OPERATIONS RESEARCH
Vol. 62, No. 2, March–April 2014, pp. 332–344
ISSN 0030-364X (print) | ISSN 1526-5463 (online)

http://dx.doi.org/10.1287/opre.2013.1250

On Markov Equilibria in Dynamic Inventory Competition

Tava Lennon Olsen
The University of Auckland Business School, Auckland 1142, New Zealand, t.olsen@auckland.ac.nz

Rodney P. Parker
The University of Chicago Booth School of Business, Chicago, Illinois 60637, rodney.parker@chicagobooth.edu

We provide a review of the types of equilibria typically found in operations management inventory papers and a discussion on when the commonly used stationary infinite-horizon (open-loop) equilibrium may be sufficient for study. We focus particularly on order-up-to and basestock equilibria in the context of inventory duopolies. We give conditions under which the stationary infinite-horizon equilibrium is also a Markov perfect (closed-loop) equilibrium. These conditions are then applied to three specific duopolies. The first application is one with stockout-based substitution, where the firms face independent direct demand but some fraction of a firm's lost sales will switch to the other firm. The second application is one where shelf-space display stimulates primary demand and reduces demand for the other firm's product. The final application is one where the state variables represent goodwill rather than inventory. These specific problems have been previously studied in both the single period and/or stationary infinite-horizon (open-loop) settings but not in Markov perfect (closed-loop) settings. Under the Markov perfect setting, a variety of interesting dynamics may occur, including that there may be a so-called commitment value to inventory.

Subject classifications: stochastic games; inventory policies; noncooperative games; Markov equilibrium; inventory games.
Area of review: Operations and Supply Chains.
History: Received April 2011; revisions received January 2013, July 2013, October 2013, November 2013; accepted November 2013.

1. Introduction

It is well known that significantly different dynamics can occur in closed-loop (with feedback) versus open-loop (with no feedback) settings for multiperiod games (see, Başar and Olsder 1999). Yet the predominant type of multiperiod game considered in the operations literature is the stationary infinite-horizon (open-loop) setting. We wish to explore the difference between closed- and open-loop equilibria in the context of dynamic inventory duopolies. Particularly, we explore when so-called order-up-to policies form both open- and closed-loop equilibria.

The appeal of order-up-to policies is well known; they are easy to implement and intuitively satisfying. They simply require a firm to know its current inventory level and its target stocking level and to order the difference. In the single-firm case, order-up-to policies are optimal under quasiconcave payoff functions. Here we study conditions for order-up-to policies to form the same equilibrium in both open-loop and closed-loop versions of a dynamic inventory duopoly, where the players may compete multiple times. In the context of such operational problems, it often makes sense to have these models be state dependent. Such a dependence is frequently associated with an assumption that all payoff-relevant information is embedded in these state variables and anything else in the past is payoff irrelevant; these are Markov games. In this paper, the firms' beginning inventory levels form the periodic state variables of a Markov game. Anticipating the new inventory levels and the competitor's response to these levels is the closure in the feedback loop.

As one example of our conditions, we consider a simple horizontal duopoly competition from the existing literature where prices are fixed and consumers have a preferred product (primary demand). Upon finding that product unavailable, they either will not buy or will switch to the competing product. Such a scenario can either model two competing retailers stocking the same product or one retailer with two competing vendors (e.g., Coca Cola and Pepsi) who manage their stock on the retailer's shelves. In addition to giving conditions under which the stationary infinite-horizon (open-loop) equilibrium is also a Markov perfect (closed-loop) equilibrium, we also highlight examples where this is not the case. In this latter case, the dynamics can be quite complex but there may be a *commitment value to inventory*. We also apply our conditions theoretically, but not numerically, to two further duopolies from the literature, in particular one where inventory stimulates demand and another where the state variable is goodwill rather than inventory.

This paper contributes to two main areas of literature: inventory duopolies and dynamic operational games. The earliest acknowledged inventory oligopoly is attributed to Kirman and Sobel (1974), who elegantly formulate and analyze two competing firms. They were the first to recognize that *intertemporal dependence* in inventory duopolies arises in two manners: via the physical inventory being carried from period to period and via the stochastic demand. They formulate a dynamic oligopoly where the firms set prices; there is no explicit flow of customers based upon stockout substitution. Further relevant literature may be found in §2.

Our primary contribution is a careful treatment of the competitive dynamics for an inventory duopoly. As mentioned above, we establish conditions for when stationary infinite-horizon (open-loop) and Markov (closed-loop) equilibria coincide. We apply our derived conditions to three illustrative models. Although similar models have been studied in the literature, to the best of our knowledge, none have studied the Markov equilibria (ME). Indeed, the literature on ME in operational games is limited and we review this further in §2. We also believe that the treatment of equilibria type has been hampered in the operations literature by a lack of specific vocabulary, and we propose the term a *stationary-strategy equilibrium* (SSE) for the common stationary infinite-horizon open-loop (Nash) equilibrium, used frequently in operations papers. Giving conditions for when this equilibrium is also a Markov equilibrium requires the use of a novel proof technique. We also believe we are the first to provide a formal definition of a basestock equilibrium for an inventory game. In Appendix B, we use numerical examples to highlight where SSE and ME solutions do not coincide. In these examples, we explore some behavior not evident when one limits attention exclusively to the SSE solution technique. It is here where we observe the commitment value to inventory among other strategic behaviors.

This paper is organized as follows. Section 2 reviews the appropriate equilibria types for an inventory duopoly, defines an SSE, further describes relevant literature for dynamic operational games, and outlines the type of behavior that may be seen when the SSE is not a Markov equilibrium. Section 3 gives conditions for existence and uniqueness of an order-up-to SSE and gives conditions under which the SSE is also an ME. Section 4 defines our three example duopolies and applies the conditions of §3 to these settings. Finally, §5 concludes the paper. The proofs of minor results as well as an illustrative numerical study may be found in the appendix, but we retain the proofs of major results in the body of paper because they add intuition to the results. The online appendix (available as supplemental material at http://dx.doi.org/100.1287/opre.2013.1250) contains further minor results.

2. Equilibrium Types and Literature

This section discusses the types of equilibria and policies that are most relevant to inventory competition. In particular, we propose a new equilibrium term, a stationary-strategy equilibrium as well as reviewing the existing concept of a Markov equilibrium. We define order-up-to and basestock policies in the context of inventory duopoly games. We also provide a brief review of papers in the inventory literature using the ME concept. We highlight the key differences between SSE and ME in inventory competition, particularly as they relate to strategic behavior, and explain how there may be a commitment value to inventory seen in ME but not SSE for the same game.

We first describe two policy types relevant for inventory duopolies. We will refer to firms i and $j \neq i$, and it may be assumed that the following definitions and all results in what follows are for $i, j = 1, 2$ and $j \neq i$. Time is given by $t = 1, 2 \ldots$ and counted forward (sometimes up to a finite horizon T); where no confusion is likely, the time subscript will be dropped.

Definition 1 (Order-Up-to Policy). A policy for firm i is a *stationary order-up-to response policy* if, in each period t, for some level y^{i*}, whenever the initial inventory is below y^{i*} firm i orders up to y^{i*} regardless of the other firm's ordering policy for that period. If both firms follow an order-up-to response policy then we call this (y^{i*}, y^{j*}) a *stationary order-up-to policy*.

Notice that an order-up-to policy makes no policy prescription for what should be done if inventory is found to be above (y^{i*}, y^{j*}), and hence papers proving the optimality or equilibrium nature of an order-up-to policy will typically impose a requirement on initial inventory. A natural extension to an order-up-to policy is the basestock policy. In the single-firm setting, a basestock policy is simply an order-up-to policy where the firm orders nothing should it find itself above the basestock level. However, in the competitive setting, additional policy description is needed when one firm is above the order-up-to level and the other firm is not. Let (x_t^i, x_t^j) be the initial inventory levels at the beginning of some period t.

Definition 2 (Basestock Policy). A policy for firm i is a *stationary basestock response policy* if, in each period t, for some level y^{i*}, firm i orders up to $\max(x_t^i, \min(R^i(x_t^j), y^{i*}))$ regardless of the other firm's ordering policy for that period, for some response function $R^i(\cdot)$ with $R^i(x^j) \geqslant y^{i*}$ for $x^j \leqslant y^{j*}$. If both firms follow a stationary basestock response policy then we call this (y^{i*}, y^{j*}) together with $(R^i(\cdot), R^j(\cdot))$ a *stationary basestock policy*.

Notice that this definition relies on the definition of a response function because should firm j be at level $x_t^j > y^{j*}$ and hence order nothing, firm i may have a value $y^i \neq y^{i*}$ that it would prefer to stock to in this case. Figure 1 depicts this situation.

The constraint in Definition 2 that $R^i(x^j) \geqslant y^{i*}$ for $x^j \leqslant y^{j*}$ ensures that if both firms are below (y^{i*}, y^{j*}), then both will order up to (y^{i*}, y^{j*}). A special case of such a basestock policy would be the more traditional definition where each firm always orders up to the same level regardless of the competitor's inventory and orders nothing if above that desired order-up-to level (i.e., $R^i(x^j) = y^{i*}\ \forall\, x^j$). Figure 1 give an example where the response function is not constant but instead decreasing in the opposing firm's inventory level (as might be expected in a substitution game). Point 1 shows initial inventory levels where both firms are below the equilibrium and hence an order-up-to policy may be followed. In points 2 and 3 one firm is above its desired inventory level (and orders nothing) and the opposing firm is below and so orders up to its best response to its opponent's inventory, which is below its desired order-up-to equilibrium level. These points show the *commitment value* of starting the period with inventory above the order-up-to level causing the opposing firm to understock (relative to its order-up-to equilibrium point), which will be beneficial for the first firm if the products are substitutes as implied by the decreasing response curves. In point 4 both firms are above both their desired inventory and their best responses to the opposing firm's initial inventory levels, and therefore, under a basestock policy as defined, both firms order nothing.

In a single period, the appropriate form of equilibrium to study is typically a simple Nash equilibrium (Nash 1950, 1951). A natural extension to the single-period Nash equilibrium is the "open-loop" equilibrium, or as it is more typically referred to in the operations management literature, a Nash equilibrium in the infinite horizon game (see, e.g., Bernstein and Federgruen 2004, for a careful treatment). We find neither of these terms as fully descriptive

Figure 1. Initial inventories are represented by the bold circles, and the arrows indicate the movement to the equilibrium solution.

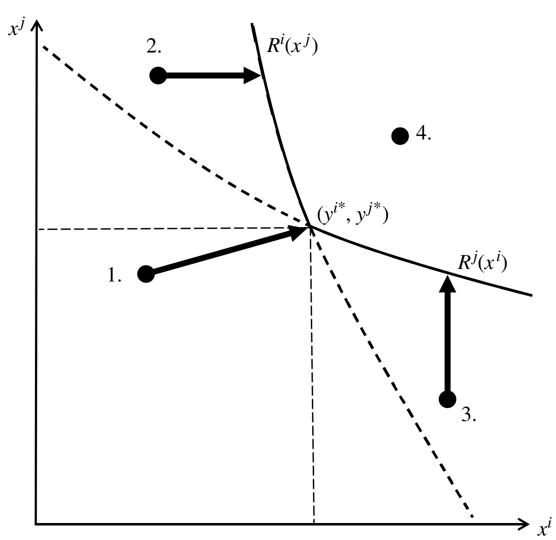

and so propose a new term, namely a stationary-strategy equilibrium, which is more specific than both open-loop and Nash equilibrium and also emphasizes the difference with closed-loop ME, which can also be played on the infinite horizon but do not assume a restriction to stationary policies. Formally, it may be defined as follows.

Definition 3 (SSE). A policy is said to be a *Stationary-Strategy Equilibrium* if, for a possibly restricted set of initial states, the policy is a pure-strategy Nash equilibrium in the infinite horizon game where the same (possibly state-dependent) strategy is played in each period.

If the initial set of states is restricted, then any proof of an SSE must include showing that the future state under the equilibrium policy also lies in the same restricted set. In effect, it assumes that both players are currently setting up their enterprise systems, that strategy decisions are made on a much longer time scale than the operational payoffs, or that firms can't observe their rival's inventory in a regular fashion. It is one of the most commonly studied multiperiod equilibria types in the operations literature having a long history (see Kirman and Sobel 1974, Heyman and Sobel 1984), primarily because of its simplicity in formulation and efficacy in analysis.

A stronger equilibrium solution concept than an SSE is a Markov equilibrium. This is a closed-loop equilibrium and requires subgame perfection in every period.

Definition 4 (ME) (Fudenberg and Tirole 1991, p. 501). A *Markov strategy* depends on the past only through the current period's state variables. A *Markov Equilibrium* or *Markov perfect equilibrium* is a profile of Markov strategies that yields a Nash equilibrium in every proper subgame.

An ME may be defined over a finite or infinite horizon, but unlike the SSE above, in the infinite horizon, subgame perfection is still required. In particular, for an inventory duopoly each player must anticipate the behavior of the other player if inventory is found to be above the equilibrium level. One of the simpler forms of ME found commonly in the operations literature is a linear game. For example, Hall and Porteus (2000) study a linear game where service competitors, or pure newsvendors, compete for market share that may diminish as a result of poor service. In this approach, the firms' payoffs are linear in the state variable (each firm's market share). See also Olsen and Parker (2008) for a review of linear game extensions of this work. A related technique is to assume separability between the firm's decisions in the probabilistic transition function, and this is the approach taken in Albright and Winston (1979), who consider firms competing via pricing and advertising, and Nagarajan and Rajagopalan (2009) who consider newsvendor stocking conditions under which the competitive scenario reduces into a single-product optimization for each firm.

Another approach to proving subgame perfection is to limit the model to two periods, and there are some papers that deal specifically with ME of inventory competition across two periods. For example, Caro and Martínez-de-Albéniz (2010) examine the effect of competition on quick response (see Table 2 in their paper for a crisp summary of the literature on inventory competition with substitutable products). Zeinalzadeh et al. (2010) consider an inventory duopoly where the firms incur a fixed cost for placing orders. They find that in two periods, under a number of conditions, ME do not exist.

The inventory literature on nonlinear nonseparable ME across greater than two periods appears scarce. Lu and Lariviere (2012) numerically compute ME in order to examine a supplier's allocation of scarce inventory to competing retailers; they show a "turn and earn" approach induces greater sales. Numerical computation is the typical approach for finding ME taken in the economics literature, and there is a large literature on efficient methods for such computation (e.g., Pakes and McGuire 1994, Borkovsky et al. 2010). A recent paper by Ifrach and Weintraub (2012) discusses more computationally attractive alternatives. Parker and Kapuściński (2011) demonstrate that the ME for a supplier-retailer relationship subject to capacity constraints is a modified echelon basestock equilibrium policy. They prove the equilibrium policy structure through construction and induction, exploiting Pareto improvement among nonunique equilibria in every period.

Unfortunately, ME found numerically are frequently unattractive from an implementation standpoint (see, e.g., Doraszelski and Satterthwaite 2010, for a discussion). They can be nonstationary where the target equilibria differ from period to period, despite being drawn from a problem with stationary parameters. Worse, there are frequently multiple equilibria with no obvious way to choose the appropriate equilibrium. Note that the conditions in §3 will guarantee a unique SSE, but we will not show ME uniqueness and cannot rule out the existence of other equilibria even under the conditions where the SSE is shown to be an ME. For the numerical example in Appendix B, we have seen up to five equilibria in a given period, although there seems to be no theoretical bound to that number.[1] The immediate significance of finding multiple equilibria is that it ostensibly destroys the potential for finding an equilibrium *policy*.

Unless there is a compelling reason to choose one equilibrium above the others, a model with multiple equilibria cannot deliver a valid equilibrium policy over multiple periods since the ME concept consists of a predictable sequence of actions from a policy, delivering a unique cost and a predictable state in the following period. The literature gives some guidance as to what may constitute a compelling reason for choosing one equilibrium over another, including *Pareto improvement*, where all players encounter a superior outcome for that choice (see Parker and Kapuściński 2011 for an example). Another reason could be a *focal point equilibrium*, where the custom or norms of the problem context suggest one equilibrium has a greater likelihood of being chosen, such as jointly choosing the time of a meeting where the equilibrium (12:00 P.M., 12:00 P.M.) could be preferred over (11:54 A.M.,12:08 P.M.). In the context of our inventory setting, the firms may both be attracted to ordering a full truckload or pallet of goods or ordering to fill a specific space requirement.

One of the most interesting characteristics of (subgame perfect) ME in our context is that there may be a *commitment value to inventory*. We found numerical examples of equilibria where a firm will stock higher than its SSE level even for symmetrical sets of parameters. This is consistent with behavior of the overstocking firm in period t attempting to credibly *commit* to an overstock in the following period. In expectation, firm i believes it can benefit in period $t+1$ if it begins that period with an inventory above the period $t+1$ target level. This may cause firm j to stock less, which will be beneficial for firm i in our substitution-based model of §3.

A review of the literature did not yield any previous articles on the commitment value of inventory in the sense we intend, despite the lengthy literature on inventory control. Saloner (1986) does describe a partial commitment of inventory in the sense that acquiring inventory (either through production or purchasing) and incurring the holding cost does endow the firm with the ability to sell up to that quantity. Our sense is somewhat different, in that a firm may order up to a "high" level in one period with the *hope* that its inventory entering the following period is similarly high. Such a hope is predicated on knowledge of the inventory transition functions. Thus, the beginning inventory level in the following period is far from guaranteed, but if it happens to be high, then the commitment to this inventory is beyond dispute: the concrete physical nature of the durable inventory means it is available for usage.

The commitment value of *capacity* has been studied in the economics literature. The general idea is that the irreversible nature of capacity allows a firm to demonstrate commitment. Tirole (1988, p. 308) describes the period of commitment associated with an investment (such as purchasing capacity) as "a period of time over which the cost of being freed from the commitment within the period is sufficiently high that it does not pay to be freed." The irreversibility of capacity typically arises from the fact that it is lumpy, industry-specific, or otherwise difficult to discard on a secondary market. The sinking of money into an investment of such a magnitude tends to show a commitment, usually towards operating in a particular market. Spence (1977), Dixit (1980), and others show how such a capacity commitment by an incumbent firm can be used to ward off potential entrants into a market, ensuring monopoly operations for the incumbent. The entrants are aware of the incumbent's investment and its durability and irreversibility implies the commitment. Prices do not have a similar deterrent effect because they can be easily changed (Tirole 1988, p. 368). Besanko and Doraszelski (2004) consider dynamic capacity competition but find that capacity differences between firms cannot be sustained because a rival can "catch up" with a leading firm.

3. A General Inventory Duopoly

This section considers equilibrium results for a general inventory game. The results derived in this section are then applied to three specific games in §4. We consider a model with two retailers i and j who face inventory (x_t^i, x_t^j) at the beginning of period t and simultaneously decide their order-up-to amounts (y_t^i, y_t^j). Define the total demand in period t at retailer i, when the retailers stock (y_t^i, y_t^j), as $\Lambda_t^i(y_t^i, y_t^j)$.[2] It is assumed that $\{\Lambda_t^i(y_t^i, y_t^j)\}$ is independent across periods.

The inventory transition function between periods for retailer i is

$$x_{t+1}^i = g^i\big(y_t^i, \Lambda_t^i(y_t^i, y_t^j)\big),$$

where $g^i(\cdot)$ is the level of physical inventory remaining after satisfying demand. Write

$$X^i(y^i, y^j) = g^i\big(y^i, \Lambda^i(y^i, y^j)\big)$$

as the random variable representing the firm i inventory left after demand in a generic period when order-up-to points of (y^i, y^j) are chosen for firms i, j and $X(y^1, y^2) = (X^1(y^1, y^2), X^2(y^2, y^1))$.

We assume that retailer i's infinite-horizon discounted expected payoff may be written as

$$\mathrm{E}\left[\sum_{t=1}^{\infty} \alpha^{t-1} \pi^i(y_t^i, y_t^j)\right]$$

for some stationary payoff function $\pi^i(y^i, y^j)$ where α is the periodic discount factor. We assume $0 \leqslant \alpha < 1$. For the one-period game with payoff function $\pi^i(y^i, y^j)$, let $R^i(y^j)$ be the best response for firm i to firm j setting an inventory level of y^j (similarly for $R^j(y^i)$).

For the sake of establishing existence of the equilibrium, certain technical assumptions about the state and strategy spaces need to be made. A minor assumption is that we treat these spaces as drawn from real space. The state space is $(x_t^1, x_t^2) \in \mathcal{S} = \mathcal{R}_+^2$ and the action space is $(y_t^1, y_t^2) \in \mathcal{A}_t = [x_t^1, \infty) \times [x_t^2, \infty)$ for $t = 0, 1, 2, \ldots$. So we treat the inventory decisions as continuous variables, a common assumption in the literature (e.g., Kirman and Sobel 1974). We also assume that this assumption leads to continuous (but not necessarily differentiable) payoff functions. Further assumptions require restrictions upon the strategy space, commonly requiring compactness. A blunt approach to this is to merely truncate the real number line (for a single decision variable) at a very high positive level (for an upper bound on physical inventory) and at zero as a lower limit (if no backlogs are allowed). This truncation of the strategy spaces is sometimes done merely by assumption or through a justification derived from the demand distribution. A reasonable approach is simply to select an extremely large quantity ("U") to which inventory levels will never realistically be raised. Thus, we augment the action space definition above to $(y_t^1, y_t^2) \in \mathcal{A}_t = [x_t^1, U] \times [x_t^2, U]$.[3]

In what follows we give properties and results that relate these properties to equilibria. These properties will be illustrated further in three specific sample applications in the following section. The first set of properties we consider is for the payoff function.

Property 1. *Payoff $\pi^i(y^i, y^j)$ is*
1. *quasiconcave in y^i*;
2. *submodular in (y^i, y^j); and*
3. *nonincreasing in y^j.*

Note that although Property 1(b) is sufficient for existence of an equilibrium we are interested in the structure of the equilibrium and thus will require Property 1(a) to ensure the equilibrium is order-up-to. Property 1(a) together with continuity and compact strategy spaces will also guarantee existence of an equilibrium. Property 1(c) ensures that firm i is always negatively affected when firm j increases its stocking level.

Property 2. *For each response function $R^i(\cdot)$,*
1. $\left|\partial R^i(z)/\partial z\right| < 1$;
2. $\partial R^i(z)/\partial z \leqslant 0$; *and*
3. $R^i(\cdot) \geqslant 0$.

Properties 2(a)–(c) will be used to ensure a unique equilibrium, submodularity, and nonnegative inventory levels, respectively. The following is a standard result (e.g., Topkis 1998, Lemma 4.2.2).

Lemma 1. *Property* 2(b) *follows directly from Property* 1(b).

Finally, we will assume the following:

Property 3. *Future inventory $X(y^i, y^j)$ has the following properties*:
1. $X^i(y^i, y^j) \leqslant y^i$;
2. *$X(y^i, y^j)$ is stochastically nondecreasing in (y^i, y^j); and*
3. *$X^i(y^i, R^j(y^i))$ is stochastically nondecreasing in y^i.*

Property 3(a) is natural for any reasonable inventory transition function. Property 3(b) ensures that future inventory is nondecreasing in both players' current inventory. Finally, Property 3(c) ensures that the effect of firm j decreasing its inventory in response to an increase in firm i's inventory is less significant for firm i's inventory than the direct effect of that increase, so that overall future inventory for firm i increases when it decides to stock more.

Properties 1–3 will be shown to hold for the three specific applications of §4. However, other more complex inventory duopoly applications that satisfy these properties appear likely. The key restriction from these policies is that the two firms' inventory levels must be substitutes in the sense of Properties 1(b) and (c). Other possible examples may include firms competing for a limited resource such as shelf space or Cournot based price competition.

Theorem 1. *Under Properties* 1(a), 2(a), *and* 3(a), *there exists a unique SSE, which is order-up-to* (*given sufficiently low initial inventories*).

Proof. Existence of an equilibrium follows from Property 1(a) together with continuity and compact strategy spaces (e.g., Fudenberg and Tirole 1991, Theorem 1.2). If the magnitude of the best response derivative is less than 1 (Property 2(a)), then uniqueness of the equilibrium is guaranteed (Vives 1999). The fact that it is order-up-to is given by the quasiconcavity of $\pi^i(\cdot)$ (Property 1(a)). Further, if initial inventory is below the SSE level, then it will also be below it in all following periods, by Property 3(a). Therefore, so long as initial inventories lie below the equilibrium levels, then following the equilibrium order-up-to levels is indeed an SSE. □

Note that although the term basestock is often used in the inventory literature, most proofs follow the above structure and therefore prove that the equilibrium is order-up-to, saying nothing about what should be done when inventory is above the SSE levels. An exception to this is Bernstein and Federgruen (2004), who show that the SSE is basestock under the average reward case. The core of their argument is that there will be a finite number of periods until inventory is below the SSE level, which are negligible in the infinite horizon when $\alpha = 1$. Such an argument does not work for the discounted cost case (because the initial periods are not negligible) and even the correct form for the response function in Definition 2 is not clear because the excess may be maintained for multiple periods and hence the one-period response function will likely not be the correct function to use. Because assuming that initial inventory is low does not seem unreasonable, we leave this as an open but mostly technical problem.

Under Properties 1(a), 2(a), and 3(a), define (y^{i*}, y^{j*}) as the unique SSE, $\pi^{i*} = \pi^i(y^{i*}, y^{j*})$ as the expected one-period profit under the SSE and $V^{i*} = \pi^{i*}/(1-\alpha)$ as the infinite-horizon discounted expected payoff under repeated playing of the SSE.

Let $V_t^i(x_t^i, x_t^j)$ be the profit-to-go function under basestock policy (y^{i*}, y^{j*}) (using the one-period best response function $R^i(\cdot)$ in Definition 2) when the terminal value for the game in period $T+1$ is equal to $V_{T+1}^i(x_{T+1}^i, x_{T+1}^j) \triangleq V^{i*}$. Assuming Properties 2(b) and (c), define $\tilde{y}^i = \lim_{y\to\infty} R^i(y)$, which exists because $R^i(\cdot)$ is nonincreasing and bounded below by zero (by Properties 2(b) and (c)). We will show that, regardless of the other firm's behavior, firm i will always stock at least $\tilde{y}^i$. Note that the less firm j stocks the better for firm i from a one-period payoff perspective because of Property 1(c).

Definition 5 (Maximum Potential Reward for Overstocking). Let $e_{T+1}^i(x^i) = 0$ for all $x^i \geqslant 0$. Then recursively define

$$e_t^i(x) = \begin{cases} 0 & 0 \leqslant x \leqslant y^{i*}, \\ \max_{z\leqslant x}\{\pi^i(z, \tilde{y}^j) - \pi^{i*} \\ \quad + \alpha \mathrm{P}(X^i(z,\infty) > y^{i*}) e_{t+1}^i(\hat{x}_{t+1}^i)\} & x > y^{i*}, \end{cases}$$

and $\hat{x}_t^i = \arg\max_x e_t^i(x)$.

Note that $e_t^i(x)$ is an upper bound on the extra reward available to firm i for an overstock of x at the beginning of period t, assuming firm j reacts in the most favorable fashion possible. The intuition behind the maximum over z in the definition is to ensure that $e_t^i(x)$ is nondecreasing; it is allowable because $e_t^i(\cdot)$ forms a bound. Note that by definition, $e_t^i(y^{i*}) = 0$. We will require the following property.

Property 4. *For $x^i \geqslant y^{i*}$ and $y^j \geqslant y^{j*}$,*

$$\pi^i(x^i, y^j) + \alpha \mathrm{P}\big(X^i(x^i, y^j) > y^{i*}\big) e_t^i(\hat{x}_t^i) \leqslant \pi^{i*}, \qquad \text{for } t = 0, 1, \ldots, T.$$

Property 4 is the key assumption for an SSE to remain an equilibrium in all periods. It states that is not worthwhile for firm i to overstock in this period purely for next period's gain, given that the firm doesn't get an immediate response from firm j. Because this requirement is dependent on t, it requires a loop over t to validate; the following property is stronger but time independent. Note that both Properties 4 and 5 depend only on the model primitives of the one-period payoff function and transition function and not on the equilibria or profit-to-go functions.

Property 5. *Define $\tilde{\pi}^i = \max_x \pi^i(x, \tilde{y}^j)$,*

$$\tilde{z}^i = \arg\max_z \left\{\pi^i(z, \tilde{y}^j) + \alpha \mathrm{P}(X^i(z,\infty) > y^{i*}) \frac{\tilde{\pi}^i - \pi^{i*}}{1-\alpha}\right\},$$

and $\tilde{p}^i = \mathrm{P}(X^i(\tilde{z}^i, \infty) > y^{i*})$. Assume, for $x^i \geqslant y^{i*}$ and $y^j \geqslant y^{j*}$,

$$\pi^i(x^i, y^j) + \alpha \mathrm{P}(X^i(x^i, y^j) > y^{i*}) \frac{\tilde{\pi}^i - \pi^{i*}}{1-\alpha\tilde{p}^i} \leqslant \pi^{i*}.$$

Notice that we can see immediately from Property 5 that if the discount rate is low, if there is a low probability that overstocked inventory remains high in the following period, or if the payoff function is not very responsive to the competitor's inventory level, then this condition, which will be shown to imply that the SSE forms an ME, is more likely to hold. Such conditions are highly intuitive.

Lemma 2. *Under Property* 1(a), *Property* 5 *implies Property* 4.

The restrictiveness of Properties 4 and 5 is tested numerically in the online appendix for the specific stockout-based substitution game of §4.1. Property 5 is found to be very restrictive but is still useful for intuition building. For Property 4 we easily found examples both where it was and was not satisfied. The regions where it held followed the intuition from Property 5 (i.e., low discount rate, etc.).

Before presenting the existence result in this section, we present the redefined Definition 4 in terms of our inventory notation.

Definition 4′ (ME Under Defined Functions). The order-up-to levels (y^{i*}, y^{j*}) form a *Markov equilibrium* if

$$\pi^i(y^{i*}, y^{j*}) + \alpha \mathrm{E}[V^i_{t+1}(X(y^{i*}, y^{j*}))] \geqslant \pi^i(y^i_t, y^{j*}) + \alpha \mathrm{E}[V^i_{t+1}(X(y^i_t, y^{j*}))] \quad (1)$$

and

$$\pi^j(y^{i*}, y^{j*}) + \alpha \mathrm{E}[V^j_{t+1}(X(y^{i*}, y^{j*}))] \geqslant \pi^j(y^{i*}, y^j_t) + \alpha \mathrm{E}[V^j_{t+1}(X(y^{i*}, y^j_t))] \quad (2)$$

for all $(y^i_t, y^j_t) \in \mathscr{A}_t = [x^i_t, U] \times [x^j_t, U]$, $x^i_t \leqslant y^{i*}$, and $x^j_t \leqslant y^{j*}$ for $t \leqslant T$. The basestock policy with levels (y^{i*}, y^{j*}) and response functions $(R^i(\cdot), R^j(\cdot))$ form a *Markov equilibrium* if conditions (1) and (2) hold without the restriction that $x^i_t \leqslant y^{i*}$, and $x^j_t \leqslant y^{j*}$ for $t \leqslant T$, and when $y^{i*}$and y^{j*} are replaced by $\max(x^i_t, \min(R^i(x^j_t), y^{i*}))$ and $\max(x^j_t, \min(R^j(x^i_t), y^{j*}))$, respectively, throughout.

The following theorem shows that under Properties 4 or 5 the SSE is an ME. The proof is inductive. In each period it is shown that if firm j follows the SSE equilibrium, then firm i has no incentive to deviate. No assumptions are made on the behavior if inventory is above the SSE level, but a bound is given for the value in these regions, thus allowing us to prove subgame perfection, i.e., that there is no incentive to deviate from the equilibrium in order to move off the equilibrium path in the following period.

Theorem 2. *Under Properties* 1–3 *and either Property* 4 *or* 5, *if initial inventory starts below* (y^{i*}, y^{j*}) *then ordering up to* (y^{i*}, y^{j*}) *in each period is a Markov equilibrium. Further, regardless of firm* j*'s initial inventory, firm* i *will always order at least* $\tilde{y}^i$.

Proof. Notice that as Property 5 implies Property 4 we need only consider the latter. Further note that Property 4 also holds for $x^i < y^{i*}$ by the following reasoning. For $x^i < y^{i*}$, $\mathrm{P}(X^i(x^i, y^j) > y^{i*}) = 0$ and hence

$$\pi^i(x^i, y^j) + \alpha \mathrm{P}(X^i(x^i, y^j) > y^{i*}) e^i_t(\hat{x}^i_t) = \pi^i(x^i, y^j) \leqslant \pi^i(x^i, y^{j*}) \leqslant \pi^{i*}$$

for $t = 0, 1, \ldots, T$. The first inequality follows from Property 1(c) and $y^j \geqslant y^{j*}$ and the second inequality follows from the definition of π^{i*}.

We proceed by induction where at each step of the induction we will prove the following:

1. $V^i_t(x^i, x^j) = V^{i*}$ for $x^i \leqslant y^{i*}$ and $x^j \leqslant y^{j*}$ and the order-up-to policy is Markov from period t onward when initial inventories are below the order-up-to levels.
2. $V^i_t(x^i, x^j) \leqslant V^{i*} + e^i_t(\hat{x}^i_t)$ for $x^i > y^{i*}$.
3. Firm i will always order at least $\tilde{y}^i$ in period t, and $V^i_t(x^i, x^j)$ is constant in x^i for $x^i < \tilde{y}^i$.

The basis is trivially established in period $T+1$ and now assumed true in period $t+1$ onward.

Part (a). Assume $0 \leqslant x^i \leqslant y^{i*}$ and $0 \leqslant x^j \leqslant y^{j*}$. We need to show that if firm j follows a y^{j*} order-up-to policy, then firm i's optimal response is to order up to y^{i*} and that $V^i_t(x^i, x^j) = V^{i*}$, which clearly yields (a) for time t. Now

$$\begin{aligned} V^i_t(x^i, x^j) &= \max_{y \geqslant x^i} \{\pi^i(y, y^{j*}) + \alpha \mathrm{E}[V^i_{t+1}(X(y, y^{j*}))]\} \\ &\leqslant \max_{y \geqslant x^i} \{\pi^i(y, y^{j*}) + \alpha [\mathrm{P}(X^i(y, y^{j*}) \leqslant y^{i*}) V^{i*} \\ &\qquad + \mathrm{P}(X^i(y, y^{j*}) > y^{i*})(V^{i*} + e^i_{t+1}(\hat{x}^i_{t+1}))]\} \\ &= \max_{y \geqslant x^i} \{\pi^i(y, y^{j*}) + \alpha (V^{i*} + \mathrm{P}(X^i(y, y^{j*}) > y^{i*}) \\ &\qquad \cdot e^i_{t+1}(\hat{x}^i_{t+1}))\} \\ &\leqslant \pi^{i*} + \alpha V^{i*} = V^{i*}, \end{aligned}$$

where the first inequality follows by the induction Hypotheses (a) and (b) and the second inequality follows from Property 4. But by definition, $\pi^i(y^{i*}, y^{j*}) + \alpha \mathrm{E}[V^i_{t+1}(X(y^{i*}, y^{j*}))] = V^{i*}$ so that ordering up to y^{i*} must be a best response for firm i to firm j ordering y^{j*} (which is feasible by the assumption on the initial inventories).

Part (b). Suppose $x^i > y^{i*}$. We need to show that so long as firm j orders at least $\tilde{y}^j$, then $V^i_t(x^i, x^j) \leqslant V^{i*} + e^i_t(x^i)$. Note that we make no claim as to firm i's response; extra conditions will be needed to show that firm i's best response is to order nothing (see Theorem EC.1 in the online appendix). Now, if firm j orders $y^j \geqslant \tilde{y}^j$,

$$\begin{aligned} V^i_t(x^i, x^j) &\leqslant \max_{y \geqslant x^i} \{\pi^i(y, \tilde{y}^j) + \alpha (V^{i*} + \mathrm{P}(X^i(y, \infty) > y^{i*}) \\ &\qquad \cdot e^i_{t+1}(\hat{x}^i_{t+1}))\} \\ &\leqslant \pi^{i*} + \alpha V^{i*} + \max_{y \geqslant x^i} \{e^i_t(y)\} \\ &\leqslant V^{i*} + e^i_t(\hat{x}^i_t). \end{aligned}$$

The first inequality follows by the same reasoning as in Part (a), the second by the definition of $e^i_t(\cdot)$, and the third by the definition of $\hat{x}^i_t$.

Part (c). Suppose that $x^i < \tilde{y}^i$. We need to show that firm i orders at least $\tilde{y}^i$ and that $V^i_t(\cdot)$ is constant over this range. But for $y < \tilde{y}^i$, for any y^j, $X^i(y, y^j) < \tilde{y}^i$ and $\mathrm{E}[V^i_{t+1}(X(y, y^j))]$ is constant in y by the induction hypothesis. Therefore,

$$\begin{aligned} V^i_t(x^i, x^j) &= \max_{y \geqslant x^i} \{\pi^i(y, y^j) + \alpha \mathrm{E}[V^i_{t+1}(X(y, y^j))]\} \\ &= \max_{y \geqslant \tilde{y}^i} \{\pi^i(y, y^j) + \alpha \mathrm{E}[V^i_{t+1}(X(y, y^j))]\}, \end{aligned}$$

where the second equality follows since $\pi^i(y, y^j)$ is nondecreasing in y for $y < \tilde{y}^i$ (by quasiconcavity). Thus, firm i orders at least $\tilde{y}^i$ and $V^i_t(x^i, x^j)$ is constant in x^i for $x^i < \tilde{y}^i$. Q.E.D

The proof of Theorem 2 is inductive in nature, which is natural for proving subgame perfection. Although a finite horizon is assumed, a similar technique could also be used in the infinite horizon, either directly or by letting $T \to \infty$ (e.g., Fudenberg and Levine 1983). The key to the inductive arguments is the definition of the $e_t^i(\cdot)$ functions, which are single variable functions in a firm's own stocking level and therefore independent of the competing firm's decisions. Bounding the payoff functions in terms of these functions allows subgame perfection to be proven.

A stronger result than the SSE order-up-to levels forming an ME is that they are actually basestock. A theorem showing that a basestock policy forms an ME, under some given conditions, is given in the online appendix. The response function used in the basestock definition (see Definition 2) is the single-period response function R^i, which is relatively intuitive. However, the extra condition required to prove this theorem is quite strong. This is because we must show not only that there is no incentive to deviate from the SSE if inventories lie below the SSE (as was guaranteed by Property 4) but also that should inventory lie above the SSE then no further order should be placed. Because the required condition is relatively unattractive, we have placed this result in the online appendix.

4. Examples

In this section we apply the results of the previous section to three specific models. The first is one of stockout-based inventory substitution, the second is a game where inventory display size influences demand, and the last is one where the state variable is not inventory but market size. We have kept our models simple both for ease of exposition and because they are designed to illustrate the theoretical results, rather than comprehensively apply them, which is left as the subject of future research. The proofs of the lemmas in this section are relatively simple and hence relegated to the online appendix.

4.1. A Stockout-Based Substitution Game

It is well documented that consumers may switch brands upon experiencing a stockout (e.g., Schary and Christopher 1979). Given this behavior, one dimension of retailer competition is the availability of stock. Of course, retail prices also strongly influence customers purchasing decisions, but once price competition has settled prices to a common level between retailers, customers next look to whether the product is on the shelf. This subsection applies the results of the previous section to a simple stockout-based substitution game from the literature.

We consider two competing firms who face both primary and secondary customer demand. Primary (or "direct") customers arrive to a firm each period and will each buy one unit of the product if it is found to be available. If the product is unavailable, then some proportion of the unsatisfied customer demand will switch products and become secondary (or "transferred") demand at the opposite firm and the rest will be lost sales.

The earliest work on stockout-based substitution appears to be the single-period model of Parlar (1988) who shows that there exists a unique Nash equilibrium and, by implication, that an order-up-to policy is Nash. Lippman and McCardle (1997) similarly show existence of a Nash solution (via supermodularity) and so a joint inventory stocking level (order-up-to policy) is Nash. Avsar and Baykal-Gürsoy (2002) extend Parlar's (1988) model to an infinite-horizon setting and show that an order-up-to policy forms an SSE. Netessine et al. (2006) also show an order-up-to policy forms an SSE, but under more general demand routings and an allowance for backlogging. Both Avsar and Baykal-Gürsoy (2002) and Netessine et al. (2006) assume initial inventories to be below the order-up-to levels in their key theorems.

The periodic reward consists of revenue minus costs. The costs are (a) holding costs for physical inventory and (b) production costs for the amount "processed" in each period; these will be described more precisely below. Also, all payoff relevant information is assumed to be embedded in the inventory levels of the retailers. This assumption, or something similar, is necessary in a Markov game formulation. We also assume that all information on costs, revenues, and demand distributions is common knowledge.

The notation for retailer i in period t is the following:

r^i = revenue per unit (retail price);
h^i = holding cost per unit per period;
c^i = production cost per unit (wholesale price);
y_t^i = inventory order-up-to level in period t (decision variable);
x_t^i = inventory level at the beginning of period t (state variable);
D_t^i = direct demand to retailer i in period t (random variable); and
γ^{ij} = the proportion of retailer i's direct customers who will transfer to retailer j.

We define the total demand in period t at retailer i, when retailer j stocks y_t^j, as

$$\Lambda_t^i(y_t^j) = D_t^i + \gamma^{ji}(D_t^j - y_t^j)^+,$$

where $x^+ = \max(x, 0)$. We assume the $\{D_t^i\}$ and $\{D_t^j\}$ form mutually independent i.i.d. sequences.

We assume $\gamma^{ij} \in [0, 1)$.[4] Firms accrue revenues only when they satisfy the customers' orders; recall that no backlogging is allowed. We assume $r^i \geqslant c^i$, which is natural. All revenue and cost parameters are nonnegative.

The inventory transition function between periods for retailer i is

$$x_{t+1}^i = (y_t^i - \Lambda_t^i(y_t^j))^+,$$

which is the level of physical inventory remaining after satisfying direct and transfer demand. So in the notation of the previous section, $g^i(y, \Lambda) = (y - \Lambda)^+$.

Retailer i's infinite-horizon expected payoff is

$$\mathrm{E}\left[\sum_{t=1}^{\infty}\alpha^{t-1}(r^i\min(y_t^i,\Lambda_t^i)-c^i(y_t^i-x_t^i)-h^i(y_t^i-\Lambda_t^i)^+)\right].$$

Using the standard translation that $x_t^i=(y_{t-1}^i-\Lambda_{t-1}^i)^+$ (e.g., Veinott 1965) and $r^i\min(y_t^i,\Lambda_t^i)=r^iy_t^i-r^i(y_t^i-\Lambda_t^i)^+$, we rewrite the infinite-horizon expected payoff as

$$\mathrm{E}\left[\sum_{t=1}^{\infty}\alpha^{t-1}((r^i-c^i)y_t^i-(r^i+h^i-\alpha c^i)(y_t^i-\Lambda_t^i)^+)\right]+c^ix_1^i.$$

We therefore define firm i's one-period expected payoff as

$$\begin{aligned}\pi^i(y^i,y^j)&=(r^i-c^i)y^i-\mathrm{E}\big[(r^i+h^i-\alpha c^i)(y^i-\Lambda^i(y^j))^+\big]\\&=(r^i-c^i)y^i-\mathrm{E}\big[(r^i+h^i-\alpha c^i)\\&\qquad\cdot(y^i-D^i-\gamma^{ji}(D^j-y^j)^+)^+\big].\end{aligned}$$

As previously, for the one-period game with payoff function $\pi^i(y^i,y^j)$, let $R^i(y^j)$ be the optimal response for firm i to firm j setting an inventory level of y^j (similarly for $R^j(y^i)$). We begin by showing that π, R, and X satisfy Properties 1 through 3 from §3.

Lemma 3. *Payoff* $\pi^i(y^i,y^j)$ *satisfies Property* 1.

Lemma 4. *The response function* $R^i(\cdot)$ *satisfies Property* 2.

Lemma 5. *Inventory* $X^i(y^i,R^j(y^i))$ *satisfies Property* 3.

These properties imply (also as in Netessine et al. 2006) by Theorem 1 that there exists a unique SSE, (y^{i*},y^{j*}), which is order-up-to (given sufficiently low initial inventories). Thus, Theorem 2 of §3 may be carried over to yield the following result.

Corollary 1. *Under Properties* 4 *or* 5, *if initial inventory starts below* (y^{i*},y^{j*}) *then ordering up to* (y^{i*},y^{j*}) *in each period is a Markov equilibrium.*

We have exercised this model numerically in Appendix B. We discuss the likelihood of Properties 4 or 5 holding and find that the former holds in a wide variety of cases whereas the latter is quite restrictive. We study strategic behaviors under the ME (when it is not an SSE). A variety of interesting behaviors are observed including the commitment value of inventory, as was discussed in §2. Extending this numerical study more generally is left as the subject of future research.

4.2. Inventory-Level-Dependent Demand Inventory Model

In this subsection we present an inventory competition model where the arriving demand may be stimulated or dampened by the presence of the primary product inventory or the other product inventory, respectively. Urban (2005) comprehensively reviews these inventory models. Although there are multi-item models in the literature, there do not seem to be competitive versions. Similar to Roy and Maiti (1998) and others (see Urban 2005), we propose a linear dependency on inventory with an additive stochastic element, namely (for $i\neq j\in\{1,2\}$),

$$\Lambda^i(y_t^i,y_t^j)=D^i+\beta_{ii}y^i-\beta_{ij}y^j$$

where D^i is a random base demand and $\beta_{ii}\geqslant 0$ and $\beta_{ij}\geqslant 0$. We will assume $\beta_{ii}+\beta_{ij}<1$ and that the demand for product i, Λ^i, is nonnegative under appropriate regulatory limitations on D^i, β_{ii}, and β_{ij}. The essence of the demand model is that the presence of more product i and less product j stimulates demand for product i. We define the inventory transition function as

$$x_{t+1}^i=(y_t^i-\Lambda_t^i(y_t^i,y_t^j))^+,$$

which is the level of physical inventory remaining after satisfying total direct demand. So in the notation of §3, $g^i(y,\Lambda)=(y-\Lambda)^+$. We assume r^i, h^i, and c^i as given previously.

Retailer i's infinite-horizon expected payoff is

$$\mathrm{E}\left[\sum_{t=1}^{\infty}\alpha^{t-1}\big(r^i\min(y_t^i,\Lambda_t^i)-c^i(y_t^i-x_t^i)-h^i(y_t^i-\Lambda_t^i)^+\big)\right]$$

and using identical logic as before, we define firm i's one-period expected payoff as

$$\begin{aligned}\pi^i(y^i,y^j)&=(r^i-c^i)y^i-\mathrm{E}\big[(r^i+h^i-\alpha c^i)(y^i-\Lambda^i(y^i,y^j))^+\big]\\&=(r^i-c^i)y^i-\mathrm{E}\big[(r^i+h^i-\alpha c^i)\\&\qquad\cdot((1-\beta_{ii})y^i-D^i+\beta_{ij}y^j)^+\big].\end{aligned}$$

The following results follow naturally from the above assumptions.

Lemma 6. *Payoff* $\pi^i(y^i,y^j)$ *satisfies Property* 1.

Lemma 7. *The response function* $R^i(\cdot)$ *satisfies Property* 2.

Lemma 8. *Inventory* $X^i(y^i,R^j(y^i))$ *satisfies Property* 3.

These properties imply by Theorem 1 that there exists a unique SSE, (y^{i*},y^{j*}), which is order-up-to (given sufficiently low initial inventories). Thus, Theorem 2 of §3 may be carried over to yield the following result.

Corollary 2. *Under Properties* 4 *or* 5, *if initial inventory starts below* (y^{i*},y^{j*}), *then ordering up to* (y^{i*},y^{j*}) *in each period is a Markov equilibrium.*

4.3. A Competitive Advertising Model

In this subsection we consider a model where the underlying state is not inventory per se, but some sort of underlying customer goodwill or market share. There are a number of duopoly models in the literature with this sort of structure (e.g., Albright and Winston 1979, Hall and Porteus 2000), but here we follow the model of Heyman and Sobel (1984, §9.5). In each period the firm decides its advertising expenditures and the subsequent effects on goodwill are cumulative but diminishing with time at rate $\theta < 1$ per period. Thus, current goodwill $y_t^i = x_t^i + z_t^i$ where z_t^i is the advertising spent in this period, and $x_{t+1}^i = \theta y_t^i$. Demand depends on both firms' current goodwill, and expected demand is denoted $\mu^i(y^i, y^j)$. Expected profit in period t is given by $r^i \mu^i(y_t^i, y_t^j) - z_t^i$ and after substitution of $y_t^i = x_t^i + z_t^i$ and rearranging,

$$\pi^i(y^i, y^j) = r^i \mu^i(y_t^i, y_t^j) - (1-\alpha\theta) y_t^i.$$

Submodularity of this game (which was not the focus of Heyman and Sobel) depends on the form of $\mu^i(\cdot)$ and here for ease of exposition we will take a simple Cobb-Douglas form for substitutes (e.g., Huang et al. 2013) as follows. Assume

$$\mu^i(y^i, y^j) = (y^i)^{\gamma} (y^j)^{-\beta},$$

where $0 < \gamma < 1$, $\beta > 0$, and the parameters γ and β are symmetric across firms for simplicity. Note that by just specifying the expected demand, we leave the precise form for demand uncertainty unstated, although clearly a number of variants are possible. We make the following two assumptions on the model. First, that β is sufficiently small so that

$$\frac{\beta}{1-\gamma}\left(\frac{r\gamma}{1-\alpha\theta}\right)^{1/(1-\gamma)} < 1.$$

Second, we assume that the action space is $[1, M]$ with M a large number. That is, the firms must keep at least one unit of goodwill at all times and goodwill is bounded above by some large value. These assumptions result in the following properties.

LEMMA 9. *Payoff* $\pi^i(y^i, y^j)$ *satisfies Property* 1.

LEMMA 10. *The response function* $R^i(\cdot)$ *satisfies Property* 2.

LEMMA 11. *Inventory* $X^i(y^i, R^j(y^i))$ *satisfies Property* 3.

These properties imply by Theorem 1 that there exists a unique SSE, (y^{i*}, y^{j*}), which is order-up-to (given sufficiently low initial goodwill). Thus, Theorem 2 of §3 may be carried over to yield the following result.

COROLLARY 3. *Under Properties* 4 *or* 5, *if initial goodwill starts below* (y^{i*}, y^{j*}), *then advertising up to* (y^{i*}, y^{j*}) *in each period is a Markov equilibrium.*

5. Conclusions and Future Research

The extant literature on inventory games does not usually make explicit that the type of equilibrium being studied is open-loop and therefore not guaranteed to be subgame perfect. We suggest the term a stationary-strategy equilibrium for the standard equilibrium, where the restriction is to stationary policies across the infinite horizon, in order to make this distinction clearer. We give conditions for when the SSE is also an ME (i.e., is also a closed-loop equilibrium). The proof technique used to show subgame perfection relies on an inductive argument showing that any deviation from the SSE in the current period does not provide sufficient reward in the following period to warrant consideration. Although relatively natural, we are not aware of other papers using this technique and believe it may be a general technique to consider for proving that an SSE is also an ME. Our conditions are sufficient but are not proven to be necessary, so future research could try to find necessary conditions.

We also show that behavior under an ME may be far from trivial, and future research should study ways to allow policy insights from the equilibrium. Until that is achieved, we feel that the SSE concept has significant value and should continue to be studied from an operations perspective. After all, there are many environments where the macro decisions on operations (e.g., the order-up-to levels for an enterprise system) are made on a longer timescale then the day-to-day fluctuations of, say, inventory. Also, the macro-level insights obtained from such games (e.g., the losses to a firm due to competition) will likely also frequently carry over to the closed-loop ME setting.

We performed a limited numerical study for a stockout-based substitution inventory model. We found no examples where both firms overstock relative to the SSE (which we had initially hypothesized might be observed) nor any examples where both firms understock (which is not surprising). We did find numerical examples where one firm overstocked and the opposing firm understocked by even further than their one-period best response to this overstock. A possible interpretation of this behavior is that the firm that is understocking relative to its SSE is attempting to *sabotage* the firm that is overstocking. The understocking firm knows that it will be at a disadvantage in the following period if the overstocking firm continues to have inventory superiority in the following period, so it understocks in the current period knowing that some of its current-period unsatisfied customers will transfer to the rival firm and draw down the rival's inventory. Thus, it can indirectly attempt to reduce the high stocking level of the other firm by deliberately channeling unmet demand toward it. Future research should study such competitive behaviors further for more comprehensive applications than our deliberately simple examples; this may need to be done numerically.

In summary, our primary contribution is a deeper understanding of the general relationship between SSE and ME in inventory duopolies and particularly for models with substitution effects. There is space for further methodological results on the analysis of ME.

Appendix A. Proof of Lemma 2

$$e^i_{t+1}(x) \leqslant \tilde{\pi}^i - \pi^{i*} + \alpha e^i_{t+2}(\hat{x}^i_{t+2}) \leqslant \sum_{s=0}^{T-t-1} \alpha^s (\tilde{\pi}^i - \pi^{i*}) \leqslant \frac{\tilde{\pi}^i - \pi^{i*}}{1-\alpha},$$

so that by the quasiconcavity of $\pi^i(\cdot)$ and the nondecreasing nature of cumulative probabilities, $\hat{x}^i_t \leqslant \tilde{z}^i$. Therefore,

$$e^i_t(x) \leqslant \tilde{\pi}^i - \pi^{i*} + \alpha \tilde{p}^i e^i_{t+1}(\hat{x}^i_{t+1}) \leqslant \sum_{s=0}^{T-t} (\alpha \tilde{p}^i)^s (\tilde{\pi}^i - \pi^{i*}) \leqslant \frac{\tilde{\pi}^i - \pi^{i*}}{1-\alpha \tilde{p}^i}. \quad \text{Q.E.D}$$

Appendix B. Numerical Examples of Markov Equilibria

In this appendix, we examine numerical examples of Markov equilibria, particularly where the equilibria do not coincide with the SSE. This highlights the fact that when the conditions of §3 do not hold, there can be interesting nonstationary, dynamic behaviors exhibited by Markov games. The numerical examples we explore are exercising the stockout-based substitution model from §4.1.

For the sake of exercising the model numerically, we consider a discrete state space with discrete demand distributions.[5] For the end of the horizon (i.e., period $T+1$), we determine the SSE equilibrium and then construct a policy that is iterated numerous times to establish a value function perpetuity of playing the SSE as a basestock policy with the one-period response function forming the response when one player is above the SSE. Thus, any deviations from the SSE are not due to the end-of-horizon effects but instead are true dynamic deviations due to the strategic value of inventory. Then for each element of the state space, we find the best response functions for each firm, constrained by the firms' initial inventories. Enumerating over these curves, we find all intersecting points that are, by definition, fixed points or equilibria. The operand of the value function is then mapped to the value function for each equilibrium. Then the inventory transition function is inserted into that period's value function and an expectation calculated, akin to a dynamic programming recursion; the process repeats for preceding period.

We deliberately searched for parameters where the ME examples would not coincide with the SSE and possibly exhibit interesting behavior. One behavior we hypothesized was that there exists a *commitment value of inventory*. Specifically, our motivation is to explore the possibility that a firm may choose to stock higher than the SSE levels in order to gain in the following period, hopeful it will enter that following period at a high stocking level and thus forcing its rival to stock at its best response (below the SSE level). The demand distributions we use in the examples will have a point mass at zero specifically to encourage this behavior. With this point mass the following examples did not satisfy Property 4 in §3. However, without the point mass (and using simple uniform demand) Property 4 is satisfied for both examples. These are not the only examples we examined, but we consider them representative of the properties we wish to highlight. The target equilibrium stocking levels for the two examples we consider can be found in Tables B.1 and B.2.[6] These target levels are desirable stocking levels but may not be feasible in every period because each firm has an initial inventory in each period that may make these target levels infeasible. The SSE target levels may be found in period $T+1$ in each table.

We are now ready to discuss some specific behaviors we observe in these examples:

• *Lack of coincidence with SSE*: In Table B.1 we see that the SSE, (29, 20), is not the equilibrium in every period. The point (29, 20) is on firm 1's best response function in period T but not on firm 2's best response function. Defining $J^2_T(y^1, y^2) = \pi^2(y^1, y^2) + \alpha \mathrm{E}[V^2_{T+1}(X(y^1, y^2))]$, we find $J^2_T(29, 19) = 4{,}897.37$, $J^2_T(29, 20) = 4{,}906.13$, $J^2_T(29, 21) = 4{,}919.65$, demonstrating that firm 2 has a unilateral incentive to deviate from $y^2 = 20$ in period T when firm 1 chooses $y^1 = 29$. This consolidates the message that the SSE solution is not always an ME.

• *Nonstationary*: We see in both tables that the target equilibria differ from period to period, despite being drawn from a problem with stationary parameters.

• *Nonuniqueness*: Table B.2's example illustrates the existence of multiple equilibria.[7] The conditions in §3 guarantee a unique SSE, but we have not shown ME uniqueness and cannot rule out the existence of others. Numerically, we have seen up to five equilibria in a given period, although there seems to be no theoretical bound to that number. As shown in Table B.2, in period T there are three equilibria, (14, 17), (17, 14), and (15, 15), and each firm prefers a different equilibrium. Without further information or problem context, it is impossible to validly select one of these equilibria as a focal point. However, although (15, 15) is not the preferred equilibrium of either firm, the SSE levels (15, 15) *might* arguably be an attractive candidate as a mutual choice, particularly when it is consistent as it is here.

Table B.1. Example for a single equilibrium.

Period	Equilibrium
$T+1$	(29, 20)
T	(28, 21)
$T-1$	(28, 20)
$T-2$	(29, 20)
$T-3$	(28, 21)
$T-4$	(28, 20)

Table B.2. Example for multiple equilibria.

Period	Choosing highest	Choosing lowest
$T+1$	(15, 15)	(15, 15)
T	(14, 17) (15, 15) (17, 14)	(14, 17) (15, 15) (17, 14)
$T-1$	(13, 19) (14, 17) (15, 16)	(16, 15) (17, 14) (19, 13)
$T-2$	(13, 19) (14, 17) (16, 15)	(12, 21) (13, 20) (15, 16)
	(20, 13) (21, 12)	(17, 14) (19, 13)
$T-3$	(13, 19) (14, 17)	(17, 14) (19, 13)
$T-4$	(13, 19) (18, 14)	(14, 18) (19, 13)
$T-5$	(13, 19) (14, 17)	(17, 14) (19, 13)
$T-6$	(13, 19) (19, 13)	(13, 19) (19, 13)
$T-7$	(13, 19) (14, 17)	(17, 14) (19, 13)
$T-8$	(13, 19) (17, 14)	(12, 21) (13, 20)
	(20, 13) (21, 12)	(14, 17) (19, 13)
$T-9$	(13, 19) (14, 17)	(17, 14) (19, 13)

- *Cycling*: There is a *cycling* behavior observed in Table B.1, where the equilibria are repeated in subsequent periods, in this case with a periodicity of three periods. This regular cycle continues as the horizon length grows. Notice these are equilibrium solutions of the order-up-to variables of the Markov model, which automatically implies this is an order-up-to equilibrium policy, albeit not a stationary or monotone one. It is also interesting to note that the equilibria in the Table B.1 always have one player on the best response curves generated by the SSE; for example, firm 1 stocking 28 is a best response to firm 2 stocking 21, yet 21 is not a best response for firm 2 to firm 1 stocking 28 but 20 is. In Table B.2's example, we again observe a *cycling* behavior. This is where an equilibrium will reappear on a regular basis. For example, in Table B.2 we see that (14, 17) appears under "Choosing highest" and (17, 14) appears under "Choosing lowest" in period $T-1$, $T-3$, $T-5$, $T-7$, etc.
- *Commitment to inventory*: As noted, there are equilibria where a firm will stock higher than its SSE level even for symmetrical sets of parameters. This is consistent with the above-conjectured behavior on the overstocking firm attempting to credibly *commit* to an overstock in the following period. For example, we find (14, 17), (17, 14), and (15, 15) in a symmetrical example where the SSE is (15, 15).[8] Thus, for firm i in period T ("Choosing highest"), it finds that it can benefit by committing to an inventory level of 17 and its rival, firm j, will respond by stocking at 14. In expectation, firm i believes it can benefit in period $T+1$ if it begins that period with an inventory above the period $T+1$ target level of 15 and capture the bulk of the primary and transfer demand. Another observation is that whenever the target levels do not correspond to the SSE levels, one firm will stock *above* its SSE level while the other will stock *below* its SSE level. This property is observed both in symmetrical and asymmetrical examples. We found no examples where both firms overstock (which we had initially hypothesized might be observed) nor any examples where both firms understock (which is not surprising).
- *Sabotage*: An alternative perspective of this "overstock-understock" behavior is that the firm that is understocking relative to its SSE is attempting to *sabotage* the firm that is overstocking. The understocking firm knows it will be at a disadvantage in the following period if the overstocking firm continues to have an inventory superiority in the following period, so it understocks in the current period knowing that some of its current-period unsatisfied customers will transfer to the rival firm and draw down the rival's inventory. Thus, it can indirectly attempt to reduce the high stocking level of the other firm by deliberately channeling unmet demand towards it. Consider an equilibrium in Table B.2 of (21, 12) in Period $T-2$ (when equilibria of (15, 16) and (17, 14) are chosen in Periods $T-1$ and T, respectively). With $R^1(12)=17$, which is lower than 21, we can see that firm 1 is stocking higher than the SSE best response level. Further, as $R^2(21)=13$, firm 2 is *understocked* relative to its SSE best response to firm 1 stocking 21. The effect of this understocking is that more of its customers will stockout and transfer to the overstocking firm, thus reducing the chance that the overstocking firm will enter the following period at a high stocking level.
- *See-saw*: One final property observed in Table B.2 is what we label as the *see-saw* effect. This is where the highest stocking level for firm 1 under "Choosing highest" appears above and below the SSE levels in alternating periods. Similarly, firm 1's lowest stocking level see-saws above and below firm 1's SSE level under "Choosing lowest." Firm 2's stocking level similarly follows such an alternating see-saw effect but will be above its SSE level when firm 1's stocking level is below its respective SSE level. Of course, these levels will not always be feasible when demand is low and the actual stocking level will then vary from these desired equilibria.

Finally, we note that these behaviors tend to crop up in the numerical examples in varying propensities but are easily found, particularly when demand has a high probability of being small. Moreover, these behaviors could not be found by limiting attention to the SSE solution concept alone but requires the truly dynamic perspective of the ME.

Supplemental Material

Supplemental material to this paper is available at http://dx.doi.org/10.1287/opre.2013.1250.

Endnotes

1. Our approach of enumerating over each firm's best response curve guarantees we find all equilibria. Borkovsky et al. (2010) provides a guide for finding multiple equilibria using the homotopy method, which is utilized in practice by Besanko et al. (2010).
2. The arguments of $\Lambda_t^i(\cdot)$ will sometimes be dropped when they are too notationally cumbersome and the meaning is clear from the context. Further, in §4.1 the first argument will be dropped because primary demand is independent of the firm's own stocking level.
3. For ease of exposition we follow the common assumption that inventory cannot be destroyed, but extensions to models where there is a disposal cost (and hence $y_t^i < x_t^i$ is feasible) seem likely to be possible.
4. The assumption that $\gamma^{ij} < 1$ will be used in Lemma 4 to show that the magnitude of the derivative of the response function is strictly less than 1. An equivalent more technical condition on the support of the density could also be used if $\gamma^{ij} = 1$ (i.e., if all customers transfer upon a stockout).
5. Note that this assumption does not match the assumption for continuous demand in §3. However, the purpose of this section is to show examples of interesting dynamics that may be seen under an ME that are not seen under the SSE, and thus we believe the assumption is appropriate because it allows complete enumeration of the equilibria.
6. The numbers in Table B.1 are created for the following data: $r^1=21$, $r^2=18$, $h^1=1$, $h^2=2$, $c^1=14$, $c^2=10$. The numbers in Table B.2 are created for the following data: $r^1=r^2=20$, $h^1=h^2=5$, $c^1=c^2=10$. Both examples have these common values: $\gamma^{12}=\gamma^{21}=1$, $\Pr(D^1=0)=\Pr(D^2=0)=0.6$, $\Pr(D^1=i)=\Pr(D^2=i)=0.01$, $i=1,\ldots,40$, $\alpha=0.99$.
7. The columns in the table illustrate the equilibria in each period when firm 1 chooses the equilibrium (in each period) with the highest and lowest stocking level. For example, moving back from the end of the horizon, in Period T when "Choosing highest" the (17, 14) equilibrium is chosen and when "Choosing lowest" the (14, 17) equilibrium is chosen.
8. Clearly, in a symmetrical example that has unequal target stocking levels for one equilibrium, there will be a corresponding equilibrium with the target levels reversed.

Acknowledgments

The authors appreciate the comments of Jeroen Swinkels; seminar participants at Northwestern University, Purdue University, Southern Methodist University, and the University of British Columbia and conference attendees at INFORMS, MSOM, and IFORS. Financial support from the University of Chicago Booth School of Business is gratefully acknowledged. The authors appreciate the support of Area Editor Jeannette Song, the associate editor, and two anonymous referees, whose comments have considerably improved the paper.

References

Albright SC, Winston W (1979) Markov models of advertising and pricing decisions. *Oper. Res.* 27(4):668–681.

Avsar ZM, Baykal-Gürsoy M (2002) Inventory control under substitute demand: A stochastic game application. *Naval Res. Logist.* 49(4):359–375.

Başar T, Olsder GJ (1999) *Dynamic Noncooperative Game Theory*, 2nd ed. (Society for Industrial and Applied Mathematics, Philadelphia).

Bernstein F, Federgruen A (2004) Dynamic inventory and pricing models for competing retailers. *Naval Res. Logist.* 51(2):258–274.

Besanko D, Doraszelski U (2004) Capacity dynamics and endogenous asymmetries in firm size. *RAND J. Econom.* 35(1):23–49.

Besanko D, Doraszelski U, Lu LX, Satterthwaite M (2010) Lumpy capacity investment and disinvestment dynamics. *Oper. Res.* 58(4):1178–1193.

Borkovsky R, Doraszelski U, Kryukov Y (2010) A user's guide to solving dynamic stochastic games using the homotopy method. *Oper. Res.* 58(4):1116–1132.

Caro F, Martínez-de-Albéniz V (2010) The impact of quick response in inventory-based competition. *Manufacturing Service Oper. Management* 12(3):409–429.

Dixit A (1980) The role of investment in entry-deterrence. *Econom. J.* 90(357):95–106.

Doraszelski U, Satterthwaite M (2010) Computable Markov-perfect industry dynamics. *RAND J. Econom.* 41(2):215–243.

Fudenberg D, Levine D (1983) Subgame-perfect equilibria of finite- and infinite-horizon games. *J. Econom. Theory* 31(2):251–268.

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

Hall J, Porteus E (2000) Customer service competition in capacitated systems. *Manufacturing Service Oper. Management* 2(2):144–165.

Heyman D, Sobel M (1984) *Stochastic Models in Operations Research*, Vol. II (McGraw Hill, New York).

Huang J, Leng M, Parlar M (2013) Demand functions in decision modeling: A comprehensive survey and research directions. *Decision Sci.* 44(3):557–609.

Ifrach B, Weintraub GY (2012) A dynamic oligopoly model for concentrated industries. Working paper, Columbia University, New York.

Kirman AP, Sobel MJ (1974) Dynamic oligopoly with inventories. *Econometrica* 42(2):279–287.

Lippman SA, McCardle KF (1997) The competitive newsboy. *Oper. Res.* 45(1):54–65.

Lu LX, Lariviere MA (2012) Capacity allocation over a long horizon: The return on turn-and-earn. *Manufacturing Service Oper. Management* 14(1):24–41.

Nagarajan M, Rajagopalan S (2009) A multiperiod model of inventory competition. *Oper. Res.* 57(3):785–790.

Nash JF (1950) Equilibrium points in n-person games. *Proc. Natl. Acad. Sci. USA* 36(1):48–49.

Nash JF (1951) Noncooperative games. *Ann. Math.* 54(2):286–295.

Netessine S, Rudi N, Wang Y (2006) Inventory competition and incentives to back-order. *IIE Trans.* 38(11):883–902.

Olsen TL, Parker RP (2008) Inventory management under market size dynamics. *Management Sci.* 54(10):1805–1821.

Pakes A, McGuire P (1994) Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model. *RAND J. Econom.* 25(4):555–589.

Parker RP, Kapuściński R (2011) Managing a non-cooperative supply chain with limited capacity. *Oper. Res.* 59(4):866–881.

Parlar M (1988) Game theoretic analysis of the substitutable product inventory problem with random demand. *Naval Res. Logist.* 35(3):397–409.

Roy TK, Maiti M (1998) Multi objective inventory model of deteriorating items with some constraints in a fuzzy environment. *Comput. Oper. Res.* 25(12):1085–1095.

Saloner G (1986) The role of obsolescence and inventory costs in providing commitment. *Internat. J. Indust. Organ.* 4(3):333–345.

Schary PB, Christopher M (1979) The anatomy of a stockout. *J. Retailing* 55(2):59–70.

Spence AM (1977) Entry, capacity, investment and oligopolistic pricing. *Bell J. Econom.* 8(2):534–544.

Tirole J (1988) *The Theory of Industrial Organization* (MIT Press, Cambridge, MA).

Topkis DM (1998) *Supermodularity and Complementarity* (Princeton University Press, Princeton, NJ).

Urban TL (2005) Inventory models with inventory-level-dependent demand: A comprehensive review and unifying theory. *Eur. J. Oper. Res.* 162(3):792–804.

Vives X (1999) *Oligopoly Pricing: Old Ideas and New Tools* (MIT Press, Cambridge, MA).

Zeinalzadeh A, Alptekinoğlu A, Arslan G (2010) Inventory competition under fixed order costs. Working paper, Pennsylvania State University, University Park.

Tava Lennon Olsen holds the Ports of Auckland Chair in Logistics and Supply Chain Management at the University of Auckland Business School. She is a past president of the MSOM Society and was awarded the Auckland Business School's Sustained Research Excellence Award.

Rodney P. Parker is an associate professor of operations management at the University of Chicago Booth School of Business. His primary research interests lie in inventory theory, supply chain management, dynamic games, healthcare operations, and finance/operations interface problems.