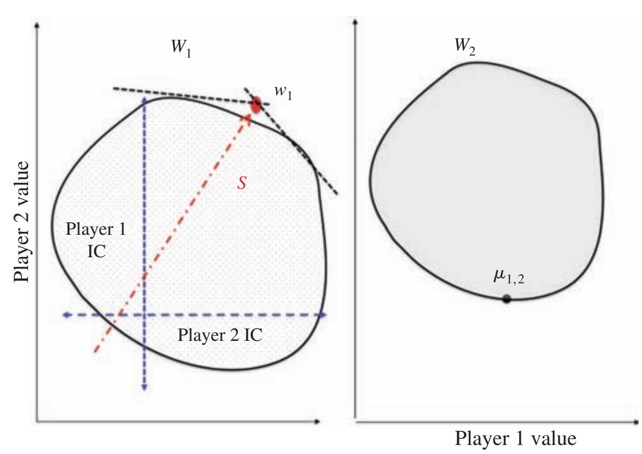

choose an action that moves the dynamic game to the aggregate state $x = 2$. The value $\mu_{1,2}$ represents the worst payoff agent 1 would receive in state $x = 2$.

Iteration and convergence criterion. The convergence criterion for computing an outer approximation of the equilibrium value correspondence is very straightforward if the search directions across iterations stay the same. In that case, each iteration produces a set of hyperplane positions determined by the weighted payoffs $\{c_x^*(s)\}$ for each state x and convergence is assumed when, for some ϵ small, the Hausdorff distance between subsequent outer approximations satisfies

$$\max_{x\in X}\left\{\max_{s\in S}\left|c_x^{*,+}(s)-c_x^*(s)\right|\right\}\leqslant \epsilon,$$

where for a specific direction $s = s^j$, $c_x^*(s^j) = s^j \cdot z_x^j$, $s^j \in S$, $z_x^j \in Z_x$, and $\{c_x^{*,+}(s)\}$ are the updated values of $\{c_x^*(s)\}$ after one outer approximation iteration.

If the search directions change across iterations, then a modified convergence criterion is needed. Following the description of our inner approximation, we provide one such criterion.

3.5. Inner Approximation of B

We now turn to the inner approximation of the B operator.

Definition 9. A mapping $B^I\colon \mathcal{P} \to \mathcal{P}$ is an inner monotone approximation of B if

1. $\forall\, W \in \mathcal{P}$, $B^I(W) \subseteq B(W)$, and
2. $\forall\, U, W \in \mathcal{P}$ if $U \subseteq W$, then $B^I(U) \subseteq B^I(W)$.

Proposition 2. *Suppose $B^I(\cdot)$ is an inner monotone approximation of $B(\cdot)$. Then, the maximal fixed point of B^I contains V. More precisely, for $W \in \mathcal{P}$, if $W \subseteq B^I(W) \subseteq V$, then $B^I(W) \subseteq B^I(B^I(W)) \subseteq \cdots \subseteq V$.*

Proof. Monotonicity of B^I implies that if $W \subseteq B^I(W)$, then $B^I(W) \subseteq B^I(B^I(W))$, etc. By definition of B^I, if $W \subseteq B^I(W)$, then $W \subseteq B(W)$, which implies that $W \subseteq V$. Then, monotonicity also implies $B^I(W) \subseteq B(W) \subseteq B(V) = V$, $B^I(B^I(W)) \subseteq B(B^I(W)) \subseteq B(V) = V$, etc.

The previous proposition establishes that applying the inner monotone approximation of the B operator repeatedly to a carefully chosen initial set yields an approximation to the equilibrium value correspondence that is contained in the equilibrium value correspondence. In APS language, the initial correspondence must be a self-generating correspondence. Unlike the outer approximation, we do not have an obvious candidate for the initial inner correspondence that leads to convergence. In practice, however, it is possible to find such a correspondence. We discuss this in more detail in Section 3.6.

Inner Monotone Approximation Algorithm for Supergames with State Variables

1. Input:
   (a) Sets of points for each state:

$$Z_x = \{z_x^1, \dots, z_x^{m_x}\},$$

which represents the correspondence W, where for each x, $W_x = \mathrm{co}(Z_x)$.

2. Choose search subgradients

$$S = \{s^1, \dots, s^l\}.$$

3. Compute new Z^+ that represents an inner approximation of $B(W)$. For each $x \in X$, $s \in S$, do:

   (a) For each action profile $a \in A$, solve the LP problem:

$$\begin{aligned} c_x^a(s) = {}& \max_{w_{g(a,x)}} s\cdot[(1-\delta)\Pi(a,x)+\delta w_{g(a,x)}], \qquad (1)\\ &\text{s.t.}\quad w_{g(a,x)} \in \mathrm{co}(Z_{g(a,x)}),\\ &\qquad IC_x(i,a,\tilde a_i,w)\geqslant 0, \quad \forall\, \tilde a_i \in A_i,\ i=1,\dots,N. \end{aligned}$$

Let $w^*_{g(a,x)}(s)$ be an arg max of (1) and define

$$v_x^a(s) = (1-\delta)\Pi(a,x) + \delta w^*_{g(a,x)}(s)$$

to be the corresponding vector of player payoffs. If (1) is not feasible, then $c_x^a(s) = -\infty$, $w^*_{g(a,x)}(s) = \varnothing$ and $v_x^a(s)$ is a vector of $-\infty$.

   (b) Record the values $v_x^a(s)$ that support the same action profiles for each x in $\{\Omega_x^a\}$:[11]

$$\Omega_x^a = \{v_x^a(s) \mid v_x^a(s) > -\infty,\ s \in S_x\}.$$

   (c) Choose action profile that maximizes the weighted value:

$$a_x^*(s) \in \arg\max_a c_x^a(s).$$

Let $v_x^*(s) = v_x^{a^*}(s)$ be the corresponding vector of payoffs.

4. Update correspondence Z: For each $x \in X$,

$$Z_x^+ = \{v_x^*(s) \mid s \in S\}.$$

The set of Z_x^+ represents an inner approximation of $B(W)$ and for each x, $W_x^+ = \mathrm{co}(Z_x^+)$.

5. Check for convergence:

Stop if the convergence criterion is satisfied, else set $W_x = W_x^+$, $Z_x = Z_x^+$, and go back to Step 2.

Iteration and convergence criterion. The convergence criterion for the outer monotone approximation is based on the Hausdorff distance between successive approximations of the outer polytope. A similar criterion can be applied to the inner monotone approximation, where the distance between $W_x^+ = \mathrm{co}(Z_x^+)$ and $W_x = \mathrm{co}(Z_x)$ can be computed, and iterations continued until the biggest of such distances amongst x is smaller than some epsilon. That is, the distance between W and W^+ is defined as

$$\max_{x\in X} d(W_x, W_x^+),$$

where the distance between W_x and W_x^+ is defined to the maximum of the distances between W_x^+ and the extreme points of W_x, and the distances between W_x and the extreme points of W_x^+, i.e.,

$$d(W_x, W_x^+) = \max\left\{\max_{z\in Z_x}\min_{w\in W_x^+}\|z-w\|, \max_{z\in Z_x^+}\min_{w\in W_x}\|z-w\|\right\}.$$

Note that the distance between W_x^+ and one extreme point of W_x, z is defined to be the shortest distance between z and all boundary and interior points of W_x^+. Thus, if z is an interior point of W_x^+, then the distance between z and W_x^+ is 0. This applies symmetrically to the definition of the distance between W_x and one extreme point of W_x^+.[12]

3.6. Error Bounds and Initial Guesses

The inner and outer monotone approximation schemes have two main differences. First, the outer monotone approximation iterations produce a sequence of correspondences that monotonically shrink toward the equilibrium value correspondence. The inner monotone approximation proceeds in the opposite way; it produces a sequence of monotonically increasing correspondences that converge toward the equilibrium value correspondence. Second, in each iteration, the value sets are constructed differently; in the outer monotone approximation, the intersection of the half spaces defined by the search normals and boundary points approximate the value sets. In the inner monotone approximation, the convex hull of the boundary points is used to construct the sets.

While the outer monotone approximation is an important tool for ruling out values from the equilibrium value sets, the inner monotone approximation is used for verifying that values are part of the equilibrium value sets. In other words, any value outside of the outer monotone approximation is *not* an equilibrium value, while any value inside the inner monotone approximation *is* an equilibrium value. These inner and outer approximations "sandwich" the true equilibrium value correspondence. Therefore we can calculate the approximation error at each finite iteration, which is unusual for iterative algorithms, by finding the distance or the area between the inner and outer correspondences. Figure 5 illustrates such errors for the $N = 2$ case. The errors are marked with the shaded areas, corresponding to the area between the inner and outer approximation boundaries. In our computed examples, the errors are very small; the inner and outer approximations look almost identical visually.

Figure 5. (Color online) Inner and outer approximations and errors

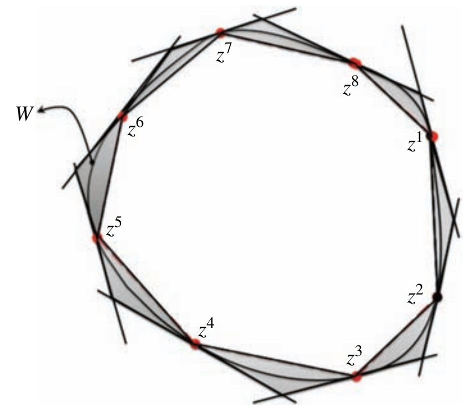

Since the B operator is a monotone operator, initial correspondence used in the inner and outer monotone approximations have to be carefully chosen. Lemma 1 establishes that using $\Omega = \times_{i=1}^{N}[\underline{\Pi}_i, \overline{\Pi}_i]$ as the initial correspondence guarantees that the outer approximation contains the equilibrium value correspondence. We do not have an analogous lemma for the inner approximation. However, in practice, we have found that first computing the outer approximation and then using a set contained within the outer approximation as the initial set for the inner approximation iterations works very well. Typically, we "shrink" the outer approximation correspondence and apply the self-generation test. In other words, we check if the shaved outer approximation, $\tilde{W}^O$, satisfies the condition $\tilde{W}^O \subseteq B^I(\tilde{W}^O)$ before proceeding with the inner approximation iterations.

3.7. Computing Equilibrium Paths

Our inner and outer approximations reveal all of the equilibrium *values* of the dynamic game, but we are also interested in the action profiles associated with these equilibrium values, and how these profiles may evolve along an equilibrium path. In the final stage of our computational procedure, we provide an algorithm to construct *sample* equilibrium paths. Each value in the equilibrium set may be supported by more than one action profile and continuation value profile, therefore, in addition to the multiplicity in equilibrium values, there is multiplicity in the action profiles that support a particular value. Our third algorithm delivers *sample* paths; it is not intended to recover strategies, which are complicated functions of histories.

To construct a sample equilibrium path, we first choose a starting equilibrium value, b_x and find an action profile a and continuation value profile $w_{g(a,x)}$ that support b_x. We then find an action profile and a continuation value profile that together support $w_{g(a,x)}$ and continue iterating in this manner. Note that the algorithm uses the information stored during the inner approximation scheme, specifically, the extreme points whose convex hull define an inner approximation and also the Ω_x^a sets that record the action profiles for the extreme points of V_x^I for each x. The precise steps are as follows.

Sample Equilibrium Path Algorithm

1. Input: Set Z_x and $b_x \in V_x^I = \text{co}(Z_x)$.
2. Find equilibrium actions and continuation values that support b_x.

(a) For each $a \in A$, check if b_x belongs to some Ω_x^a. If yes, then compute $w_{g(a,x)} = \delta^{-1}[(b_x - (1-\delta)\pi(a,x)]$. Record $v_x = b_x$ as the equilibrium value in path, set $b_x = w_{g(a,x)}$, and go back to Step 1 to continue path. If no, proceed to Step 2b.

(b) For each $a \in A$, calculate $w_{g(a,x)} = \delta^{-1}[(b_x - (1-\delta)\pi(a,x)]$ and check constraints:

$$w_{g(a,x)} \in V_{g(a,x)}^I$$

and

$$IC_{x,V_{g(a,x)}^I}(i, a, \tilde{a}_i, w_{g(a,x)}) \geqslant 0, \quad \forall\, \tilde{a}_i \in A_i,\ i = 1, \ldots, N.$$

If there is an a and associated $w_{g(a,x)}$ that can support b_x, then record $v_x = b_x$ as the equilibrium value in path, set $b_x = w_{g(a,x)}$, and go back to Step 1. If not, proceed to Step 2c.

(c) Solve for $0 \leqslant \lambda_i \leqslant 1$, $i = 1, \ldots, m_x$ such that

$$b_x = \sum_i^{m_x} \lambda_i z_x^i \quad \text{and} \quad \sum_i^{m_x} \lambda_i = 1,$$

where $\text{co}(\{z_x^i\}) = V_x^I$.

Then, pick a value z_x^i with $\lambda_i > 0$, according to a randomization device. Replace $b_x = z_x^i$, and go back to Step 1.

It is possible (and quite probable) that any b_x can be implemented by more than one pure strategy action profile. In this case, we choose an a according to some criterion. Examples of criteria that can be used include maximizing total current profits, or minimizing current profits, or symmetry in current payoffs. If the algorithm proceeds to Step 2c, then public randomization is used to support the particular b_x value by making a random choice from the set $\{z_x^i \mid \lambda_i > 0\}$.

4. Application: Dynamic Oligopoly with Endogenous Capacity

This section provides an example of a dynamic game with state variables to illustrate our numerical method. Specifically, it extends a standard textbook oligopoly model to a dynamic setting with capacity investment. This particular choice of application is not arbitrary; similar models have been studied extensively in game theory and industrial organization literatures, but due to the intractability of these models, only partial characterization of the solution has been possible. Despite these difficulties, the literature on oligopolistic firm behavior has provided, amongst others, the following two important insights. First, the resulting outcome is influenced by the strategic variable firms employ: price versus quantities. Second, regardless of the strategic variable firms employ, results are highly sensitive to whether a dynamic or a static model is used. In the oligopoly game analyzed in this section, we allow firms to use multiple strategic variables (quantity and capacity) and the environment is dynamic.

One of the most compelling reasons for studying dynamic oligopoly models with endogenous capacity is to understand how strategic capacity choices can affect equilibrium outcomes, especially when building and reducing capacity come with nontrivial costs. Capacity can limit both incentives to deviate (e.g., to undercut rivals) and retaliation possibilities. For example, in the static models employed by Spence (1979), Fudenberg and Tirole (1983), increased capacity enables the incumbent firm to threaten the would-be entrant with lower profits in the postentry equilibrium, thus discouraging entry. Benoit and Krishna (1991), however, show that when the model is dynamic, a completely different finding emerges: *lower* capacity could deter entry because higher capacity increases the prospects for tacit collusion.[13] Because solutions to even the simplest of such models are quite difficult to characterize, especially for arbitrary discount factors, much of the literature on dynamic oligopoly imposes restrictions on the environment, choice variables or equilibria. The most popular amongst these are allowing firms to choose capacities only in the initial period, concentrating on symmetric stationary equilibria or choosing high discount factors. These restrictions can limit the scope of applications and the range of outcomes—for example, a realistic study of firm splits, mergers, transfers, and competition require an ability to consider asymmetric costs or capacities—and in most cases, they are not enough to rule out multiplicity of equilibria.[14]

Our method is not intended to reduce the multiplicity problem that arise in dynamic games, but to uncover all SPE value sets and how these sets change when the primitives of the underlying game are modified. The following results are illustrative of such comparative static exercises, and of sample equilibrium paths to demonstrate our method, and therefore are not meant to be a comprehensive analysis of dynamic oligopoly with capacity constraints. We first describe the particular infinite-horizon dynamic oligopoly model in detail and then provide numerical solutions based on our algorithms.

The environment is as follows. Firms, indexed by i, with $i \in \{1, \ldots, N\}$ participate in a market for a finished good with linear demand. Each firm's technology is augmented by the incorporation of a capacity constraint, which depends on the total number of machines it possesses. If firm i has $k_i \in K = \{0, 1, \ldots, nk\}$ machines, then it can choose an output from the set $Q_i = [0, \bar{Q}(k_i)]$, where the maximum output it can produce in the current period is bounded above by $\bar{Q}(k_i)$. Firms can exit the market by choosing $k = 0$ and a choice of $k > 0$ from the state $k = 0$ is interpreted as entry. Firm i has sales of $q_i \in Q_i$, and per unit cost c_i. The cost of maintaining a machine for a period is given by c_M. In every period, each firm can alter its capital/capacity stock by purchasing and installing a new machine at a cost of c_F. Machines are indivisible and become operational one period after the cost of investment is incurred. There is a market for the resale of capital at a price of p_S.

Let $C(k_i, k_i')$ denote the cost of maintaining k_i machines this period and altering the number of machines to k_i' in the next period. Then, $C(\cdot)$ takes the following form:

$$C(k_i, k_i') = \begin{cases} c_M \cdot k_i + c_F \cdot (k_i' - k_i) & \text{if } k_i' \geqslant k_i \\ c_M \cdot k_i - p_S \cdot (k_i - k_i') & \text{if } k_i' \leqslant k_i. \end{cases}$$

Firm i's current profits are given by

$$\Pi_i(q, k_i, k_i') = q_i(p(q) - c_i) - C(k_i, k_i'),$$

where c_i is the per unit production cost for firm i. The market price, $p(q)$ is given by a linear demand curve and q represents the vector of quantities $\{q_1, \ldots, q_N\}$, chosen by each firm:

$$p(q) = \max\left\{\Gamma - \gamma \sum_{i}^{N} q_i, 0\right\}, \tag{2}$$

where Γ and γ are two parameters. The capacity constraint takes the following linear form,

$$\bar{Q}(k_i) = \rho k_i, \tag{3}$$

so that the incremental increase in maximum output a firm can produce is set to ρ.

The stage game action spaces for this problem are sets of outputs and capital stocks (i.e., machines), for each firm. The state space is the set of feasible capital stocks for all firms. Thus, $A_i = Q_i \times K$ and $X = K^N$. Firm strategies are collections of functions that map from histories of outputs and capital stocks into current output and capital choices. As usual, each firm attempts to maximize its average discounted profits.

In the next subsections, we show a variety of numerical results from four different parameterizations of our capacity game with two firms. Table 1 reports the parameter values for each of the four cases. Case 1 represents our benchmark parameterization with $nk = 2, \ldots, 5$ and features reversible investment. Case 2 is used to investigate the effects of irreversibility of investment. Case 3 represents the game with a low discount factor and irreversible investment. Case 4 is the asymmetric cost case. The parameters that govern the inverse demand function (Γ and γ) and the capacity constraint (ρ) are set to $\Gamma = 6$, $\gamma = 0.3$, and $\rho = 3$ for all of the computed examples. Additionally, we provide an example from a game with three firms. Further details about the dynamic game and our numerical procedure, including the solution to the monopolist's problem for each case, as well as action grids, number of hyperplanes, and running times for our benchmark case are also provided.

Table 1. Cases and parameter values

	Investment	nk	δ	c_1	c_2	c_F	c_M	s_P
Case 1	Reversible	2–5	0.8	0.9	0.9	2.5	1.5	1.5
Case 2	**Irreversible**	2–5	0.8	0.9	0.9	2.5	1.5	N/A
Case 3	Irreversible	3	**0.4**	0.9	0.9	2.5	1.5	N/A
Case 4	Irreversible	3	0.8	**0.1**	0.9	2.5	1.5	N/A

Table 2 provides the solutions to the monopolist's problem for each case. In Case 4, we assume the monopolist has per unit cost of 0.1. Also, highlighted are the steady states, assuming that the monopolist is endowed with 0 capacity in the initial period and builds up its capacity after entry.

Before we proceed with the oligopoly results, we define the terminology we use for certain types of strategic behavior. *Tacit collusion* refers to equilibria with strategy profiles that generate the monopoly outcome. If the joint profits of the firms equal, the monopoly profits in that state, and the production/investment levels collectively match the monopoly choices, we refer to these outcomes as tacit collusion. When the monopoly profits and production are shared equally by the firms, we refer to it as the *symmetric collusive* outcome. *Joint profit maximization* is a type of collusion, typically associated with cartels or explicit coordination, rather than tacit collusion.[15]

4.1. Cases 1 and 2: Reversible vs. Irreversible Investment

Case 1 serves as a benchmark for comparison across various parameterizations. We also use it to display and discuss certain features (initial correspondences, B operator, action grids, etc.) of our algorithm.

Initial Correspondence. In Case 1, because investment is reversible and entry and exit are possible, in equilibrium, firms cannot be forced to earn profits less than zero. The best they can earn in each state is the monopoly profit associated with that state. Therefore we make sure the initial correspondence for the outer approximation includes the correspondence $\{[0 - \epsilon, v^m(k) + \epsilon]^2\}_{k=1}^{nk}$. The margin of ϵ is added to

Table 2. Monopoly values and choices

Case 1				Case 2				Case 3				Case 4			
k	q	k'	v^m	k	q	k'	v^m	k	q	k'	v^m	k	q	k'	v^m
0	0	2	12.44	0	0	2	12.44	0	0	2	3.72	0	0	3	17.94
1	3	2	15.16	1	3	2	15.16	1	3	2	11.88	1	3	3	21.14
2	6.0	2	16.80	2	6.0	2	16.80	2	6.0	2	16.80	2	6	3	23.26
3	8.5	2	17.18	3	8.5	3	17.18	3	8.5	3	17.18	3	9.0	3	24.30
4	8.5	2	17.18	4	8.5	4	15.68	4	8.5	4	15.68	4	9.9	4	23.01
5	8.5	2	17.18	5	8.5	5	14.18	5	8.5	5	14.18	5	9.9	5	21.51
6	8.5	2	17.18	6	8.5	6	12.68	6	8.5	6	12.68	6	9.9	6	20.01

ensure that numerical inaccuracy does not eliminate potential equilibrium values from the initial correspondence. The monopoly values $\{v^m\}$ for each state and case, including Case 1, can be found in Table 2.

Search subgradients. In the two-dimensional case, we typically use 36 subgradients to span each set. This can be done by setting search subgradients to $(\cos(\theta), \sin(\theta))$, where $\theta \in \{0, \dots, 2\pi\}$. When we want to have the subgradients uniformly distributed, we set the increments between each θ value to $2\pi/35$. In most cases, we keep the search subgradients the same across states and iterations. We have also used an adaptive scheme for the search subgradients, adding more of them in the directions of large changes in the boundaries between iterations, reducing them in the directions of no or very little changes in boundaries between iterations to increase accuracy and speed. All the results reported in Section 4, however, are from the uniformly distributed subgradients.

B operator. The main part of our algorithm is the implementation of the B operator. It involves solving a series of LP problems to find the new values that are incentive compatible and that are supported by continuation values from W. The number of LP problems per iteration depends on the number of aggregate states, number of discrete actions (quantities produced and capacity investment), and number of search subgradients. For example, for our benchmark case, Case 1, we set $nk = 3$, number of quantities to 91, and the number of search subgradients to 36. Without taking advantage of symmetry, this results in more than 19 million LP problems per iteration. For this particular example, the outer approximation took 82 iterations for convergence and the inner approximation had 65 iterations. With 640 cores, the outer approximation took 3.2 minutes and the inner approximation 2.6 minutes to converge.

Figure 6 displays the equilibrium value set, for state (1, 1), at different iterates of the inner approximation algorithm. The initial equilibrium value set, which is the innermost polygon that corresponds to iteration 0, is small and is chosen to lie within the equilibrium value set.[16] With the repeated application of the B^I operator, the value sets enlarge and monotonically converge (plotted are the initial set, iterations 1 and 20) to the value set labeled as the inner approximation. Figure 7 shows the inner and outer approximations for

Figure 6. (Color online) Monotonic convergence

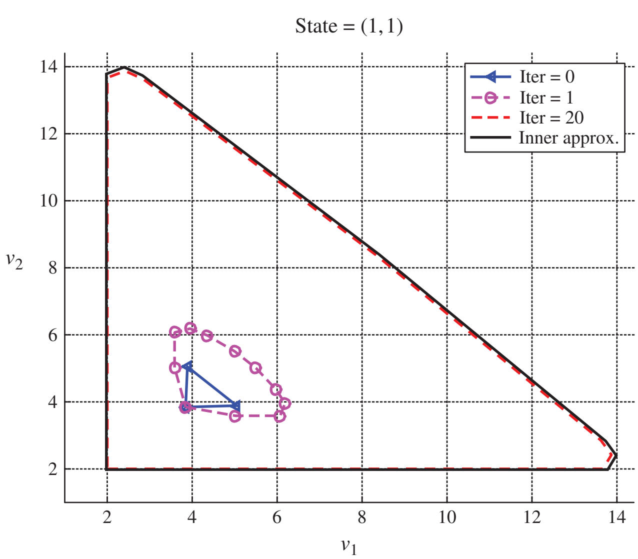

Figure 7. (Color online) Inner vs. outer

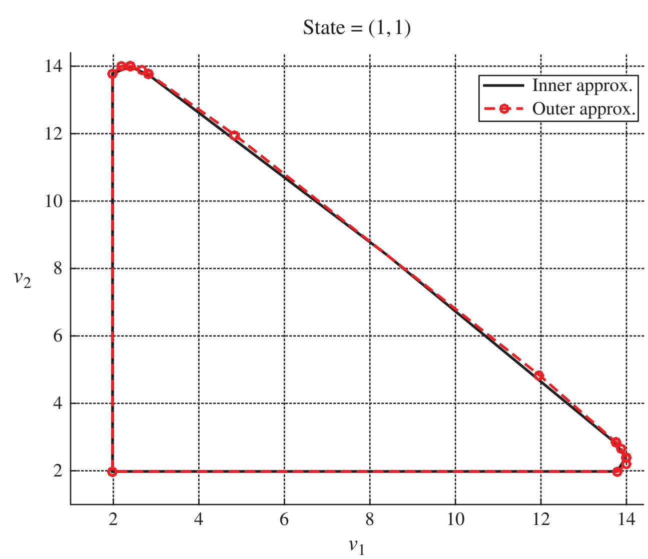

Figure 8. (Color online) Case 1: Values

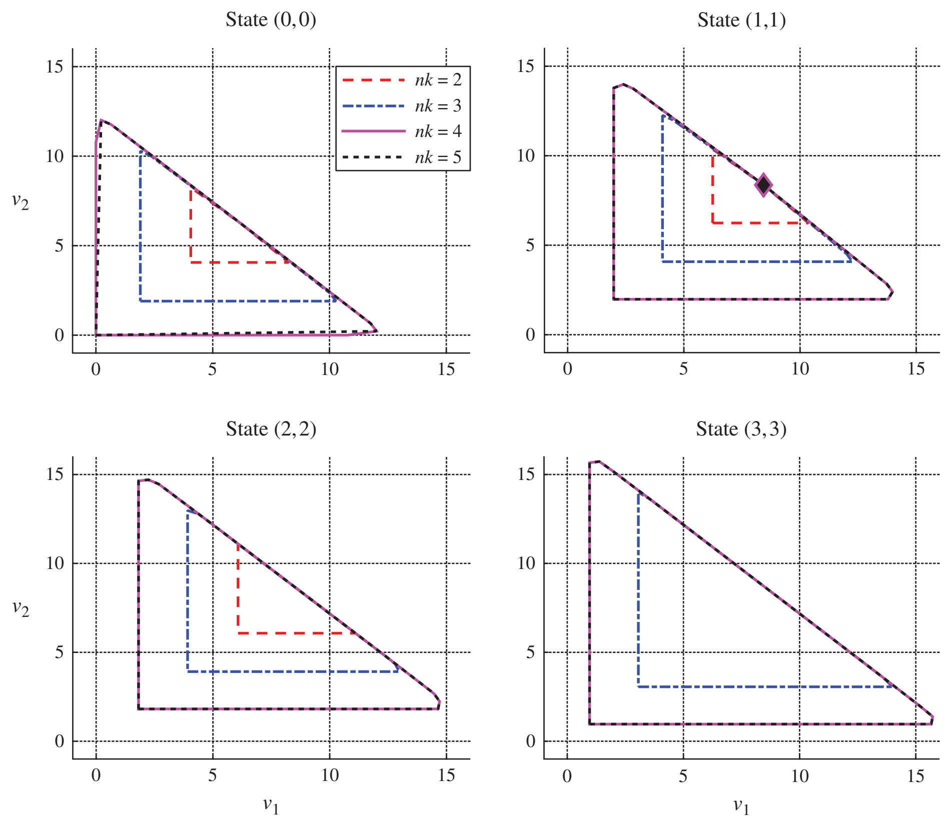

Case 1, state $(1,1)$ together. The two approximations are nearly identical, with the exception of small differences in the northeast region. The figure confirms the accuracy of our results since the inner approximation is contained within the outer approximation, and the error region between the two value sets is very small. Although Figures 6 and 7 show the results from state $(1,1)$ only, the results are similar for all the other aggregate capacity states.

In Figure 8, we show the Case 1 inner approximations for the symmetric states $(0,0)$, $(1,1)$, $(2,2)$, and $(3,3)$ for maximum allowed capacity ranging from $nk = 2$, to $nk = 5$. We are interested in the effect of nk on equilibrium values because industrial policies such as rules about mergers and acquisition, import quotas, or environmental regulations may limit the amount of *maximum* capacity a firm can hold. We do, however, ensure that the choice of nk does not make the optimal monopoly capacity and production choices infeasible, although in equilibrium, they may not be achieved by a single firm. As Table 2 shows, a monopolist facing the same market in Case 1 would choose two machines and produce six units of output, earning a payoff of 16.8 in steady state. If monopoly profits were to be equally shared, each firm would have one capacity, produce three units of output and earn 8.4. This payoff is marked with a diamond in state $(1,1)$. In this example, tacit collusion is possible in equilibrium since all values on the frontier (northeast direction) sum to monopoly profits for all $nk \in \{2,3,4,5\}$. The ability or inability to accumulate more capacity does not rule out tacit collusion. However, with the ability to build more capacity, the maximum and minimum values an individual firm can achieve does change. As nk increases, the equilibrium value sets enlarge, until the lowest equilibrium values (nearly) reach zero. Since negative values are not supported in equilibrium when free entry and capital reversal are allowed, the equilibrium value sets for $nk = 5$ cannot extend anymore, and hence overlap with the $nk = 4$ sets.

As mentioned before, the best symmetric stationary equilibrium in Case 1 involves one machine and output $q = 3$ per firm, which would deliver an equilibrium value of 8.4 for each firm. Therefore, in our discussion of the numerical results involving Case 1, we refer to any aggregate capital state exceeding two machines and aggregate production exceeding six units as overcapacity and overproduction. One interesting question that arises is whether firms can reach this best symmetric equilibrium from zero initial capacity. In Figure 9, we display an equilibrium path with nodes numbered to reflect the time sequence of equilibrium play, starting from the lowest computed value (node 1) in the game with zero machines for each firm. The path to cooperation involves a brief period of overaccumulation of capacity as firms move to the state

Operations Research, 2017, vol. 65, no. 2, pp. 337–356, © 2017 INFORMS

Figure 9. (Color online) Case 1: Sample path

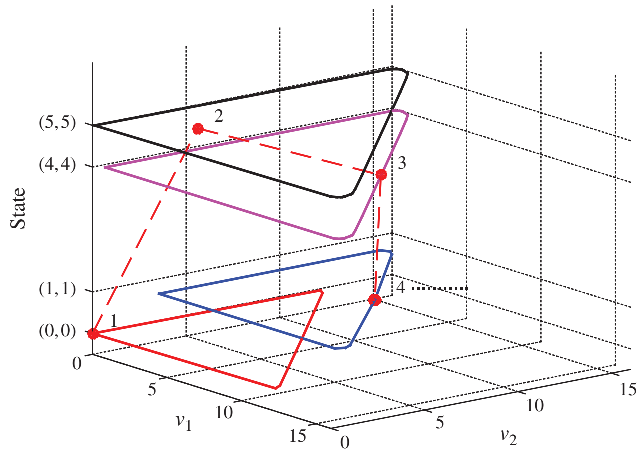

(5, 5) and then (4, 4), but quickly reaches the symmetric best payoff in the (1, 1) state. Note that although there may be equilibrium paths that start from node 1 and not reach this symmetric best SPE, one way to get to that best symmetric payoff from a much worse payoff involves costly excessive investment. Besanko et al. (2010) find similar equilibrium behavior in a dynamic model of an oligopolistic industry with lumpy capacity and lumpy investment or disinvestment although they focus on Markov-perfect equilibria. They display equilibrium paths in which excess capacity in the short run often go hand-in-hand with capacity coordination in the long run.

Figure 9 features a sample path of firms reaching the best stationary symmetric equilibrium after a period of excess capacity. Because investment is reversible, excess capacity for a brief period is not very costly for the firms, especially if it helps ensure a higher payoff later on. A question of interest is how irreversibility of investment alters the equilibrium value correspondence and outcomes. Figure 10 compares the equilibrium values in four symmetric states for the reversible (Case 1) and irreversible (Case 2) cases.[17] Player payoffs are very similar in states (0, 0) and (1, 1), but the cost of irreversibility becomes apparent when firms are in higher capacity states. The mandatory maintenance cost hurts the firms and the best they can achieve in these large capacity states is well below the best they can achieve in the flexible investment case. One would expect that given the irreversibility of investment, firms would shy away from overinvestment in capacity. The equilibrium path in Figure 11, however, shows the firms moving from zero capacity to five machines at the same time. Interestingly, they quickly move (see nodes 2–4) to the best symmetric equilibrium value (node 4) in state (5, 5) and stay there. At node 4, there is not only overcapacity, but also overproduction with aggregate output equal to 8.5 instead of 6.

Figure 10. (Color online) Irreversible vs. reversible investment: Value sets

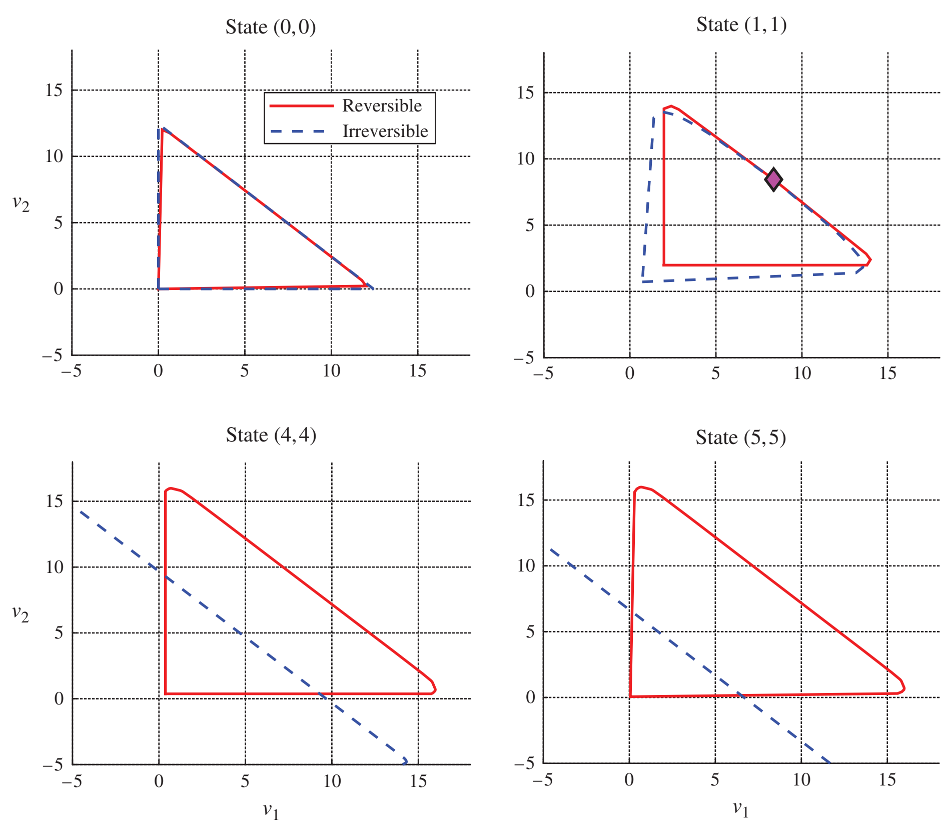

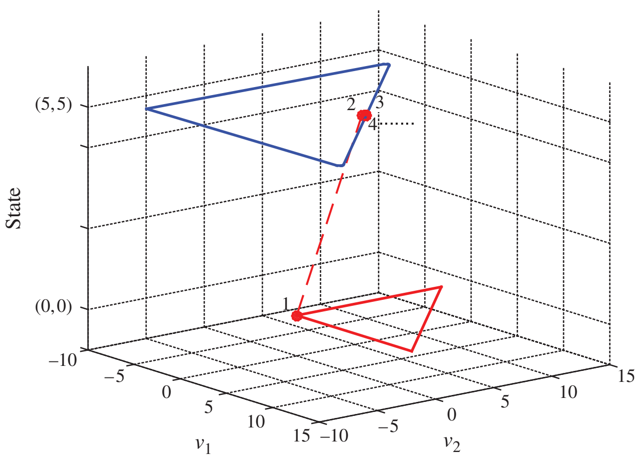

Figure 11. (Color online) Irreversible case sample path 1

Although Case 1 produces equilibrium outcomes that involve tacit collusion (see Figure 9) following a brief period of overcapacity by both firms, it also has an equilibrium path in which the road to collusive behavior involves one firm acting as a monopolist for a period of time before the second firm enters. This path is displayed in Figure 12. Firm 2 enters first with three machines, so the state is $(0,3)$ for a brief period, followed by some disinvestment by Firm 2, therefore the state moves to $(0,2)$. While in state $(0,2)$, Firm 2's continuation value slightly decreases while Firms 1's increases (see nodes 9 to 16), which is supported by Firm 1 entering with 1 machine, and Firm 2 disinvesting further to move to state $(1,1)$ (see node 17). The payoffs in node 17 in Figure 12 are jointly equal to monopoly profits, and the firms stay there for the remainder of the game. Given that the values associated with node 1 are positive $(0.22, 12.01)$, it is not surprising that at some point firm 1 enters and earns positive profits, but it is surprising that a path to collusion involves a period of monopoly power for one

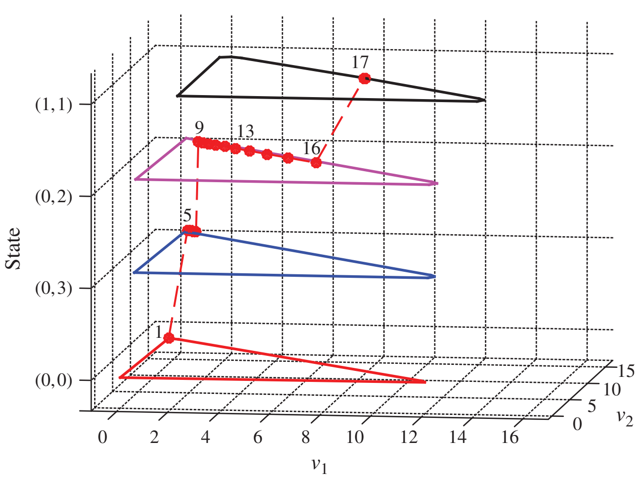

Figure 12. (Color online) Reversible case sample path

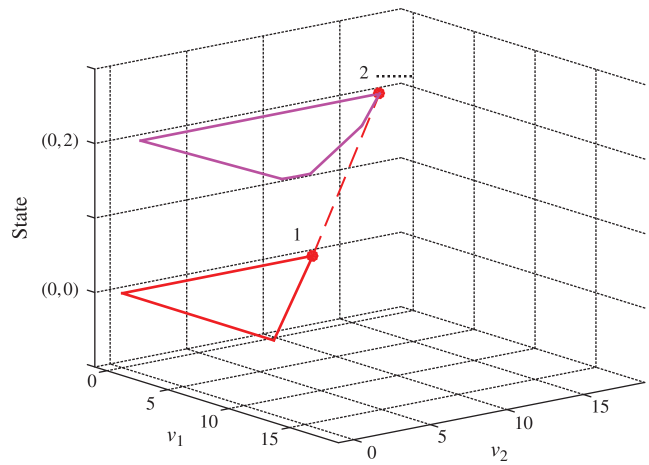

Figure 13. (Color online) Irreversible case sample path 2

of the firms. Entry by the competing firm is deferred rather than deterred. When investment is irreversible, however, a payoff of 0 is in the value set in state $(0,0)$, which implies that there is an equilibrium outcome in which one firm never enters. Figure 13 shows one such path, where Firm 2 enters with two machines first, deters entry, and operates permanently as a monopoly.

Our algorithm is designed to deliver all SPE values and to construct sample equilibrium paths. It is not intended to be used to uncover strategies, which are complicated objects of histories. In the dynamic oligopoly literature (and elsewhere), it is common practice to restrict attention to Markov-perfect equilibria (MPE) for tractability, especially since Markov strategies do not condition on entire histories. One would ideally like to know if such restrictions on strategies alter firm competition and cooperation in significant ways, however, there is no available method that reliably computes *all* MPE of infinite-horizon dynamic games, therefore comparisons between all MPE values to all SPE values cannot be made. Nevertheless, we can determine if a particular equilibrium path is *not* MPE. For example, the path in Figure 12 is not part of an MPE, because in states $(0,2)$ and $(0,3)$, firms' action profile includes staying in the same state, but also switching to another state deterministically.[18]

4.2. Cases 2 and 3: Low vs. High Discounting

One of the advantages of our algorithm is the ability to compute equilibrium value correspondences for any discount factor and not confine ourselves to large discount factors to rely on folk theorem types of results. For illustrative purposes, we have computed the irreversible investment case (Case 2) using $\delta = 0.4$ (Case 3). This particular example also shows that our algorithm can handle equilibrium value sets that are reduced to singletons, while delivering some insights about how discounting can affect firm behavior and equilibrium

Figure 14. (Color online) Low vs. high discounting

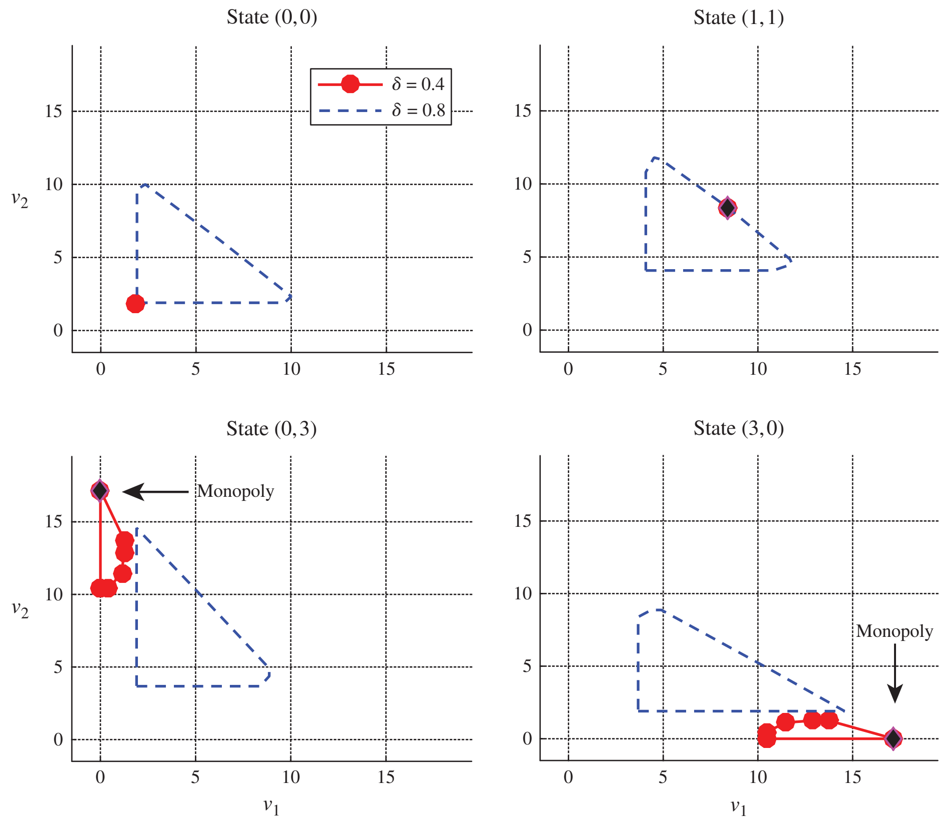

outcomes. Discounting also typically affects the convergence speed in iterative methods aimed at solving dynamic problems, but we did not find it a significant constraint on our ability to compute value sets efficiently. We used 36 subgradients, 91 quantities, maximum 3 machines for each firm in computing Case 2 and Case 3 using 640 cores. The $\delta = 0.8$ case took 83 (outer) and 76 (inner) iterations for convergence, for a total computing time of 15 minutes. The $\delta = 0.4$ case took 27 (outer) and 20 (inner) iterations for convergence, for a total time of 8 minutes.

Figure 14 shows the equilibrium value sets for four different states ($nk = 3$) for high and low discounting with irreversible investment. The equilibrium value set for states $(0,0)$ and $(1,1)$ in for $\delta = 0.4$ are singletons. In state $(0,0)$, the only equilibrium value for $\delta = 0.4$ overlaps with the lowest equilibrium value of $\delta = 0.8$. In state $(1,1)$, in contrast, the only equilibrium value is the best symmetric equilibrium value. There is an equilibrium in which starting from the single value in state $(0,0)$, firms immediately jump to the best symmetric equilibrium in state $(1,1)$ and stay there forever. In other words, firms enter into the market simultaneously and move to best symmetric equilibrium and stay there for the remainder of the game. This particular path shows that cooperative behavior is possible even when firms are impatient and investment is irreversible.

Figure 14 also shows the value sets from two asymmetric states: $(0,3)$ and $(3,0)$. There are a couple of features to note. One, the equilibrium value sets for low and high δ do not overlap. Second, when firms are less patient, there is an equilibrium in which the incumbent firm, if it has large enough capacity advantage, can operate as a monopolist forever. The values associated with monopoly profits are shown on the figure. When firms are patient (and $nk = 3$), however, no firm acts as a monopolist forever. The equilibrium values in the $\delta = 0.8$ case do not include zero value for one firm and monopoly payoff for the other firm. Both firms enter the market at one time or another.[19]

Although our method is not designed to uncover strategies, the computed value sets can be used to provide insight into equilibrium behavior. Figure 15 shows the equilibrium values sets from all 16 aggregate states for $\delta = 0.4$ and $\delta = 0.8$. States $(0,0), (1,0), (0,1)$, and $(1,1)$ have a single equilibrium value each. The stationary equilibrium that delivers the best equilibrium in state $(1,1)$, which is the single equilibrium value in that state, is the continuation value for each of the single values in states $(0,0)$, $(1,0)$, and $(0,1)$. States $(0,2)$ and $(0,3)$ ($(2,0)$ and $(3,0)$) all have equilibrium values sets that include monopoly profits for Firm 2 (Firm 1) for $\delta = 0.4$. In all other states, equilibrium values are above 0. In other words, in this particular game, firms

Figure 15. (Color online) Low vs. high discounting

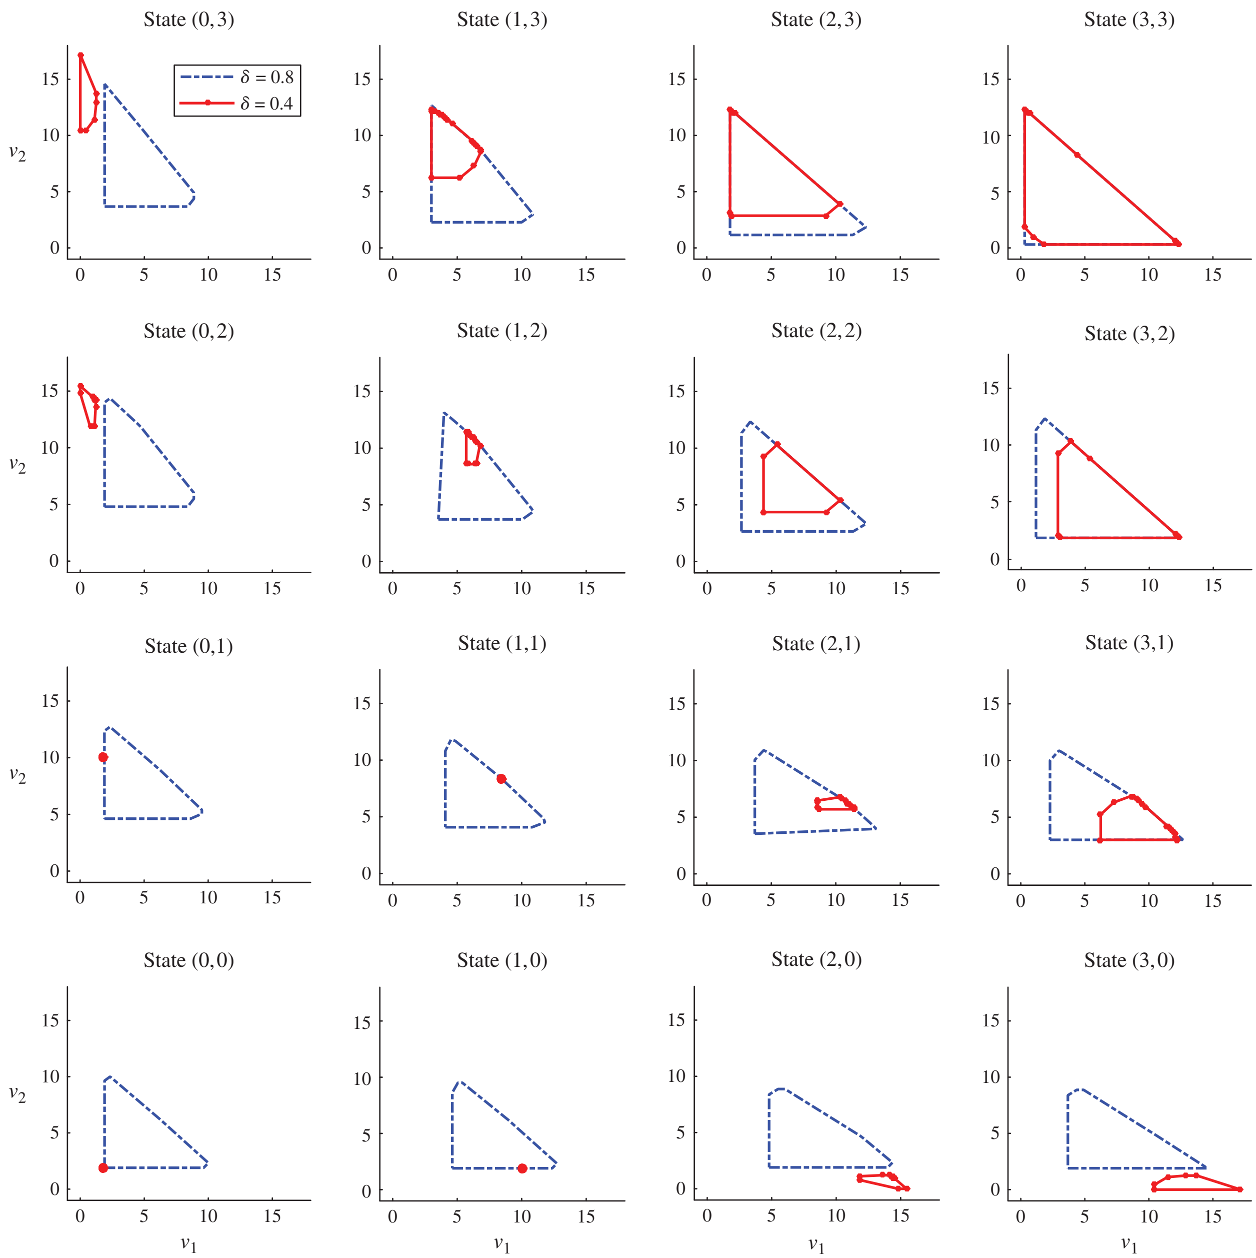

can deter entry, earn, and maintain monopoly profits indefinitely when they start with a large capacity advantage, and future profits are discounted heavily.

4.3. Case 4: Asymmetric Cost

In this example, we compute the equilibrium value correspondence of an asymmetric case with irreversible investment. Firm 1 has lower per unit production cost, $c = 0.1$, relative to Firm 2 whose per unit cost is $c_2 = 0.9$. Figure 16 displays equilibrium values from four different capacity states for Case 2 and the asymmetric cost case, Case 4, for $nk = 3$. As one would expect, in all four states, the best equilibrium payoffs for Firm 1 increase, while the best equilibrium payoffs for Firm 2 decrease. The lowest equilibrium payoff for Firm 2 does not decrease, however. This is due to the fact that Firm 1, although more cost efficient, hits its capacity constraint and cannot decrease the price further to deliver lower profits to Firm 2. The cost advantage of Firm 1 is tampered by the costly and limited capacity investment.

The select numerical results from Cases 1 to 4 show how the equilibrium value sets are altered with changes in investment reversibillity, asymmetric cost, discounting and maximum capacity firms can attain. Some of the insights gained from the comparative statics exercises for our particular game are the following. Payoffs associated with tacit collusion is part of the equilibrium values in all of the cases. Overproduction

Figure 16. (Color online) Low vs. high unit cost

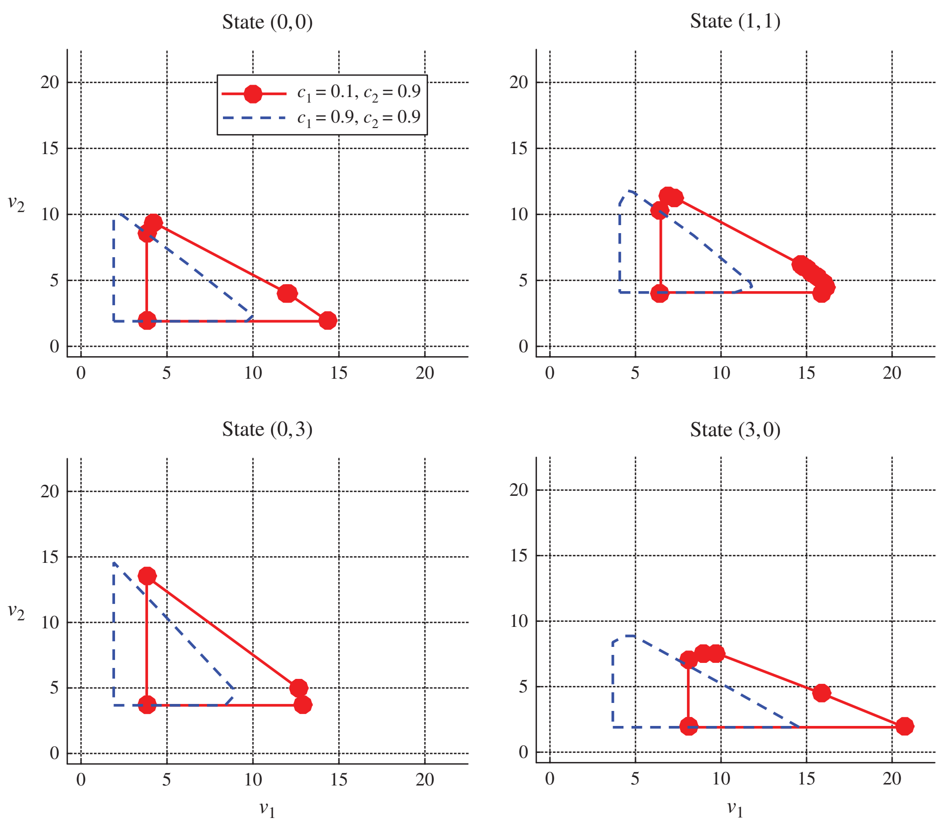

and overinvestment is present in the reversible and irreversible investment cases, albeit in the reversible case it is a temporary phase, while it can be a permanent state in the irreversible investment case. Firms can deter entry by their competitors indefinitely if they have a large enough capacity advantage, face a cap in the maximum capacity they can obtain and are sufficiently impatient, but not when they are patient. When firms face a larger maximum capacity (nk in our setting), they increase their ability to deliver lower profits to their competitor, by having the ability to increase production enough to lower the price significantly. A lower per unit production cost delivers higher equilibrium values for a firm, but does not deter entry if that cost advantage is constrained by the maximum capacity it can build. It is important to note that these are results from cases with positive maintenance costs, which naturally deters excessive capacity investment. Eliminating the maintenance cost would alter equilibrium values and behavior.[20]

4.4. Three Firms

Our method is not constrained to handle only two firms. Our algorithm is written for any N, as long as N is finite. On the practical side, there is a curse of dimensionality, as in all dynamic models with multiple strategic agents and large state spaces. However, one of the appealing features of our algorithm is that it is highly parallelizable, therefore the computational burden can be significantly reduced by parallel programming. In Figures 17 and 18, we show the equilibrium value sets for two aggregate states from a three-firm version of our capacity game. In this particular example, the maximum number of machines per firm is set

Figure 17. (Color online) Three firms: State (1, 1, 1)

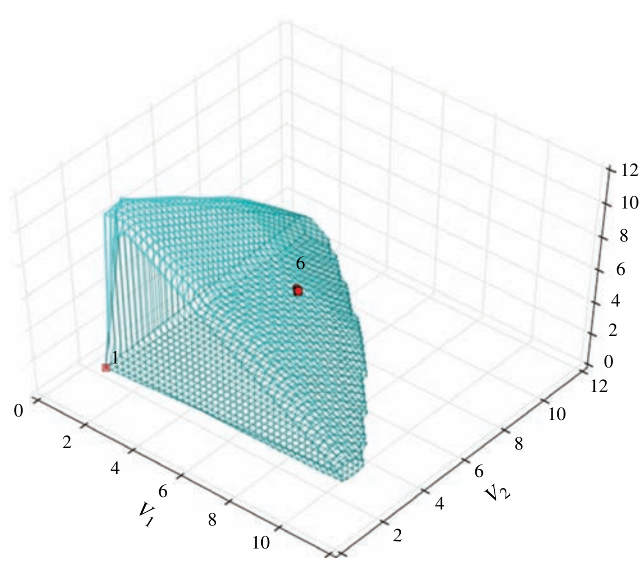

Figure 18. (Color online) Three firms: State (2, 2, 2)

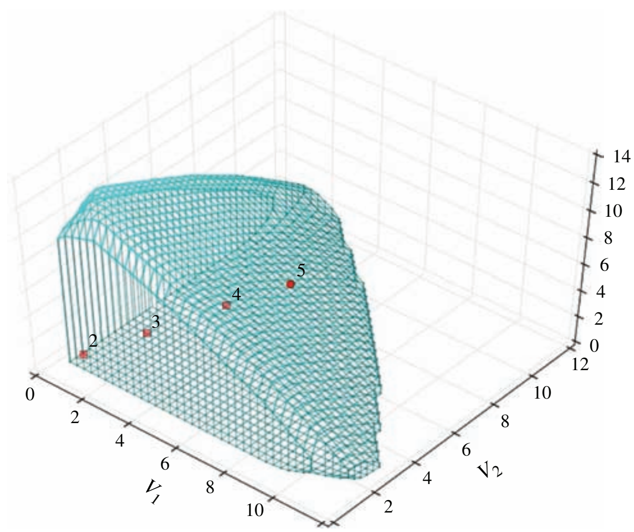

to two, therefore $nk = 2$. Otherwise, the parameters are the same as in our benchmark case, Case 1. For this particular example, we used 133 different search subgradients and 61 quantity choices per firm. This translates to almost one billion LP problems per iteration. It took 72 iterations for the outer approximation to converge, and 76 iterations for the inner approximation. With 8,000 cores, the outer approximation took 38 minutes, and the inner approximation 42 minutes.

Figures 17 and 18 also display an equilibrium path starting from the lowest symmetric value in state (1, 1, 1). The path nodes are numbered corresponding to the order of moves starting from node 1 in (1, 1, 1). Firms move to state (2, 2, 2), to an equilibrium value close to the lowest symmetric one (node 2), but then gradually move toward higher value nodes (nodes 3–5), and then jump back to the Pareto frontier of state (1, 1, 1). This path is similar to the path for Case 1 (See Figure 9), in that, the path to cooperation involves a period overinvestment and overproduction. It's also important to note that this particular path is not part of an MPE, since the equilibrium play in state (2, 2, 2) involves a period of staying in the same state and then moving to state (1, 1, 1) deterministically.

5. Conclusion

Most strategic games admit many equilibria. In the dynamic case, the multiplicity is often much more severe. One way to simplify the analysis of dynamic games is to restrict strategies or to concentrate on certain types of equilibria. While such restrictions can make problems more tractable, especially in macroeconomic policy games where conditioning policies on entire histories may be viewed as too costly or infeasible, they can eliminate equilibria and confine them to fewer applications in other settings.

In this paper, we provide an alternative approach to analyzing dynamic games. Instead of attempting to eliminate the multiplicity, we provide a numerical method for computing *all* SPE of dynamic games. Our method is based on the theory of repeated games and provides an iterative scheme, which maps convex-valued correspondences into convex-valued correspondences monotonically, and with a suitably chosen initial correspondence, delivers an accurate approximation to the true equilibrium value correspondence of the underlying game. Our algorithm has three parts. The first provides an outer approximation to the equilibrium value correspondence. This outer approximation is constructed such that any value outside of the approximation is *not* an equilibrium value. The second provides an inner approximation to the equilibrium value correspondence; any value contained within this inner approximation *is* an equilibrium value. Together, the two approximations "sandwich" the true equilibrium value correspondence and deliver a practical check of approximation accuracy. The third part of our algorithm delivers sample equilibrium paths.

We then apply our algorithm to an oligopoly competition game with endogenous production capacity and use it to compare equilibrium value sets and outcomes across different cost, discounting, maximum capacity, and investment flexibility cases. This application demonstrates how our method can be used to gain insights into equilibrium behavior in dynamic games even when multiplicity of equilibria is an issue.

Acknowledgments

The authors thank the two anonymous referees and the associate editor for their comments.

References

Abreu D, Pearce D, Stacchetti E (1986) Optimal cartel equilibria with imperfect monitoring. *J. Econom. Theory* 39(1):251–269.
Abreu D, Pearce D, Stacchetti E (1990) Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica: J. Econometric Soc.* 58(5):1041–1063.
Atallah MJ (1983) A linear time algorithm for the Hausdorff distance between convex polygons. *Inform. Processing Lett.* 17(4):207–209.
Baldauf M, Judd KL, Mollner J, Yeltekin S (2014) Supergames with states. Mimeo, Stanford University, Stanford, CA.
Benoit JP, Krishna V (1987) Dynamic duopoly: Prices and quantities. *Rev. Econom. Stud.* 54(1):23–35.
Benoit JP, Krishna V (1991) Entry deterrence and dynamic competition: The role SOF capacity reconsidered. *Internat. J. Indust. Organ.* 9(4):477–495.
Bernheim BD, Ray D, Yeltekin Ş (1999) Self-control, saving, and the low asset trap. Mimeo, Stanford University, Stanford, CA.
Bernheim BD, Ray D, Yeltekin Ş (2015) Poverty and self-control. *Econometrica* 83(5):1877–1911.
Besanko D, Doraszelski U, Lu LX, Satterthwaite M (2010) Lumpy capacity investment and disinvestment dynamics. *Oper. Res.* 58(4-part-2):1178–1193.
Conklin J, Judd KL (1996) Computing value correspondences for repeated games with state variables. Mimeo, Stanford University, Stanford, CA.
Cronshaw MB, Luenberger DG (1990) Computing value correspondences for two-player games with state variables. Mimeo, Stanford University, Stanford, CA.

Cronshaw MB, Luenberger DG (1994) Strongly symmetric subgame perfect equilibria in infinitely repeated games with perfect monitoring and discounting. *Games Econom. Behav.* 6(2):220–237.

Davidson C, Deneckere R (1986) Long-run competition in capacity, short-run competition in price, and the Cournot model. *Rand J. Econom.* 17(3):404–415.

Davidson C, Deneckere R (1990) Excess capacity and collusion. *Internat. Econom. Rev.* 31(3):521–541.

Fudenberg D, Maskin E (1986) The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* 54(3):533–554.

Fudenberg D, Maskin E (1991) On the dispensability of public randomization in discounted repeated games. *J. Econom. Theory* 53(2):428–438.

Fudenberg D, Tirole J (1983) Capital as a commitment: Strategic investment to deter mobility. *J. Econom. Theory* 31(2):227–250.

Fudenberg D, Tirole J (1991) *Game Theory* (MIT Press, Cambridge, MA).

Green EJ, Marshall RC, Marx LM (2014) Tacit collusion in oligopoly. *The Oxford Handbook Internat. Antitrust Econom.* 2:464–497.

Judd KL, Yeltekin S, Conklin J (2003) Computing supergame equilibria. *Econometrica* 71(4):1239–1254.

Kitti M (2016) Subgame perfect equilibria in discounted stochastic games. *J. Math. Anal. Appl.* 435(1):253–266.

Kreps DM, Scheinkman JA (1983) Quantity precommitment and Bertrand competition yield Cournot outcomes. *The Bell J. Econom.* 14:326–337.

Sleet C, Yeltekin Ş (2006) Optimal taxation with endogenously incomplete debt markets. *J. Econom. Theory* 127(1):36–73.

Sleet C, Yeltekin Ş (2007) Recursive monetary policy games with incomplete information. *J. Econom. Dynam. Control* 31(5): 1557–1583.

Sorin S (1986) On repeated games with complete information. *Math. Oper. Res.* 11(1):147–160.

Spence AM (1979) Investment strategy and growth in a new market. *The Bell J. Econom.* 10(1):1–19.

Sevin Yeltekin is a professor of cconomics at Tepper School of Business, Carnegie Mellon University. Her main research interests include fiscal policy, dynamic policy games, public finance and computational economics. She has published in a variety of academic journals, including *Econometrica*, *Review of Economic Studies*, *Journal of Economic Theory*, *Journal of Monetary Economics*, *International Economic Review*, *American Economic Journal: Macroeconomics*, and *Review of Economic Dynamics*.

Yongyang Cai is an associate professor at the Department of Agricultural, Environmental and Development Economics, The Ohio State University. Prior to joining OSU, he was a Senior Research Scientist at the Becker Friedman Institute, University of Chicago, and a Visiting Fellow at the Hoover Institution, Stanford University. His main research interests include computational economics, climate change economics, decision making under uncertainty, and dynamic stochastic modeling. He has published in a variety of academic journals, including *Nature Climate Change*, *Proceedings of the National Academy of Sciences*, *Computers & Operations Research*, *Journal of the European Economic Association*, and *Quantitative Economics*.

Kenneth L. Judd is the Paul H. Bauer Senior Fellow at the Hoover Institution. He is an expert in the economics of taxation, imperfect competition, and mathematical economics. His current research focuses on developing computational methods for economic modeling and applying them to tax policy, antitrust issues, macroeconomics, and policies related to climate change. His work has also been published in the *Journal of Economic Dynamics and Control*, *Journal of Public Economics*, *Journal of Political Economy*, *RAND Journal*, *Journal of Finance*, *Journal of Economic Theory*, *Brookings Papers of Economic Activity*, *American Economic Review*, and *Econometrica*. His book *Numerical Methods in Economics* was published by MIT Press in 1998. He is a fellow of the Econometric Society and served as a member of the Economics Panel of the National Science Foundation (1986–1988).