Capacity Allocation over a Long Horizon: The Return on Turn-and-Earn

March 19, 2011

Abstract

We consider a supply chain in which a supplier sells products to multiple retailers. When orders from the retailers exceed the supplier's capacity, she must employ an allocation mechanism to balance supply and demand. In particular, we consider a commonly used allocation scheme in the automobile industry: turn-and-earn, which uses past sales to allocate capacity. In essence, dealers "earn" allotment of a vehicle after they sell one. In contrast to turn-and-earn, fixed allocation ignores past sales and gives retailers equal share of the capacity. Earlier work has demonstrated that turn-and-earn induces more sales in a two-period setting compared to fixed allocation. The question remains unanswered whether turn-and-earn induces similar behaviors over a long horizon when retailers possess private demand information. We construct a dynamic stochastic game of order competition over an infinite horizon to track the order dynamics of the supply chain. We obtain a richer set of equilibrium behaviors than existing models predict. Instead of a symmetric allocation outcome, we observe that sales leadership may arise in equilibrium and that retailers with different past sales adopt distinct ordering strategies to respond to demand uncertainty. Transient sales dynamics suggest that sales leadership may not be persistent and can be eliminated by the occurrence of extremely low demand. This provides a theoretical explanation for several behavioral observations of some U.S. automobile dealers. In addition to the sales-inducing effect, interestingly, turn-and-earn also causes the retailers to absorb local demand variability.

Key words: Capacity Allocation; Turn-and-Earn; Dynamic Stochastic Game; Markov Perfect Equilibrium.

1 Introduction

Many supply chains share two features: The upstream supplier has limited capacity and there are multiple downstream retailers selling to the public. Limited capacity implies that the supplier will not always be able to fulfill all the orders she receives. Multiple retailers imply that when capacity binds, one retailer's order can only be filled if someone else's order is shorted; that is, some will win while others lose. Although they create headaches for retailers, potential shortages may offer opportunities for the supplier. Just who will win and who will lose depends on the allocation mechanism the supplier chooses. An allocation mechanism is a procedure for converting an infeasible set of orders into a feasible set of capacity assignments. There are a large number of procedures that the supplier could use and the choice matters. How the supplier doles out scarce capacity can have a non-trivial impact on how the retailers order and act, making the supplier's allocation scheme a lever for influencing the actions of her supply chain partners.

We examine these issues by considering turn-and-earn, an allocation scheme frequently used in the automobile industry. Under a turn-and-earn system, retailers earn higher allocations by selling (i.e., turning) more units. Variations of turn-and-earn are used by most car companies in the United States[1] (Lawrence 1996) and some manufacturers are using it in developing markets as well (Shirouzu 2006). Having products on allocation is not particular to the auto industry. Indeed, over the years products ranging from computers (Zarley & Damore 1996) and pharmaceuticals (Hwang & Valeriano 1992) to paper towels and liquid detergent (Associated Press 1997) have been on allocation.

Turn-and-earn has some clear virtues. First, it "moves the metal," providing a clear incentive for dealers to sell cars rapidly as opposed to increasing the price and holding out for top dollar. When it launched, the Buick Enclave was a hot-selling vehicle and was initially allocated under a controlled procedure not dependent on sales rate. In January of 2008 it was moved to turn-and-earn in order to "maximize growth." (LaReau 2008) Turn-and-earn also assures that units are sent to markets where they are most needed. When Toyota entered the Chinese market, cars were first allocated according to an annual plan (the system Toyota uses in Japan). When reality deviated from the plan, some dealers had excess stock while others had nothing to sell. The firm has since

[1] Another common scheme is balanced days supply under which stock is allocated so that all dealers in a region have the same days of inventory. Note that this will behave similarly to turn-and-earn since a dealership can increase its allocation by increasing its sales rate.

moved to a turn-and-earn system (Shirouzu 2006).

On the flip side, turn-and-earn has some acknowledged shortcomings. It can lead to high-pressure sales tactics (Lavin 1994). It has been blamed for inducing dealers to fraudulently distort their sales totals (Lynch 1997). It also has a propensity to lock in market shares. A dealer with a small allocation may find itself caught in a Catch 22: it cannot sell more without boosting its allocation but the only way to increase its allocation is to sell more (Sawyers 1999). A lagging dealer cannot catch up unless high-selling dealers reduce their sales rate. This can play out in several ways. For example, Pontiac tied the allocation of the Sunfire to the sales of the Sunbird, the model it replaced. While the Sunbird had sold well in the Midwest, it had fared poorly in California. When the Sunfire proved popular in California, Pontiac had limited flexibility to allocate more vehicles to that market (Child 1995). Toyota has had similar issues, but in reverse. Its turn-and-earn mechanism had for a long time directed more stock to the West Coast as opposed to the Midwest. A recession-induced sales drop allowed Midwest dealers to catch up. In the words of one Kansas dealer "we've been asking for more inventory for two years and now we finally got it." (Rechtin 2008)

Academic research on turn-and-earn has been limited and mostly focuses on its sales-inducing effect (see §2 for a detailed discussion). Some interesting characteristics of the turn-and-earn system have been overlooked. For example, how would a dealer with a large allocation behave differently from a dealer with a small allocation? Do asymmetric allocations tend to be locked in as in the Pontiac example? Or, are they transitory so that dealers with small allocations eventually overcome the Catch 22 during a time of slow sales as in the Toyota example? How does local demand variation experienced by retailers get propagated in the supply chain under turn-and-earn?

To address these questions, we adopt the Markov perfect equilibrium (MPE) framework of Ericson & Pakes (1995) and construct a dynamic stochastic game of order competition over an infinite horizon. We fix the supplier's capacity and wholesale price throughout the game but allow the retailers to submit orders in each period after demand is realized. Demand varies from period to period with uncertainty coming from two sources: common demand shocks that are publicly observed and local demand shocks that are specific and only known to individual retailers. We prove that the retailers' decision rule is a threshold policy dependent on their realized local demand. However, the complexity of the dynamic stochastic game with incomplete information precludes analytical closed-form solutions and necessitates numerical analysis. Pure-strategy MPE exists for our game, and we solve for them using an iterative algorithm. Our numerical approach borrows the techniques developed in the field of computational economics. As far as we are aware, this paper is among the first to adopt those techniques to study a supply chain problem.

Using computed MPE we show that turn-and-earn induces the retailers to increase sales for the demand states when the supplier's capacity is abundant, boosting the supplier's sales. Moreover, we obtain a richer set of equilibrium behaviors than existing models predict. Instead of a symmetric allocation outcome, we observe that sales leadership may arise in equilibrium and that retailers with different past sales adopt distinct ordering strategies to respond to demand uncertainty. Sales leadership can be maintained as long as demand is not extremely low and the expected return of a larger allocation in the future is sufficiently high – a situation that may have occurred in the Pontiac case. Nevertheless, occurrence of extremely low demand makes sales leadership too costly to maintain. The resulting excess capacity provides an opportunity for sales laggards to catch up – a situation that likely mirrors the Toyota case. Finally, we demonstrate that turn-and-earn can cause the retailers to alter their orders in a way that reduces the demand variability placed on the supplier. That is, the retailers absorb local demand fluctuations.

In what follows, we first review the related literature in §2 and then lay out the base model in §3. We characterize the structure of the MPE in §4. We compute the MPE and analyze the numerical results in §5. We consider several extensions of the base model in §6 and conclude in §7. All proofs are relegated to the Online Appendix.

2 Literature Review

How to allocate scarce supply is relevant in any distribution system with limited capacity or inventory whether the system is run by a single decision maker or made up of independent firms (Zipkin 2000). Here we focus on allocation schemes in supply chains made up of independent self-interested firms. Academic research interest in this area can be traced to Lee, Padmanabhan & Whang (1997)'s seminal work on the bullwhip effect. They acknowledge that turn-and-earn reduces the bullwhip effect. A well accepted explanation is that the bullwhip effect is counteracted because turn-and-earn separates allocation from current orders and thus mitigates order inflation.

Work subsequent to Lee et al. (1997) roughly breaks into two streams. One focuses on a single sales period and considers the performance of different allocation schemes that depend only on current orders. Cachon & Lariviere (1999*b*) characterize which allocation schemes induce order inflation and which induce the retailers to order truthfully. They show that supply chain performance may improve with order inflation since it reduces double marginalization. If the supplier is able to choose its capacity level, the retailers may benefit from an allocation scheme that induces order inflation since it results in the supplier building additional capacity. This work assumes that retailers are local monopolists. Furuhata & Zhang (2006), Furuhata, Perrussel & Zhang (2007), and Chen, Li & Zhang (2007) study related issues when retailers compete for customers. Similar allocation schemes have also been studied in a queueing system that allocates demand among multiple servers, e.g., Gilbert & Weng (1998) and Cachon & Zhang (2007).

The second stream focuses on turn-and-earn allocation. Cachon & Lariviere (1999*a*) consider a two-period model in which past sales may be used as a basis for future capacity allocation between two retailers. Demand is stochastic but identical across markets. Turn-and-earn induces the retailers to increase sales in the first period in an attempt to secure a better allocation in the second period. In equilibrium, the retailers play to a draw – because they are completely symmetric, neither can gain an advantage over the other, and they receive the same allocation. In a recent paper, Purohit & Vernik (2008) present a variation of Cachon & Lariviere (1999*a*) in which the supplier sells two substitutable products through the retailers. Only one of the products is possibly capacity constrained. They find that it is generally better to have the second-period allocation of the scarce product depend on first-period sales of both products.

Our model also has parallels with Cachon & Lariviere (1999*a*). We too consider a supplier selling to multiple retailers and model turn-and-earn in a similar fashion. Our model differs from Cachon & Lariviere (1999*a*) along three key dimensions. First, we consider an infinite horizon. In a two-period model, retailers simply maximize their current profit in the second period. Thus, turn-and-earn induced asymmetric allocation, if any, would end in the second period. With an infinite horizon, however, we can examine the persistence of asymmetric allocation over time.

Second, our model has three demand states. At first glance, this seems a minor tweak, but in reality it adds significant depth to the model. For allocation to be interesting, the supplier's capacity must bind in some demand state but not in others. With just two demand states, if the retailers increase their sales in the low demand state, asymmetric allocations would persist forever – if an allocation is worth fighting for even low demand, it is worth maintaining under all demand conditions. With three demand states, however, we can not only have a medium demand state when leadership is worth fighting for, but also have a low demand state with market conditions sufficiently poor that makes a sales lead not worth maintaining. Effectively, our model allows for overall market slumps (like that experienced by the dealers in the Toyota example) that re-level the playing field.

Third, our model allows for local variation in demand. A common demand state determines whether or not the supplier has sufficient capacity but random local shocks determine how much each retailer wants to sell. This allows asymmetric allocations to arise naturally even though the retailers follow symmetric strategies. Moreover, it allows us to evaluate the impact of turn-and-earn on demand variance propagation in the supply chain. We show that turn-and-earn induces the retailers to absorb local variability in their markets instead of passing it on to the supplier.

Finally, this paper borrows the computational methodologies developed in the Markov perfect industry dynamics literature, e.g., Ericson & Pakes (1995). Some recent papers in the operations management literature have also used similar computational methods to solve dynamic games, e.g., Besanko, Doraszelski, Lu & Satterthwaite (2010), Weintraub, Benkard & Roy (2010), Mookherjee & Friesz (2008), and Perakis & Sood (2006). While the Ericson-Pakes MPE framework has received much attention in the industrial organization literature, the present paper is arguably the first to apply the framework to a supply chain context.

3 The Model

3.1 A Dynamic Stochastic Game of Order Competition

Consider a supplier and two retailers, indexed by i, $i = 1, 2$. We will consider N retailers in §6.1. The supplier has fixed capacity K and fixed wholesale price w. Setting these two parameters exogenously allows us to ignore the supplier's problem and focus on the strategic interactions between the retailers. The retailers' markets are geographically isolated and thus all retailers would behave like monopolists in their own market if they did not have to compete for the supplier's capacity. We consider an infinite horizon game between the retailers.

Demand. Let p_i and q_i denote retailer i's price and sales quantity, respectively. He faces a linear demand $p_i = a_i - bq_i$, where a_i is stochastic and varies from period to period. b is constant across periods and identical for both retailers. We assume nonnegative prices, i.e., $p_i \geq 0$. Such stochastic linear demand functions have been commonly used in the operations management literature, e.g., Cachon & Lariviere (1999a), Chod & Rudi (2005). The stochastic component of the demand function, a_i, represents the demand state and is the sum of two parts: $a_i = \alpha + \varepsilon\theta_i$. The first part, α, is the common demand state publicly observed, drawn from set $\mathcal{A} \equiv \{\alpha_u\}_{u=1,\ldots,U}$, where $U \in \mathbb{N}$, $\alpha_u > 0$, and $\alpha_u < \alpha_v$ for any $u < v$, $u, v = 1, ..., U$. At the end of each period, the common demand transitions to another state according to transition matrix $P \equiv [p_{uv}]$ and the resulting demand process forms a discrete time Markov chain. (Figure 8 in the Online Appendix illustrates a three-state common demand process.)

The second part of the demand state, $\varepsilon\theta_i$, is the product of retailer i's *privately* observed local demand state, θ_i, and a scale factor, ε, which is identical for all retailers. θ_i is a random variable independently and identically distributed across periods and retailers. It has a continuous distribution function $F(\cdot)$ with support $[\underline{\Theta}, \bar{\Theta}]$ and $\mathbb{E}(\theta_i) = 0$, where $|\underline{\Theta}|, |\bar{\Theta}| < \infty$. We assume $\varepsilon > 0$ so that it scales the variability of local demand. We further assume ε is small enough such that $\alpha + \varepsilon\underline{\Theta} \geq 0$, i.e., demand intercepts are nonnegative for the entire support of θ_i. Admittedly, this is a crude way of modeling local demand dispersion, but it allows us to capture the essence of demand information asymmetry without sacrificing simplicity.

State of the World. Let s_i denote retailer i's sales quantity in the previous period. The state of the world consists of three parts: sales of the previous period $\mathbf{s} = (s_1, s_2)$, common demand state α, and local demand state θ_i. Notice that vector $\mathbf{s}$ is bolded. We will use boldface letters to denote all vectors introduced hereafter. The starting point of the game is exogenously chosen and denoted by $(\mathbf{s}^0, \alpha^0)$.

Sequence of Events. In any period, the following sequence of events occur: 1) The retailers learn the sales state and common demand state. Each retailer is also privately informed with his local demand state; 2) The retailers submit their orders $\mathbf{x} = (x_1, x_2)$; 3) The supplier allocates capacity $\mathbf{q} = (q_1, q_2)$ according to a posted allocation mechanism, where $q_i \equiv q_i(\mathbf{s}, \mathbf{x})$.

We assume that the retailers play *Markovian* strategies and derive Markov perfect equilibrium (MPE). The restriction to Markovian strategies implies that every retailer's action in any period depends only on the current state, i.e., $x_i \equiv x_i(\mathbf{s},\alpha, \theta_i)$, and thus excludes strategies that depend on the entire history. The benefits and the behavioral assumptions of using Markovian strategies are discussed in Maskin & Tirole (2001). Since the horizon is infinite and the influence of past play is captured in the current state, there is a one-to-one correspondence between subgames and states. Hence, any MPE in our model is subgame perfect. We further assume that retailer i chooses a smaller order quantity when two order quantities yield identical capacity allocation. This is without loss of generality because in this case larger orders have no impact on the equilibrium sales.

In any period, retailer i sells q_i and earns profit $\pi_i(\mathbf{s},\alpha, \theta_i, \mathbf{x}) \equiv (\alpha + \varepsilon\theta_i - bq_i(\mathbf{s}, \mathbf{x}) - w)\, q_i(\mathbf{s}, \mathbf{x})$. Note that we do not restrict the sign of π_i, i.e., retailer i can earn a negative profit. We assume that the retailers have to sell all they are allocated. Hence, the next period's sales state $\mathbf{s}' \equiv (s'_1, s'_2) = (q_1, q_2)$. This assumption is not unrealistic given that we consider no inventory carryover from period to period. This is also the assumption made by Cachon & Lariviere (1999a) in their order competition game without inventory. The retailers' discount factors are identical and equal to β. Retailer i maximizes his total discounted profit, denoted by $V_i(\mathbf{s}, \alpha, \theta_i)$, which is his value function. Further, we define retailer i's ex ante value function: $V_i(\mathbf{s}, \alpha) \equiv \int V_i(\mathbf{s}, \alpha, \theta_i) dF(\theta_i)$.

3.2 Allocation Mechanisms

We consider two allocation mechanisms: turn-and-earn and fixed allocation. Under each mechanism, a retailer's *maximum allocation* is capped by a rule specific limit, i.e., the *guaranteed allocation*, or the capacity unclaimed by the other retailer, whichever is larger. The guaranteed allocation is the sum of the reserved capacity plus an equal share of the unreserved capacity.

3.2.1 Turn-and-Earn

Under turn-and-earn, the supplier reserves some current-period capacity equal to $|s_1 - s_2|$ for the previous-period sales leader. The scheme then allocates the rest equally among the two retailers. Hence, the leader's guaranteed allocation is $|s_1 - s_2| + (K - |s_1 - s_2|)/2$. The laggard's guaranteed allocation is $(K - |s_1 - s_2|)/2$. Suppose retailer 1 is the leader. Given orders $\mathbf{x}$ and past sales $\mathbf{s}$, the allocations are

$$
\begin{aligned}
q_1(\mathbf{s}, \mathbf{x}) &= \min\left\{x_1, |s_1 - s_2| + \frac{K - |s_1 - s_2|}{2} + (K - x_2)^+\right\}, \\
q_2(\mathbf{s}, \mathbf{x}) &= \min\left\{x_2, \frac{K - |s_1 - s_2|}{2} + (K - x_1)^+\right\}.
\end{aligned}
$$

The way we model turn-and-earn follows exactly that of Cachon & Lariviere (1999a), thereby allowing us to make a direct comparison with their results. This definition is probably the simplest form that captures the essence of this allocation scheme. Nevertheless, in practice turn-and-earn is more complicated.[2] In §6, we consider two variants of the turn-and-earn scheme: proportional turn-and-earn and moving average turn-and-earn.

3.2.2 Fixed Allocation

Under fixed allocation, each retailer has no reserved capacity, and is guaranteed with an equal share of the supplier's capacity. The guaranteed allocation for either retailer is $K/2$. The allocation for

[2] For a more detailed discussion see www.f150online.com/forums/2004-present/127481-how-works-priority-codes-allocation-scheduling.html

retailer i is

$$q_i(\mathbf{s}, \mathbf{x}) = \min\left\{x_i, \frac{K}{2} + (\frac{K}{2} - x_j)^+\right\},\ i \neq j.$$

By linking current allocation to past sales, turn-and-earn introduces intertemporal dependence in retailers' ordering behavior. In contrast, fixed allocation disregards past sales, and thus provides a *myopic* benchmark for the retailers' order quantities.

3.3 The Pareto Mechanism and the First Best

To evaluate the performance of turn-and-earn, we derive allocations under the Pareto mechanism and the first best. Cachon & Lariviere (1999*b*) define the Pareto mechanism as the allocation that would maximize the sum of the retailers' current profits given the wholesale price. From the Pareto mechanism, we can determine the first best outcome for the supply chain by setting the wholesale price equal to the cost of production. For simplicity, we normalize the cost of production to zero. In every period, the Pareto allocation is determined by the following maximization problem.

$$\max_{\mathbf{q}\in\mathbb{R}^2_+} \sum_{i=1}^{2}(\alpha + \varepsilon\theta_i - bq_i - w)q_i \text{ s.t. } q_1 + q_2 \leq K$$

Solving the problem yields for $i, j = 1, 2,\ i \neq j$,

$$q_i^P(\alpha, \theta_1, \theta_2) = \begin{cases} \frac{\alpha+\varepsilon\theta_i - w}{2b} & \text{if } \theta_1 + \theta_2 \leq \frac{2}{\varepsilon}(bK - \alpha + w) \\ \frac{K}{2} + \frac{\varepsilon(\theta_i - \theta_j)}{4b} & \text{if } \theta_1 + \theta_2 > \frac{2}{\varepsilon}(bK - \alpha + w) \end{cases} \quad (1)$$

It is straightforward to show that the Pareto mechanism gives identical expected total sales as fixed allocation. (Notice that under both mechanisms the retailers order their monopolist quantities when capacity is not tight while the total sales quantity is equal to capacity when it is tight.) However, fixed allocation may not distribute capacity among the retailers efficiently when capacity is tight because it treats all retailers equally. Suppose every retailer's optimal sales quantity exceeds his guaranteed allocation $K/2$ under fixed allocation, then $\frac{\varepsilon(\theta_i - \theta_j)}{4b}$ is its deviation from efficiency, as shown by Eq.(1). Thus, while fixed allocation delivers the same sales quantity as the Pareto scheme, it may fail to maximize revenue. Compared to the first best, the Pareto mechanism leads to lower sales due to double marginalization.

3.4 State Discretization and Lumpy Orders

The dynamic stochastic game with incomplete information does not yield closed-form solutions under turn-and-earn. We resort to an iterative algorithm to solve for the MPE. To simplify the computational procedure, we discretize the state space by restricting that orders must take on one of $M+1$ values listed in an increasing order: $0, \delta, ..., M\delta$, where $\delta = K/M$ is the *batch* size and measures the lumpiness of order. This assumption is without loss of generality because there is no restriction on the value of δ. On the other hand, order lumpiness finds its roots in practice. Retailers order in bulk to avoid incurring fixed costs of ordering frequently. Batch ordering policies are also commonplace in the inventory management literature, e.g., continuous-review (R, Q) policies.

4 Properties of the MPE

Retailer i's optimal order quantity depends on the current system state and his own local demand state. We use an indicator function $\mathbf{1}[x_i(\mathbf{s}, \alpha, \theta_i) = m\delta]$, $m = 0, ..., M$, to denote retailer i's ordering action. Integrating the indicator function over the support of θ_i gives the probability of ordering m batches: $\sigma_{im}(\mathbf{s}, \alpha) \equiv \int \mathbf{1}[x_i(\mathbf{s}, \alpha, \theta_i) = m\delta] dF(\theta_i)$. Because his local demand state is private, retailer i's equilibrium order quantity, although a pure strategy, is perceived by his opponent as a probability distribution over the space of order quantities. Since there are $M+1$ possible order quantities, we use $\boldsymbol{\sigma}_i(\mathbf{s},\alpha)$ to denote the order distribution with $M+1$ components: $\boldsymbol{\sigma}_i(\mathbf{s}, \alpha) \equiv \{\sigma_{i0}(\mathbf{s}, \alpha), ..., \sigma_{iM}(\mathbf{s}, \alpha)\}$. In equilibrium, when retailer i decides to order m batches, the decision is based on the fact that he knows his opponent's order quantity up to a probability distribution determined by $\boldsymbol{\sigma}_j(\mathbf{s}, \alpha)$. When retailer i's order is below his guaranteed allocation, his allocation is equal to his order, and thus a deterministic number. Otherwise, his allocation becomes stochastic and depends on the opponent's actual order. Therefore, from retailer i's perspective, his expected allocation is a random variable with a probability distribution jointly determined by the opponent's order distribution $\boldsymbol{\sigma}_j(\mathbf{s}, \alpha)$ and the allocation mechanism.

To characterize the MPE, we start by writing retailer i's Bellman equation as follows:

$$V_i(\mathbf{s}, \alpha, \theta_i) \equiv \max_{x_i \in \{0\delta, ..., \bar{m}\delta\}} \mathbb{E}\left[\pi_i(\mathbf{s},\alpha, \theta_i, x_i, x_j)\right] + \beta \mathbb{E}\left[V_i(\mathbf{s}', \alpha') | \mathbf{s}, \alpha, x_i, x_j\right],$$

where $\bar{m} \equiv \max\{m \in \mathbb{Z} | \alpha + \varepsilon\theta_i - bm\delta \geq 0\}$. This upper bound ensures that the price is nonnegative. For simplicity, we write x_i with the understanding that it is a function of $(\mathbf{s}, \alpha, \theta_i)$. The maximization problem yields retailer i's best response for a given strategy profile of retailer j, i.e., $x_j(\mathbf{s}, \alpha, \theta_j)$. Because retailer j's order distribution, represented by $\boldsymbol{\sigma}_j(\mathbf{s}, \alpha)$, is discrete with $M+1$ possible order quantities, we can express the expectations in the Bellman equation as a sum weighted by

$\sigma_{jl},\, j \neq i,\, l = 0, ..., M$:

$$V_i(\mathbf{s}, \alpha, \theta_i) = \max_{x_i \in \{0\delta, ..., \bar{m}\delta\}} \sum_{l=0}^{M} \sigma_{jl}(\mathbf{s}, \alpha) \left\{ \pi_i(\mathbf{s}, \alpha, \theta_i, x_i, l\delta) + \beta \sum_{\alpha' \in \mathcal{A}} P(\alpha'|\alpha) V_i\left(\mathbf{q}(\mathbf{s}, x_i, l\delta), \alpha'\right) \right\} \quad (2)$$

A realization of the rival retailer's order quantity $x_j = l\delta$ occurs with probability $\sigma_{jl}(\mathbf{s}, \alpha)$. The summation in the Bellman equation reflects the fact that retailer i's total discounted profit is an expectation over the order distribution of his opponent.

Solving the maximization problem in Eq.(2) yields that retailer i's optimal order quantity follows a threshold policy depending on his private local demand state – it is optimal for him for order m batches if his local demand state is between certain thresholds. This property is stated formally in Proposition 1 next. For convenience, we introduce two notations: $\eta_{i0}(\mathbf{s}, \alpha) = \underline{\Theta}$, $\eta_{i,M+1}(\mathbf{s}, \alpha) = \bar{\Theta}$.

Proposition 1 *Suppose that $\mathbf{q}(\mathbf{s}, \mathbf{x})$ is weakly increasing in $\mathbf{x}$. For any $m \in \{0, ..., M\}$, it is optimal for retailer i to order $m\delta$ if and only if $\eta_{im}(\mathbf{s}, \alpha) \le \theta_i \le \eta_{i,m+1}(\mathbf{s}, \alpha)$, where $\eta_{im}(\mathbf{s}, \alpha) \in [\underline{\Theta}, \bar{\Theta}]$ is a function of $V_i(\mathbf{s}, \alpha)$ and $\boldsymbol{\sigma}_j(\mathbf{s}, \alpha)$. Accordingly, his optimal order distribution $\boldsymbol{\sigma}_i(\mathbf{s}, \alpha)$ is given by*

$$\sigma_{im}(\mathbf{s}, \alpha) = \int \mathbf{1}[\eta_{im}(\mathbf{s}, \alpha) \le \theta_i \le \eta_{i,m+1}(\mathbf{s}, \alpha)] dF(\theta_i), m = 0, ..., M.$$

The threshold property stems from two facts. First, a retailer's myopic sales quantity is increasing in his local demand state θ_i. Second, capacity allocation is weakly increasing in order quantity, which is true for both turn-and-earn and fixed allocation. Taken together, these two monotone properties yield that it is better for retailer i to order a larger quantity when his local demand increases because it improves both the current profit and the future capacity allocation. Proposition 1 also ensures that $\boldsymbol{\sigma}_i(\mathbf{s}, \alpha)$ represents a proper discrete distribution function. $\eta_{im}(\mathbf{s}, \alpha)$ characterizes the optimal equilibrium policy of retailer i for a given value function and the strategy profile of his opponent. The expressions of $\eta_{im}(\mathbf{s}, \alpha)$'s, given in the proof, allow us to solve the Bellman equation.

4.1 Equilibrium Definition

Because of the private local demand information that each retailer possesses, our definition of MPE is more involved than that of the stochastic games with complete information described in Fudenberg & Tirole (1991). From the perspective of retailer j, retailer i acts as if he plays a mixed strategy determined by an order distribution even though he plays a pure strategy conditional on his "type" θ_i. Instead of stating the equilibrium strategy profile as a function of types, we derive the ex ante formulation of the Bellman equation by integrating the value function of retailer i over the support of θ_i: $V_i(\mathbf{s},\alpha) \equiv \int V_i(\mathbf{s},\alpha,\theta_i)dF(\theta_i)$. This ex ante formulation is equivalent to retailer i maximizing his utility conditional on θ_i for all $\theta_i \in [\underline{\Theta}, \bar{\Theta}]$.

We have shown that retailer i's optimal ordering policy is a threshold rule dependent on his local demand state θ_i. Integration of retailer i's maximization problem over the support of θ_i can be transformed into choosing an M-dimensional vector $(\eta_{i1}(\mathbf{s},\alpha), ..., \eta_{iM}(\mathbf{s},\alpha))$ that determines the probability distribution of retailer i's optimal orders. Let $\underline{\eta}_{im}(\mathbf{s},\alpha) \equiv \max\left\{\underline{\Theta},\ \min\left\{\bar{\Theta},\ (bm\delta-\alpha)/\varepsilon\right\}\right\}$, $m = 1, ..., M$. This threshold is the lower bound of θ_i above which selling m batches yields a non-negative price. By imposing this lower bound on every component of $(\eta_{i1}(\mathbf{s},\alpha), ..., \eta_{iM}(\mathbf{s},\alpha))$, we make sure negative prices never arise in optimality. Hence, integrating Eq.(2) over the support of θ_i yields

$$\begin{aligned} V_i(\mathbf{s},\alpha) &= \int \max_{\eta_{im}(\mathbf{s},\alpha)\in[\underline{\eta}_{im}(\mathbf{s},\alpha),\bar{\Theta}],\ m=1,...,M} \sum_{l=0}^{M} \sigma_{jl}(\mathbf{s},\alpha) \sum_{m=0}^{M} \mathbf{1}\left[\eta_{im}(\mathbf{s},\alpha) \le \theta_i \le \eta_{i,m+1}(\mathbf{s},\alpha)\right] \\ &\quad \cdot \left\{\pi_i(\mathbf{s},\alpha,\theta_i,m\delta,l\delta) + \beta \sum_{\alpha'\in\mathcal{A}} P(\alpha'|\alpha) V_i\left(\mathbf{q}(\mathbf{s},m\delta,l\delta),\alpha'\right)\right\} dF(\theta_i) \\ &= \max_{\eta_{im}(\mathbf{s},\alpha)\in[\underline{\eta}_{im}(\mathbf{s},\alpha),\bar{\Theta}],\ m=1,...,M} \sum_{l=0}^{M} \sigma_{jl}(\mathbf{s},\alpha) \left\{\sum_{m=0}^{M} \int_{\eta_{im}(\mathbf{s},\alpha)}^{\eta_{i,m+1}(\mathbf{s},\alpha)} \pi_i(\mathbf{s},\alpha,\theta_i,m\delta,l\delta)dF(\theta_i)\right. \\ &\quad \left. + \beta \sum_{k=0}^{M} \sum_{\alpha'\in\mathcal{A}} \left(F(\eta_{i,m+1}(\mathbf{s},\alpha)) - F(\eta_{im}(\mathbf{s},\alpha))\right) P(\alpha'|\alpha) V_i\left(\mathbf{q}(\mathbf{s},m\delta,l\delta),\alpha'\right)\right\} \end{aligned} \quad (3)$$

The order distribution $\boldsymbol{\sigma}_i$, determined by the optimal $(\eta_{i1}(\mathbf{s},\alpha), ..., \eta_{iM}(\mathbf{s},\alpha))$ is the ex ante policy function perceived by retailer i's opponent and the supplier. The above integrated Bellman equation provides the basis for defining the MPE.

Definition 1 *A Markov perfect equilibrium of the order competition game involves value function V_i and policy function $\boldsymbol{\sigma}_i$ such that for all $i, j = 1, 2,\ i \ne j$,*

(i) given $\boldsymbol{\sigma}_j$, V_i solves the Bellman equation (3) and

(ii) given $\boldsymbol{\sigma}_j$ and V_i, $(\eta_{i1}(\mathbf{s},\alpha), ..., \eta_{iM}(\mathbf{s},\alpha))$ solves the maximization problem in Eq.(3) for all $(\mathbf{s},\alpha)$, and $\boldsymbol{\sigma}_i(\mathbf{s},\alpha)$ is determined by $(\eta_{i1}(\mathbf{s},\alpha), ..., \eta_{iM}(\mathbf{s},\alpha))$ according to $\sigma_{im}(\mathbf{s},\alpha) = \int \mathbf{1}[\eta_{im}(\mathbf{s},\alpha) \le \theta_i \le \eta_{i,m+1}(\mathbf{s},\alpha)]dF(\theta_i)$, $m = 0, ..., M$.

Such an MPE exists and is computable according to Doraszelski & Satterthwaite (2010).[3]

[3] Uniqueness is generally not guaranteed in such dynamic stochasic games. However, in our computational study, we have not found any multiple equilibria even when we initiate the iterative algorithm from different starting points.

4.2 Fixed Allocation

The above equilibrium characterization applies to any allocation mechanism that satisfies that $\mathbf{q}(\cdot)$ is weakly increasing in a retailer's order quantity, which is true for fixed allocation. Therefore, we can derive the solutions for fixed allocation by setting $V_i\left(\mathbf{q}(\mathbf{s},m\delta, x_j), \alpha'\right) = V_i\left(\mathbf{q}(\mathbf{s},m'\delta, x_j), \alpha'\right)$ for any $m, m' = 0, ..., M$ in Eq.(3) because current-period orders do not affect a retailer's future payoff. Nevertheless, due to the property that fixed allocation does not depend on past sales, we can derive the closed-form solution for the MPE without invoking the threshold property. We simplify the state of the world to include only the common and local demand states. Because all retailers would behave myopically, it is useful to derive their *myopic sales quantity*:

$$q^M(\alpha, \theta) = \left[\frac{(\alpha + \varepsilon\theta - w)^+}{2b}\right], \tag{4}$$

where $[\cdot]$ is the operator that finds the nearest integer number of batches that maximizes the single period profit.

Proposition 2 *Under fixed allocation, a Markov perfect equilibrium is a profile of order distributions* $(\boldsymbol{\sigma}^F, \boldsymbol{\sigma}^F)$, *where* $\boldsymbol{\sigma}^F$ *is an* $M+1$ *dimensional vector given by* $\sigma_m^F(\alpha) = \int \mathbf{1}[x^F(\alpha, \theta) = m\delta]dF(\theta)$, $m = 0, ...M$, $\alpha \in \mathcal{A}$, *where* $x^F(\alpha, \theta) = \min\left\{q^M(\alpha, \theta), \frac{K}{2} + \left(\frac{K}{2} - q^M(\alpha, \underline{\Theta})\right)^+\right\}$.

The second term in the expression of $x^F(\alpha, \theta)$ is retailer i's maximum allocation which is a sum of two parts. The first part is his guaranteed allocation while the second is the maximum residual capacity unclaimed by the other retailer, which occurs when he experiences the lowest local demand. Other equilibria exist, for example, setting $x^F(\alpha, \theta) = q^M(\alpha, \theta)$ above gives a different equilibrium. Nevertheless, these equilibria yield the same sales quantities. The one we specify in Proposition 2 has the smallest order quantities.

4.3 Symmetric Equilibria and State-to-State Transitions

Notice that the definitions of sales quantity is symmetric in the sense that $q_1(s_1, s_2, x_1, x_2) = q_2(s_2, s_1, x_2, x_1)$. Given a sales quantity, differences in retailer profits result from idiosyncratic local demand shocks. Since the distribution of local demand shocks is identical for all retailers, it is reasonable to suppose that their equilibrium behavior is symmetric. From now on, we focus on symmetric MPE. Specifically, a symmetric equilibrium satisfies $V_1(s_1, s_2, \alpha) = V_2(s_2, s_1, \alpha)$ and $\boldsymbol{\sigma}_1(s_1, s_2, \alpha) = \boldsymbol{\sigma}_2(s_2, s_1, \alpha)$ for all α. Thus, we drop the retailer-denoting subscripts and define

$V(\mathbf{s},\alpha) \equiv V_1(\mathbf{s},\alpha)$ and $\boldsymbol{\sigma}(\mathbf{s},\alpha) \equiv \boldsymbol{\sigma}_1(\mathbf{s},\alpha)$. Using this definition, we can derive retailer 2's value and policy functions as $V_2(\mathbf{s},\alpha) = V(s_2, s_1, \alpha)$ and $\boldsymbol{\sigma}_2(\mathbf{s},\alpha) = \boldsymbol{\sigma}(s_2, s_1, \alpha)$.

With the equilibrium policy functions we compute the state-to-state transition probability:

$$P(\mathbf{s}', \alpha' | \mathbf{s}, \alpha) = \sum_{k=0}^{M} \sum_{l=0}^{M} \sigma_{1k}(\mathbf{s},\alpha)\sigma_{2l}(\mathbf{s},\alpha)\mathbf{1}[\mathbf{s}' = \mathbf{q}(\mathbf{s}, k\delta, l\delta)]P(\alpha'|\alpha). \tag{5}$$

We further use these probabilities to compute the transient and limiting distributions of equilibrium sales. Let P be the $(M+1)^2U \times (M+1)^2U$ transition matrix for the Markov process of the equilibrium sales dynamics. The transient distribution over states in period T is given by $\mu^{(T)} = \mu^{(0)}P^T$, where $\mu^{(0)}$ is the $1 \times (M+1)^2U$ initial distribution. The limiting distribution over states solves $\mu^{(\infty)} = \mu^{(\infty)}P$. The limiting distribution of orders, denoted by $\mu^{(\infty)}_{\text{orders}}$, can be conveniently calculated from $\mu^{(\infty)}$. Let P_{orders} be given by

$$P_{\text{orders}}(\mathbf{x}, \alpha' | \mathbf{s}, \alpha) = \sum_{k=0}^{M} \sum_{l=0}^{M} \sigma_{1k}(\mathbf{s},\alpha)\sigma_{2l}(\mathbf{s},\alpha)P(\alpha'|\alpha).$$

Then $\mu^{(\infty)}_{\text{orders}} = \mu^{(\infty)}P_{\text{orders}}$.

5 Numerical Results

In this section, we characterize the limiting and transient characteristics of the MPE. We aim to suggest answers to three questions. First, does turn-and-earn induce more sales over a long horizon? Second, how does sales leadership, i.e., asymmetric allocation, change over time? Third, how does local demand variation affect the variability of the supplier's sales?

5.1 Computational Algorithm and Parameterization

Before we delve into the results, it is instructive to describe the computational algorithm of our dynamic programs. To compute the symmetric MPE, we use a variant of the best response iterative algorithm described in Pakes & McGuire (1994). The algorithm takes a value function $\hat{V}(\mathbf{s},\alpha)$ and a policy function $\hat{\boldsymbol{\sigma}}(\mathbf{s},\alpha)$ as its input and generates updated value and policy functions as its output. Each iteration proceeds as follows: First, we use the results in Proposition 1 to compute retailer 1's order distribution $\boldsymbol{\sigma}(\mathbf{s},\alpha)$ using $\hat{V}(\mathbf{s},\alpha)$ and retailer 2's order distribution $\hat{\boldsymbol{\sigma}}(s_2, s_1, \alpha)$. Second, we compute the updated $V(\mathbf{s},\alpha)$ and $\boldsymbol{\sigma}(\mathbf{s},\alpha)$ using Eq.(3). The iteration is completed by assigning $V(\mathbf{s},\alpha)$ to $\hat{V}(\mathbf{s},\alpha)$ and $\boldsymbol{\sigma}(\mathbf{s},\alpha)$ to $\hat{\boldsymbol{\sigma}}(\mathbf{s},\alpha)$. The algorithm terminates once the relative changes in the value and the policy functions from one iteration to the next are below a pre-specified tolerance. For our numerical study, we set it at 10^{-4}. All programs are written in Matlab 7.

We assume that $F(\cdot)$ is a Beta distribution rescaled over support $[-1, 1]$ with parameters $(3, 3)$ (see the Online Appendix for the specification of the distribution function). This distribution has a single mode at and is symmetric around 0. We have considered other distribution functions with a finite support, such as uniform and triangular. We choose this distribution over uniform distribution because we intend to use one that puts more weight around the mean. This distribution also has an advantage over a triangular one because of its smoothness, which generally facilitates convergence of our dynamic program.

To help analyze the effect of demand transition probabilities, we apply the concept of stochastic ordering of a stochastic process. A transition matrix can be ranked according to the stochastic first order (denoted by $\leq_{st}$) and so does a stochastic process. With identical initial distribution, demand process X' with transition matrix P' is first-order stochastically larger than demand process X with transition matrix P if (i) $P \leq_{st} P'$ and (ii) P or P' is $\leq_{st}$-monotone (Muller & Stoyan 2002). A sufficient condition for (i) is $P[i,*] \leq_{st} P'[i,*]$ for any i and a sufficient condition for (ii) is $P[i,*] \leq_{st} P[i+1,*]$ for any i. For our computational study, we consider three common demand states, $\mathcal{A} = \{\alpha_1, \alpha_2, \alpha_3\}$, where $\alpha_1 < \alpha_2 < \alpha_3$. We will refer to them as low, medium, and high demand, respectively. We use a demand process characterized by the following transition matrix with $1 > \rho_1 \geq \rho_2 > 0$:

$$P = \begin{bmatrix} 1-\rho_1 & \rho_1 & 0 \\ 1-\rho_1 & \rho_1-\rho_2 & \rho_2 \\ 0 & 1-\rho_2 & \rho_2 \end{bmatrix}.$$

Checking conditions (i) and (ii) yields that the demand process becomes first-order stochastically larger as ρ_1 or ρ_2 increases. To simplify exposition, we focus on a special case by setting $\rho_1 = \rho_2$ and denote $\rho \equiv \rho_1$. A larger ρ implies the probability of switching to a higher (lower) demand state increases (decreases) and thus a first-order stochastically larger demand. We explore the cases of $\rho_1 > \rho_2$ in the Extensions.

We set the order batch size $\delta = 2$. Like many discrete dynamic programs, ours suffers from the curse of dimensionality: the state space of the game expands exponentially with the number of states and actions. We thus have to trade-off between a finer state space (i.e., reducing the size of δ) and the flexibility in the order quantities (i.e., allowing the retailers to order any number of batches up to their maximum allocation). Because we are interested in finding out how much the equilibrium orders are different from their myopic counterparts, we maintain an admittedly crude

α	5	9	13	17	21	25	29	33	37	41
q_0^M	δ	2δ	3δ	4δ	5δ	6δ	7δ	8δ	9δ	10δ

Table 1: Expected Myopic Sales Quantities Under No Capacity Constraint

state space but allow full order flexibility to reduce computational burden.[4]

We set the slope of the linear demand function $b = 1$, the constant wholesale price $w = 1$. The discount factor β is set at 0.952, corresponding to a 5% interest rate. We vary the values of K, $\mathcal{A}$, ρ, and ε as each of the following experiments indicates. Expected myopic sales quantities, q_0^M, can be calculated by simply setting $\theta = 0$ in Eq.(4): $q_0^M \equiv q^M(\alpha, 0)$, which are tabulated in Table 1. We will refer back to these numbers in our subsequent analysis.

5.2 Equilibrium Orders and Sales

We characterize steady-state properties of the MPE by computing limiting distributions of the Markov process associated with the dynamic game. We analyze two sets of limiting distributions: one that is averaged over all demand states and the other that is conditional on a common demand state. For the remainder of the paper in all figures and tables, we use "Ave" to denote the former and L, M, and H to denote the limiting distribution conditional on the low, medium, and high demand state, respectively. Consider an example where $K = 8\delta$, $\mathcal{A} = \{5, 9, 21\}$, $\rho = 0.9$, and $\varepsilon = 0.01$. The steady-state equilibrium order quantities for both retailers are $(\delta, 2\delta, 4\delta)$ under fixed allocation and $(\delta, 4\delta, 4\delta)$ under turn-and-earn. The difference lies in the medium demand state – turn-and-earn induces both retailers to order two more batches than the myopic quantity.

When we increase local demand variability by increasing ε, dispersion of equilibrium orders occurs. It is helpful to illustrate order distributions using a three dimensional graph. Figure 1 displays the limiting distributions of orders conditional on the common demand state for $\varepsilon = 4$. The horizontal axes represent the number of batches of equilibrium orders by retailers 1 and 2. The vertical axis represents the probability of that pair of orders. The supplier's capacity, 8δ, determines the highest possible order and thus the maximum of the horizontal axes. Given the Markov process of the game, there exists a single closed communicating class consisting of all the possible combinations of orders. Therefore, the sum of the heights of all vertical bars in a limiting

[4] Because of order lumpiness, capacity allocations may become fractional number of batches. We round a fractional number to the nearest two integer numbers with weighted probabilities such that the expectation is equal to the fractional number (see the Online Appendix for details).

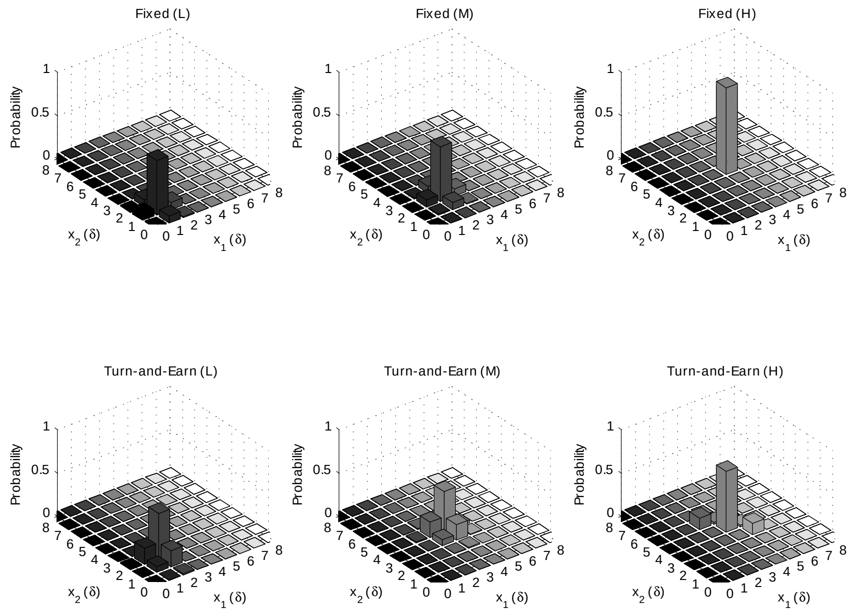

Figure 1: Conditional Limiting Distributions of Orders. $K = 8\delta$, $\mathcal{A} = \{5, 9, 21\}$, $\rho = 0.9$, and $\varepsilon = 4$.

distribution graph add up to 1.

Despite dispersion of local demand, the modal order quantities remain unchanged under fixed allocation: $(\delta, 2\delta, 4\delta)$. However, they change to $(2\delta, 4\delta, 4\delta)$ under turn-and-earn, as shown in Figure 1. The differences in the order quantities in the low and medium demand states again suggest that turn-and-earn induces the retailers to increase orders and thus sales, consistent with the findings of Cachon & Lariviere (1999*a*). However, our equilibria exhibit a novel feature – two ex ante identical retailers may end up with asymmetric allocations. Idiosyncratic local demand shocks embedded in our model give rise to the asymmetry. An important implication of this property is the natural occurrence of sales leaders and laggards, which leads to richer strategic dynamics than those in Cachon & Lariviere (1999*a*).

Because it induces retailers to increase sales beyond their myopic level, turn-and-earn leads to higher total sales than the Pareto and fixed allocation. We also know that the Pareto mechanism leads to less total sales than the first best because of double marginalization. Then, how does turn-and-earn compare to the first best allocation? Who benefits from turn-and-earn? To answer these questions, we take the example used in Figure 1 and calculate expected sales and profits under first best, Pareto, fixed, and turn-and-earn allocation. The results are summarized in Table 2. In this example, turn-and-earn induces higher sales than the first best and the highest sales increase occurs

	Expected Total Sales		Supplier's Profit		Retailers' Profits		System Profits	
	M	Ave	M	Ave	M	Ave	M	Ave
First Best	4.50	7.59	–	–	–	–	63.36	194.53
Pareto	4.00	7.54	8.00	15.08	54.86	179.40	62.86	194.22
Fixed	4.00	7.54	8.00	15.08	30.34	173.97	38.34	189.05
Turn-and-Earn	7.44	7.89	14.88	15.79	5.85	169.51	20.73	185.30

Table 2: Expected Sales (δ) and Profits. $K = 8\delta, \mathcal{A} = \{5, 9, 21\}, \rho = 0.9$, and $\varepsilon = 4$. Note: Profits Are Normalized to Per Period.

in the medium demand state. Obviously, the supplier's profit is higher under turn-and-earn than under the Pareto and fixed allocation for a given wholesale price. However, turn-and-earn leads to the lowest total retailer profits and system profits. This suggests that when the wholesale price is fixed, turn-and-earn alleviates the double marginalization problem by incentivizing retailers to increase sales, but may induce more sales than optimal for the system.

The retailers sell more in the low and medium demand states in order to gain an allocation advantage should demand go up and a capacity shortage occur in the future. Indeed, because in this example $\alpha_3 = 21$, the expected myopic sales quantity is 5δ for both retailers, leading to a capacity shortage in the high demand state given that the supplier's capacity is 8δ. It appears that both capacity tightness and demand characteristics may affect this incentive. To systematically explore this notion, we vary four key parameters of the model: K, α_2, α_3, and ρ.

Figure 2 plots the expected total sales in steady state as a function of K parameterized by ρ. The myopic sales quantities are $(\delta, 2\delta, 5\delta)$ in the three common demand states, which implies that (with no local demand dispersion) below 2δ capacity is always tight and above 10δ capacity is always ample. As a result, turn-and-earn's sales-inducing incentive effect is noticeable only when capacity is mildly tight, i.e., tight for some higher demand states and ample for some lower demand states, as seen in the top-left panel of the figure. This result extends Cachon & Lariviere (1999a)'s two-period and two-demand-state result to a setting of multiple periods and multiple demand states.

Conditional on the low and medium demand states, the top-right and bottom-left panels of the figure further illustrate that the expected total sales under turn-and-earn are increasing in ρ, suggesting an increase in turn-and-earn's incentive effect. In contrast, for fixed allocation, because retailers sell their myopic quantities, the conditional expected sales do not depend on ρ. In the high demand state (the bottom-right panel), capacity is always tight and thus making the incentive

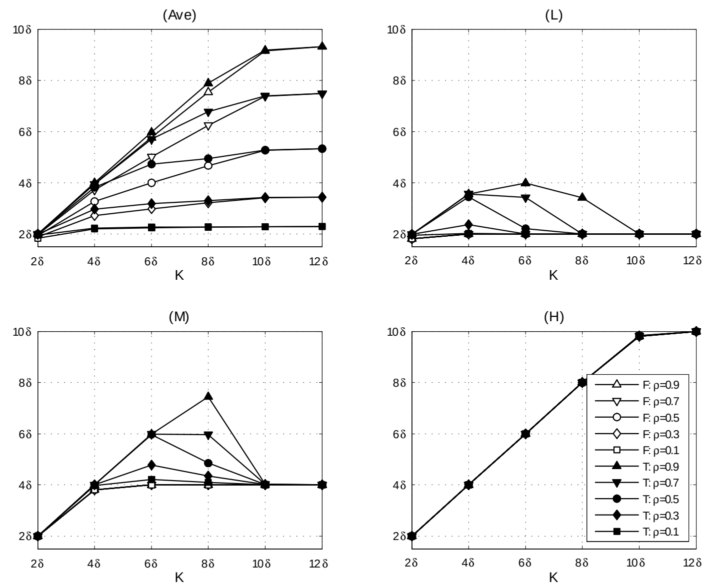

Figure 2: Expected Total Sales as a Function of K Parameterized by ρ Under Fixed (F) and Turn-and-Earn (T) Allocation: $\mathcal{A} = \{5, 9, 21\}$, $\varepsilon = 4$.

effect of turn-and-earn moot. The key to turn-and-earn's incentives is the retailers' forward-looking behavior. They exchange today's loss of profits for tomorrow's larger allotment of capacity. The sacrifice made by the retailers today can only be recouped when future demand is higher. When ρ increases, the demand process becomes first-order stochastically larger, and thus the probability of experiencing higher demand in the future is raised, making over-selling today more profitable. Figure 2 also shows that when averaged over all demand states, $\rho = 0.5$ and 0.7 give the most significant increase in sales under turn-and-earn when compared to those under fixed allocation. These results are not contradictory – larger ρ means capacity is more likely to be tight under either allocation scheme, making turn-and-earn's effect less obvious when averaged across all demand states.

While increasing ρ makes higher demand states more likely to occur, increasing α_2 or α_3 raises the demand levels of higher demand states. Both make a larger capacity allocation more attractive to obtain when excess capacity exists because it entails a potentially larger profit should demand increase and capacity become constrained in the future. Indeed, compared to fixed allocation, the increase in retailers' sales in the low demand state under turn-and-earn is increasing in α_2, as shown in the upper right panel of Figure 3. Similarly, the increase in retailers' sales in the low and medium demand states is increasing in α_3 (see Figure 9 in the Online Appendix).

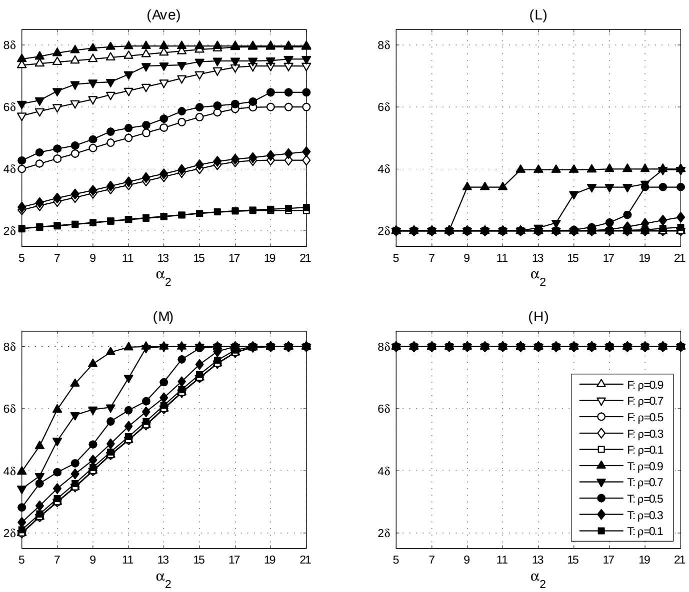

Figure 3: Expected Total Sales as a Function of α_2 Parameterized by ρ Under Fixed (F) and Turn-and-Earn (T) Allocation: $K = 8\delta$, $\mathcal{A} = \{5, \alpha_2, 21\}$, $\varepsilon = 4$.

5.3 Dynamics of Sales Leadership

An important objective of our paper is to study the impact of longer horizon on the competitive ordering behavior of retailers. How would the retailers behave differently in a low-demand period depending on their past sales? Would a sales leader defend his allocation by ordering more than his myopic quantity? Would a sales laggard opportunistically take advantage of excess capacity to catch up? By characterizing the transient dynamics of the MPE, we are able to shed some light on these questions.

We start the game from an asymmetric sales position with one retailer having larger past sales than the other, $s_1^0 = 3\delta$ and $s_2^0 = \delta$, i.e., the leader has a two-batch sales lead over the laggard.[5] The initial common demand state is medium, i.e., $\alpha^0 = \alpha_2$. Our results show that whether the leader would defend his initial leadership in a lower demand state (i.e., medium or low) critically depends on two parameters: demand transition probability and demand levels. Table 3 displays the retailers' expected sales after 1 and 25 periods under different values of ρ and α_3.

Given $\alpha_2 = 13$, the myopic sales quantity is 3 batches for the medium demand state. Table

[5]It is worth noting that only the sales gap, i.e., $s_1^0 - s_2^0$, matters to the order dynamics under turn-and-earn allocation. Therefore, our following results apply to any initial sales states with retailer 1 being 2 batches in lead.

	$\rho = 0.7,\ \alpha_3 = 37$	$\rho = 0.7,\ \alpha_3 = 21$	$\rho = 0.1,\ \alpha_3 = 37$
$T = 1$	$(5\delta, 3\delta)$	$(4\delta, 4\delta)$	$(3\delta, 3\delta)$
$T = 25$	$(5\delta, 3\delta)$, $(4\delta, 4\delta)$, $(2\delta, 2\delta)$	$(4\delta, 4\delta)$, (δ, δ)	$(4\delta, 4\delta)$, $(3\delta, 3\delta)$, (δ, δ)

Table 3: Modal Sales States (s_1, s_2) under Turn-and-Earn after T Periods When the Game Starts from $(\mathbf{s}^0, \alpha^0) = (3\delta, \delta, \alpha_2)$. $K = 8\delta, \mathcal{A} = \{5, 13, \alpha_3\}, \varepsilon = 0.01$.

3 shows that the leader may aggressively defend his leadership by increasing current sales beyond his myopic level to 5δ when $\rho = 0.7, \alpha_3 = 37, T = 1$. This indicates that at the medium demand state the sales leader decides to sell 2 batches more than his myopic quantity. Under the capacity constraint of 8 batches, the leader effectively sustains his leadership by limiting his opponent's sales to 3 batches. On the other hand, the laggard may also increase sales to catch up if there exists excess capacity. In the case of $\rho = 0.7, \alpha_3 = 21, T = 1$, at the medium demand state both the leader and the laggard sell 4 batches. In this case, the strategic benefit of aggressive defence for the leader is much reduced as α_3 decreases from 37 to 21. Nonetheless he still sells 1 batch more than his myopic level to prevent the laggard from gaining an advantage. As a result, both retailers end up on equal footing. When the chance of going into a low demand state is very high (e.g., $\rho = 0.1, \alpha_3 = 37, T = 1$), both the leader and the laggard would give up completely and sell their myopic quantity.

Though the game starts from the medium demand state, after one period there is a positive probability with which the low demand state will be reached, for which the myopic sales quantity is δ. Thus, too much excess capacity exists, implying that leadership is too costly to sustain. Therefore, once the low demand state is reached, any prior allocation asymmetry would be eliminated. When $\rho = 0.7, \alpha_3 = 37$, $(2\delta, 2\delta)$ becomes one of the modal sales states after 25 periods, indicating the equilibrium sales in the low demand state. It is interesting to note that this is still one batch higher than the myopic level for each retailer. This depicts the "fighting" between the two retailers: despite the low current demand, it is worthwhile to fight for sales leadership as long as future demand is sufficiently attractive ($\alpha_3 = 37$ as opposed to 21) and is sufficiently likely to occur ($\rho = 0.7$ as opposed to 0.1). In these two contrasting cases, (δ, δ) is the sales in the low demand state. Nevertheless, no retailer is able to get ahead because of too much excess capacity and thus the supplier is the sole winner of the fight.

It is worth noting that in the absence of the low demand state, sales leadership can be retained permanently for the example of ($\rho = 0.7, \alpha_3 = 37$) shown earlier. Because the leader's equilibrium

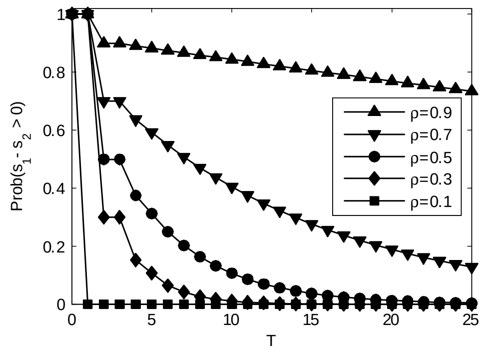

Figure 4: Probability of $s_1 - s_2 > 0$ after T Periods When the Game Starts from $(\mathbf{s}^0, \alpha^0) = (3\delta, \delta, \alpha_2)$. $K = 8\delta$, $\mathcal{A} = \{5, 13, 37\}$, $\varepsilon = 0.01$.

sales quantity is 5 batches in both the high and medium demand states, without the low demand state, the leader would be able to hold out to his allocation advantage forever. Hence, the possibility of extremely low demand in the future prevents sales leadership from becoming permanent in the long run. This *semi-defensible* leadership is a distinct feature of our model with multiple demand states. In contrast, a two-state demand model, as has been used in Cachon & Lariviere (1999*a*), yields only defensible or indefensible leadership. How long would this semi-defensible sales leadership persist depends on demand transition probabilities. To show this, we plot the probability of leadership persistence over time in Figure 4 using transient distributions of equilibrium sales. As ρ increases, i.e., as the demand process becomes first-order stochastically larger, the probability of preserving leadership is higher over time.

In summary, α_3 and ρ together determine the value of leadership. The magnitude of high demand affects the future profits of a larger capacity allotment, while demand transition probabilities decide how long the leader can hold out in the medium demand state and come back to the high demand state. These findings imply that in general the retailers find it advantageous to maintain sales leadership for a short run. A laggard may be punished for a while, but sales leadership is too costly to maintain in the long run when extremely low demand is likely to occur in the future. In other words, an economic recession may act as a sales equalizer that eventually re-levels the playing field, as in the Toyota example discussed in the Introduction.

	Fixed				Turn-and-Earn				Change (%)			
ρ	L	M	H	Ave	L	M	H	Ave	L	M	H	Ave
0.1	0.322	0.107	0	0.599	0.321	0.101	0	0.606	−0.2	−6.0	0	1.1
0.3	0.322	0.107	0	0.636	0.322	0.091	0	0.643	0	−15.4	0	1.2
0.5	0.322	0.107	0	0.478	0.322	0.102	0	0.473	0	−4.8	0	−1.0
0.7	0.322	0.107	0	0.289	0.278	0	0	0.257	−13.6	−100.0	0	−11.0
0.9	0.322	0.107	0	0.113	0.276	0.002	0	0.055	−14.4	−98.5	0	−51.6

Table 4: Coefficients of Variation of Total Sales. $K = 8\delta, \mathcal{A} = \{5, 13, 21\}, \rho = 0.5$.

5.4 Sales Variability

We have shown that turn-and-earn induces the retailers to increase sales. A related question is how does it affect the variability of sales. With significant local demand shocks, retailers may in equilibrium order and thus sell different amounts for the same common demand state. To systematically compare sales variability, we compute the coefficient of variation (CV) of expected total sales in steady state for a range of ρ values and summarize in Table 4.

Conditional on a common demand state, Table 4 shows that turn-and-earn consistently leads to lower sales variability. To explain this, we plot the histogram of total sales in Figure 5 for the case of $\rho = 0.7$ from Table 4. The left plot illustrates that total sales spread symmetrically around 6 batches under fixed allocation. In contrast, the right plot shows that the distribution of sales shifts to the right and concentrates at 8 batches, the supplier's capacity cap, thereby reducing the variability. This suggests that turn-and-earn's incentive to increase sales seems to have an inherent side effect – the capacity constraint becomes more likely to bind and sales variability is thus reduced.

Averaged over all demand states, Table 4 shows that the effect of turn-and-earn on sales variability is a bit mixed. Two opposite effects take place here. First, turn-and-earn reduces sales variation at each demand state, as explained earlier. Second, as turn-and-earn increases sales for the medium demand state (most strongly), sales are less evenly distributed across demand states. For the two cases where turn-and-earn increases average sales variability ($\rho = 0.1$ and 0.3), the second effect apparently dominates.

It is worth commenting on the fact that average CVs do not always decrease under turn-and-earn while conditional CVs do in our examples. For a supply chain where the common demand state is

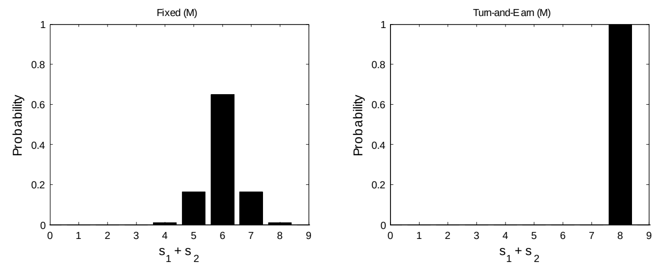

Figure 5: Limiting Distribution of Total Sales (δ) Conditional on the Medium Demand State under Fixed and Turn-and-Earn Allocation. $K = 8\delta$, $\mathcal{A} = \{5, 13, 21\}$, $\rho = 0.7$, $\varepsilon = 4$.

public knowledge, the state-dependent variability-reducing effect of turn-and-earn is more relevant. Such situations are conceivable because common demand factors such as macroeconomic conditions and product characteristics may uniformly affect all retailers while each individual retailer may experience idiosyncratic local demand shocks.

6 Extensions

We consider four extensions in this section to achieve two objectives: First, how do retailers behave differently when the model context changes, e.g., in the presence of more rival retailers (Section 6.1), or in a life cycle model (Section 6.2)? Second, are the main insights robust to modifications in the definition of turn-and-earn? Section 6.3 explores the impact of basing turn-and-earn on the proportional rather than the absolute difference of past sales, while Section 6.4 studies a turn-and-earn scheme using the moving average of past sales. We continue to use the same set of parameter values for the numeric examples unless otherwise noted.

6.1 N Retailers

Our model extends to N retailers in a straightforward way. In this section, we briefly explain how the analytics can be modified to incorporate N retailers and relegate the details to the Online Appendix. The Bellman equation of retailer i becomes

$$
\begin{aligned}
V_i(\mathbf{s}, \alpha, \theta_i) \quad \equiv \quad & \max_{x_i \in \{0\delta, \ldots, \bar{m}\delta\}} \mathbb{E}\left[\pi_i(\mathbf{s}, \alpha, \theta_i, x_i(\mathbf{s}, \alpha, \theta_i), \mathbf{x}_{-i}(\mathbf{s}, \alpha, \cdot))\right] \\
& + \beta \mathbb{E}\left[V_i(\mathbf{s}', \alpha') | \mathbf{s}, \alpha, x_i(\mathbf{s}, \alpha, \theta_i), \mathbf{x}_{-i}(\mathbf{s}, \alpha, \cdot)\right],
\end{aligned}
$$

where $\mathbf{x}_{-i}(\mathbf{s}, \alpha, \cdot) \equiv (x_1(\mathbf{s}, \alpha, \theta_1), ..., x_{i-1}(\mathbf{s}, \alpha, \theta_{i-1}), x_{i+1}(\mathbf{s}, \alpha, \theta_{i+1}), ..., x_N(\mathbf{s}, \alpha, \theta_N))$.

We need the following algorithm to implement the turn-and-earn allocation mechanism. Let r_i be retailer i's reserved capacity and g_i his guaranteed allocation. Index the N retailers in *increasing* order of sales. Given sales $\mathbf{s}$ and orders $\mathbf{x}$, allocation $\mathbf{q}$ is determined iteratively:

Step 0. $\hat{\mathbf{s}} = \mathbf{s},\ \hat{K} = K,\ \hat{N} = N,\ \mathbf{q} = Null.$

Step 1. For $i = 1, ..., \hat{N},$ let $r_i = \min\left\{\hat{s}_i - \hat{s}_1, \left(\hat{K} - \sum_{j>i}^{\hat{N}}(\hat{s}_j - \hat{s}_1)\right)^+\right\},$
and $g_i = r_i + \frac{1}{\hat{N}}\left(\hat{K} - \sum_{j=1}^{\hat{N}} r_j\right);$ If $x_i \leq g_i,$ then $q_i = x_i,\ \hat{K} = \hat{K} - x_i.$

Step 2. If $x_i > g_i$ for all i, then $q_i = g_i$ for $i = 1, ..., \hat{N}$ and stop. Else, go to Step 3.

Step 3. Keep only the retailers with $Null$ allocation. Index them in increasing order of sales, and update $\hat{\mathbf{s}}$ and $\hat{N}$ accordingly.

Step 4. If $\hat{N} = 0,$ stop. If $\hat{N} = 1,$ $q_1 = \min\left\{x_1, \hat{K}\right\}$ and stop. Else, go to Step 1.

This algorithm extends the turn-and-earn mechanism defined by Cachon & Lariviere (1999*a*) (and used in §3 earlier) to the N retailer setting: when $N = 2$, r_i becomes $|s_i - s_j|$, thus identical to the definition used in that paper. Because a given retailer may order less than his guaranteed allocation, the iterative procedure is necessary to reallocate unclaimed capacity to those retailers who order more than their guaranteed allocation. (Similar issues arise with other allocation schemes. For example, see the discussion of uniform allocation in Cachon & Lariviere (1999*b*).)

The equilibrium properties characterized by Proposition 1 applies to the case of N retailers (see the Online Appendix for the modified proof). In the following, we numerically solve for the MPE for the case of $N = 3$. An immediate question is does more retailers imply intensified competition for the limited capacity of the supplier? To make the equilibrium results comparable to the case of two retailers, we set the supplier's capacity proportional to the number of retailers. Table 5 lists the expected sales per retailer with different values of ρ under turn-and-earn. A noticeable increase in per-retailer sales can be observed across all demand states, but it is not uniformly distributed. The medium demand state, where we have seen the strongest incentive effect of turn-and-earn in the two-retailer case, claims the highest increase in sales per retailer.

It is worth noting that solving the multi-retailer game imposes a substantial computational burden. Even for the three-retailer case, when $M = 6$, as used for the results in Table 5, the dimensionality of the transition matrix becomes 1029×1029. For a comparable four-retailer setting ($M = 8$), the dimensionality increases to 19683×19683. Therefore, we are computationally constrained to explore beyond three retailers. But the consistency in the results of the three-retailer case indicates that our main results would likely to extend to the setting of more retailers. For future

	2 Retailers				3 Retailers				Change (%)			
ρ	L	M	H	Ave	L	M	H	Ave	L	M	H	Ave
0.1	1	1.303	2	1.041	1	1.319	2	1.043	0	1.23	0	0.19
0.3	1	1.437	2	1.23	1	1.515	2	1.251	0	5.43	0	1.71
0.5	1	1.794	2	1.598	1.001	1.896	2	1.633	0.10	5.69	0	2.19
0.7	1.112	1.900	2	1.872	1.119	1.914	2	1.877	0.63	0.74	0	0.27
0.9	1.711	1.948	2	1.992	1.711	1.969	2	1.994	0	1.08	0	0.10

Table 5: Expected Sales (δ) per Retailer: $\mathcal{A} = \{5, 6, 13\}, K = 4\delta$ (2 retailers), $K = 6\delta$ (3 retailers), $\varepsilon = 4$.

research, recent development in the field of computational economics offer promising opportunities for studying dynamic games with a large number of firms, e.g. Pakes & McGuire (2001) and **?**.

6.2 Life Cycle

Automobiles, like many other products, have a life cycle with successive generations of the same product. This affects how turn-and-earn is implemented in practice. In the Pontiac example, the manufacturer tied the capacity allocation of a new model, Sunfire, to the sales of Sunbird, an old model. To proxy for a life cycle, we consider a demand process illustrated in Figure 6. We assume that the demand intercept, α_2, is identical in state 2^O and 2^N. Thus the difference in those two medium demand states is solely about whether it is likely to stay in a high demand environment or a whether it is likely to be stuck in a low demand environment. For fixed allocation or the first best, this would not matter – the retailers' ordering policies would not change from what we currently have. However this will matter for turn-and-earn. If p_{ON} and p_{NO} are sufficiently low. That is, once you slip to having an old product, you will be stuck with it for a while.

We consider a demand transition matrix such that $p_{1O} = p_{ON} = p_{N3} = p_{33} = r$ and $p_{11} = p_{O1} = p_{NO} = p_{3N} = 1 - r$, where $r \in (0, 1)$. This transition matrix ensures that as r increases the associated demand process becomes first-order stochastically larger. Further, to compare the equilibrium results of the life cycle model to the base model, we *calibrate* their demand processes to yield identical limiting distributions of α by choosing appropriate r, ρ_1, and ρ_2. Table 6 shows the expected total sales of five pairs of such demand processes. For each row in the table, despite identical distributions of demand levels, the retailers' incentives to oversell beyond the myopic level in the medium demand state are quite different for the two models. For the life cycle model, selling

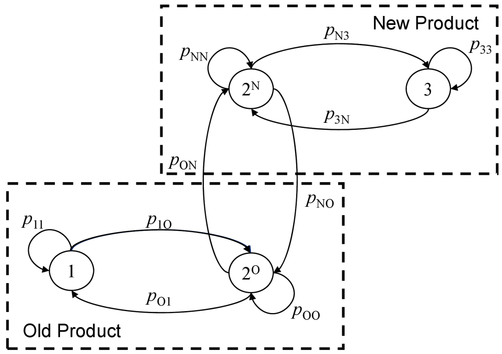

Figure 6: Demand Process for the Life Cycle Model

	Life Cycle							Base Model			
r	L	M(Old)	M(New)	H	Ave	ρ_1	ρ_2	L	M	H	Ave
0.1	2	6	6.102	8	2.448	0.110	0.011	2	6.010	8	2.448
0.3	2	6	6.355	8	3.766	0.380	0.114	2	6.117	8	3.770
0.5	2	6	6.853	8	5.713	0.667	0.333	2	6.405	8	5.702
0.7	2	6.100	8	8	7.514	0.886	0.620	2.004	7.812	8	7.653
0.9	3.968	7.999	8	8	7.995	0.989	0.890	3.968	8	8	7.995

Table 6: Expected Total Sales (δ). $\mathcal{A} = \{5, 13, 21\}, K = 8\delta, \varepsilon = 4$.

a new product in the medium demand state (i.e., 2^N) encourages the retailers to sell a lot more than selling an old product (i.e., in state 2^O). The difference can go as high as 30% when $r = 0.7$. In the medium demand state of the base model, the retailers' sales quantities are somewhere in between those of state 2^N and 2^O in the life cycle model. Averaged across all demand states, the expected total sales are fairly similar across the two models (with 2% difference at most).

Taken together, these results suggest that the retailers can take advantage of additional demand information in determining how they adjust their sales rates. Here, knowing that the product is old (or new) allows the retailers to scale back (or increase) their sales rates. Whether this benefits the retailers or suppliers depends on the specific problem. Here moderate values of r (e.g., 0.5) result in the supplier seeing higher sales as the retailers significantly increase sales of a new product in the medium demand state. For higher values of r (e.g., 0.7), new product sales are capped by capacity so the reduction in sales of old products dominates and the supplier sees lower total sales.

6.3 Proportional Turn-and-Earn

Under the turn-and-earn scheme defined earlier, retailers' guaranteed allocation depends on the absolute difference in their past sales. Alternatively, one could implement a scheme in which the guaranteed allocation depends on the relative difference in past sales. This leads to proportional turn-and-earn. Under this mechanism, the supplier reserves $(s_i/\hat{s})K$ for retailer i and thus $r_i = g_i = (s_i/\hat{s})K$, where $\hat{s} = \sum_{j=1}^{N} s_j$. (Note that this differs from the proportional scheme considered by Lee et al. (1997) which is based on orders as opposed to past sales.) The equilibrium structure under proportional turn-and-earn can be characterized by the results in Proposition 1. Because proportional turn-and-earn is based on past sales, intuition suggests that it would also have a similar effect as turn-and-earn. Our numerical study shows that both mechanisms induce similar expected total sales (see Figure 10 in the Online Appendix).

While similar in expected total sales, the distribution of sales among the retailers can be quite different under the two mechanisms. Figure 7 illustrates an example where proportional turn-and-earn induces more asymmetry in sales than turn-and-earn. The limiting sales distribution indicates that under turn-and-earn the likelihood of $(4\delta, 4\delta)$, a symmetric state, is much higher than those asymmetric states $(3\delta, 5\delta)$ and $(5\delta, 3\delta)$. Under proportional turn-and-earn, the likelihood of those three states are much similar, indicating an *increase* in the likelihood of the asymmetric states. This observation is consistent across many experiments. The intuition behind the result goes back to how reserved capacity is calculated under the two mechanisms. Recall that under turn-and-earn the difference in guaranteed allocation of two retailers is $|s_1 - s_2|$. This is identical to that of proportional turn-and-earn when $s_1 + s_2 = K$. However, when $s_1 + s_2 < K$, the difference in guaranteed allocation under proportional turn-and-earn is $K|s_1 - s_2|/(s_1 + s_2)$ and thus is greater than $|s_1 - s_2|$, thereby creating a larger disparity in capacity allocation.

6.4 Moving Average Turn-and-Earn

Our base model considers a turn-and-earn scheme using sales data of the immediate past period. To factor in a longer history of past sales, we construct an alternative scheme using the moving average of past sales. Let $s_i^{(-t)}$ represent retailer i's sales quantity t periods ago. A moving average of order L is $\bar{s}_i = \frac{1}{L}\sum_{t=1}^{L} s_i^{(-t)}$. Turn-and-earn can be redefined by replacing s_i with $\bar{s}_i$. The state of the world becomes $(\mathbf{s}^{(-1)}, ..., \mathbf{s}^{(-t)}, \alpha)$. We set $L = 2$, although larger order of moving average can be computed at the expense of a considerable computational burden. Table 7 shows the differences in

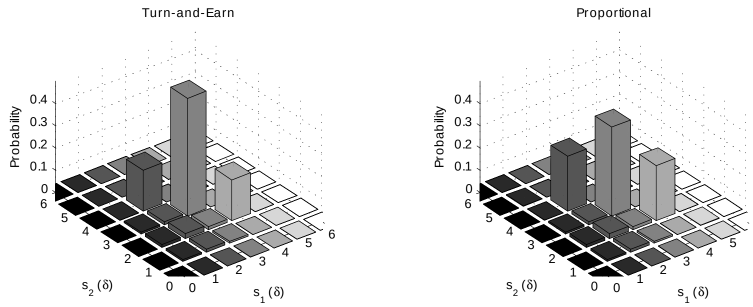

Figure 7: Limiting Distribution of Sales. $K = 6\delta$, $\mathcal{A} = \{5, 13, 21\}$, $\rho = 0.7$, $\varepsilon = 4$.

ρ_1	ρ_2	Fixed Allocation	Turn-and-Earn	Moving Average Turn-and-Earn	Change (%)
0.5	0.1	2.351	2.393	2.363	−1.23
0.5	0.3	2.566	2.713	2.587	−4.64
0.5	0.5	2.840	3.196	2.840	−11.13
0.7	0.5	3.037	3.478	3.038	−12.67
0.9	0.5	3.193	3.700	3.193	−13.70

Table 7: Expected Total Sales (δ): $K = 4\delta, \mathcal{A} = \{5, 6, 13\}, \varepsilon = 4$.

expected total sales for the two versions of turn-and-earn. Overall, moving average turn-and-earn induces lower sales than turn-and-earn. Moreover, as ρ_1 and/or ρ_2 increases, i.e., as the demand process becomes first-order stochastically larger, the reduction of sales increases. Because moving average turn-and-earn makes capacity allocation stickier, a sales leader incurs less cost to defend his leadership while a laggard incurs more cost to catch up. Hence, using moving average of past sales *dampens* the incentive effect of turn-and-earn.

7 Conclusion

We have constructed a dynamic stochastic game of order competition to study the incentive effect of turn-and-earn, a commonly used capacity allocation scheme in the automobile industry. We apply the Markov perfect equilibrium framework developed in the field of computational economics to characterize the strategic ordering behavior of retailers. Earlier results in the literature, specifically by Cachon & Lariviere (1999*a*), are limited to a two-period, two-retailer, and two-demand-state setting with identical retailer demand. Using computed MPE, we show that turn-and-earn remains effective in inducing more sales in a setting with long horizon, multiple retailers, multiple demand states, and private local demand shocks. Further, sales leadership in general is defensible in the short run under certain conditions, but indefensible in the long run when extremely low demand occurs. Finally, we find that turn-and-earn induces the retailers to absorb local demand shocks and thus reduces the supplier's sales variability. Although we have made some simplifying assumptions, such as state discretization and lumpy orders, and our results are numerical, we have found a rich set of equilibrium characteristics that are not captured by earlier analytical results derived from a simpler model.

An extension worth considering for future work is retailer competition for customers. When retailers compete in the same market, the incentive for earning a larger allocation would be enhanced because a retailer gains a larger market share if his competitor has lower inventory. Another interesting extension is to introduce asymmetry in demand, for example, one retailer has a larger demand intercept than the other for a given common demand state. Such asymmetry in market size would lead to different marginal losses in current profits due to increasing sales beyond the myopic level, and thus affect retailer ordering behavior. Neither of these extensions have been considered in the two-period setting. Incorporating them may hinder analytical tractability and even complicate computational study.

References

Associated Press (1997), 'Shortage has P&G allocating detergent', *Associated Press Newswire* **August 6**.

Besanko, D., Doraszelski, U., Lu, L. X. & Satterthwaite, M. (2010), 'Lumpy capacity investment and disinvestment dynamics', *Operations Research* **58**(4), 1178–1193.

Cachon, G. P. & Lariviere, M. A. (1999*a*), 'Capacity allocation using past sales: When to turn-and-earn', *Management Science* **45**(5), 685–703.

Cachon, G. P. & Lariviere, M. A. (1999*b*), 'Capacity choice and allocation: Strategic behavior and supply chain performance', *Management Science* **45**(8), 1091–1108.

Cachon, G. P. & Zhang, F. (2007), 'Obtaining fast service in a queueing system via performance-based allocation of demand', *Management Science* **53**(3), 408–420.

Chen, F., Li, J. & Zhang, H. (2007), Retail competition, capacity allocation, and supply performance. Working paper, Columbia University.

Child, C. (1995), 'Pontiac wants to boost Sunfire sales in California', *Automotive News* **October 19**.

Chod, J. & Rudi, N. (2005), 'Resource flexibility with responsive pricing', *Operations Research* **53**(3), 532–548.

Doraszelski, U. & Satterthwaite, M. (2010), 'Computable Markov-perfect industry dynamics', *Rand Journal of Economics* **41**(2), 215–243.

Ericson, R. & Pakes, A. (1995), 'Markov-perfect industry dynamics: A framework for empirical work', *Review of Economic Studies* **62**, 53–82.

Fudenberg, D. & Tirole, J. (1991), *Game theory*, MIT Press, Cambridge.

Furuhata, M., Perrussel, L. & Zhang, D. (2007), Mechanism design for capacity allocation under price competition. Working paper, University of Western Sydney.

Furuhata, M. & Zhang, D. (2006), 'Capacity allocation with competitive retailers', *Proceedings of the 8th Interantional Conference on Electronic Commerce* pp. 31–37.

Gilbert, S. M. & Weng, Z. K. (1998), 'Incentive effects favor nonconsolidating queues in a service system: The principal-agent perspective', *Management Science* **44**(12), 1662–1669.

Hwang, S. & Valeriano, L. (1992), 'Marketers and consumers get the jitters over servere shortages of nicotine patches', *Wall Street Journal* **May 22**, B–1.

LaReau, J. (2008), 'GM to increase Enclave output', *Automotive News* **January 21**.

Lavin, D. (1994), 'Youwannadeal? Bucking Detroit trend, Landmark Chevrolet still uses the hard sell', *Wall Street Journal,* July 8 .

Lawrence, D. (1996), 'GM test allocation change: More makers scrap turn-and-earn plan', *Automotive News* **February 12**.

Lee, H. L., Padmanabhan, V. & Whang, S. (1997), 'Information distortion in a supply chain: The bullwhip effect', *Management Science* **43**(4), 546–558.

Lynch, S. (1997), *Arrogance and Accords*, Pecos Press, Dallas, Texas.

Maskin, E. & Tirole, J. (2001), ‘Markov perfect equilibrium, I: Observable actions’, *Journal of Economic Theory* **100**, 191–219.

Mookherjee, R. & Friesz, T. L. (2008), ‘Pricing, allocation, and overbooking in dynamic service network competition when demand is uncertain’, *Production and Operations Management* **17**(4), 455–474.

Muller, A. & Stoyan, D. (2002), *Comparison Methods for Stochastic Models and Risks*, John Wiley & Sons, England.

Pakes, A. & McGuire, P. (1994), ‘Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model’, *Rand Journal of Economics* **25**(4), 555–589.

Pakes, A. & McGuire, P. (2001), ‘Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the ’Curse’ of Dimensionality’, *Econometrica* **69**(5), 1261–1281.

Perakis, G. & Sood, A. (2006), ‘Competitive multi-period pricing for perishable products: A robust optimization approach’, *Mathematical Programming* **107**, 295–335.

Purohit, D. & Vernik, D. (2008), Turn-and-earn in a product line: The impact of product substitutability. Working paper, Duke University.

Rechtin, M. (2008), ‘Toyota division inventories soar to record high’, *Automotive News* **April 28**.

Sawyers, A. (1999), ‘Turn, Earn: a slow burn’, *Automotive News* **October 19**, 34.

Shirouzu, N. (2006), ‘Foreign models: In Chinese market, Toyota’s strategy is made in U.S.A.’, *Wall Street Journal* **May 26**.

Weintraub, G. Y., Benkard, C. L. & Roy, B. V. (2010), ‘Markov perfect industry dynamics with many firms’, *Econometrica* **76**(6), 1375–1411.

Zarley, C. & Damore, K. (1996), ‘Backlogs plague HP: Resellers place phantom orders to get more products’, *Computer Reseller News* **May 6**, 247.

Zipkin, P. H. (2000), *Foundations of Inventory Management*, McGraw-Hill, USA.

Online Appendix to Capacity Allocation over a Long Horizon: The Return on Turn-and-Earn

I. Proofs

Proposition 1 Proof. We develop the proof through three lemmas.

Lemma 1 *Suppose that $\mathbf{q}(\mathbf{s},\mathbf{x})$ is weakly increasing in $\mathbf{x}$. For any $m \in \{0, ..., M\}$, it is optimal for retailer i to order $m\delta$ if and only if $\theta_i \geq \underline{\theta}_{im}(\mathbf{s},\alpha)$ and $\theta_i \leq \bar{\theta}_{im}(\mathbf{s},\alpha)$, where*

$$\underline{\theta}_{im}(\mathbf{s},\alpha) = \begin{cases} \underline{\Theta} & \text{if } m = 0 \\ \max\left\{\frac{bm\delta-\alpha}{\varepsilon}, \max\limits_{m' \in \{0,...,m-1\}} g(m, m', \mathbf{s}, \alpha)\right\} & \text{if } m = 1, ..., M \end{cases}$$

and

$$\bar{\theta}_{im}(\mathbf{s},\alpha) = \begin{cases} \bar{\Theta} & \text{if } m = M \\ \min\limits_{m' \in \{m+1,...,M\}} \left\{\max\left\{\frac{bm'\delta-\alpha}{\varepsilon}, g(m, m', \mathbf{s}, \alpha)\right\}\right\} & \text{if } m = 0, ..., M-1 \end{cases}$$

The expression for $g(\cdot)$ is given by Eq.(7) and (8). Accordingly his optimal order distribution $\boldsymbol{\sigma}_i(\mathbf{s},\alpha)$ is given by

$$\sigma_{im}(\mathbf{s},\alpha) = \int \mathbf{1}[\underline{\theta}_{im}(\mathbf{s},\alpha) \leq \theta_i \leq \bar{\theta}_{im}(\mathbf{s},\alpha)] dF(\theta_i), \ m = 0, ..., M.$$

Proof. Consider the optimization problem of retailer i:

$$V_i(\mathbf{s},\alpha,\theta_i) = \max_{x_i(s,\alpha,\theta_i)\in\{0\delta,...,\bar{m}\delta\}} \sum_{l=0}^{M} \sigma_{jl}(\mathbf{s},\alpha) \left\{\pi_i(\mathbf{s},\alpha,\theta_i,x_i,l\delta) + \beta \sum_{\alpha'\in\mathcal{A}} P(\alpha'|\alpha) V_i\left(\mathbf{q}(\mathbf{s},x_i,l\delta),\alpha'\right)\right\}$$

where

$$\bar{m} \equiv \max\{m \in \mathbb{Z} | \alpha + \varepsilon\theta_i - bm\delta \geq 0\}.$$

For $m = 1, ..., M-1$, $x_i^* = m\delta$ if and only if $m \leq \bar{m}$ and

$$\begin{aligned} &\sum_{l=0}^{M} \sigma_{jl}(\mathbf{s},\alpha) \left\{\pi_i(\mathbf{s},\alpha,\theta_i,m\delta,l\delta) + \beta \sum_{\alpha'\in\mathcal{A}} P(\alpha'|\alpha) V_i(\mathbf{q}(\mathbf{s},m\delta,l\delta),\alpha')\right\} \\ \geq &\max_{x_i\in\{0\delta,...,M\delta\}} \sum_{l=0}^{M} \sigma_{jl}(\mathbf{s},\alpha) \left\{\pi_i(\mathbf{s},\alpha,\theta_i,x_i,l\delta) + \beta \sum_{\alpha'\in\mathcal{A}} P(\alpha'|\alpha) V_i(\mathbf{q}(\mathbf{s},x_i,l\delta),\alpha')\right\}. \end{aligned}$$

The boundary cases, i.e., $m = 0$ and $m = M$ can be easily derived following the same logic and is omitted for brevity. The second inequality is equivalent to for any $m' \neq m$, $m' \in \{0, ..., m-1, m+1, ..., M\}$,

$$\theta_i \varepsilon B(m, m', \mathbf{s}, \alpha) \geq -\Delta(m, m', \mathbf{s}, \alpha). \tag{6}$$

where

$$\begin{aligned}
\Delta(m, m', \mathbf{s}, \alpha) &= \sum_{l=0}^{M} \sigma_{jl}(\mathbf{s}, \alpha) \Big\{ \pi_i(\mathbf{s}, \alpha, 0, m\delta, l\delta) + \beta \sum_{\alpha' \in \mathcal{A}} P(\alpha'|\alpha) V_i \left(\mathbf{q}(\mathbf{s}, m\delta, l\delta), \alpha' \right) \\
&\quad - \pi_i(\mathbf{s}, \alpha, 0, m'\delta, l\delta) - \beta \sum_{\alpha' \in \mathcal{A}} P(\alpha'|\alpha) V_i \left(\mathbf{q}(\mathbf{s}, m'\delta, l\delta), \alpha' \right) \Big\}, \\
B(m, m', \mathbf{s}, \alpha) &= \sum_{l=0}^{M} \sigma_{jl}(\mathbf{s}, \alpha) (q_i(\mathbf{s}, m\delta, l\delta) - q_i(\mathbf{s}, m'\delta, l\delta)).
\end{aligned}$$

$\Delta(\cdot)$ represents the payoff difference between two arbitrary order quantities, say $m\delta$ and $m'\delta$, when retailer i's private demand shock is set to zero. $B(\cdot)$ represents the difference in the expected allocations of ordering m and m' batches. Next, we characterize the threshold function denoted by $g(\cdot)$. Consider two cases:

Case 1. If $B(m, m', \mathbf{s}, \alpha) \neq 0$, then

$$g(m, m', \mathbf{s}, \alpha) = -\frac{\Delta(m, m', \mathbf{s}, \alpha)}{\varepsilon B(m, m', \mathbf{s}, \alpha)}. \tag{7}$$

Because $q_i(\mathbf{s}, x_i, x_j)$ is weakly increasing in x_i. Therefore when $m > m'$, we have $q_i(m\delta, l\delta, \mathbf{s}) \geq q_i(m'\delta, l\delta, \mathbf{s})$. Thus, Eq.(6) is equivalent to $\theta_i \geq g(m, m', \mathbf{s}, \alpha)$. When $m < m'$, we have $q_i(\mathbf{s}, m\delta, x_j) \leq q_i(\mathbf{s}, m'\delta, x_j)$. Thus, Eq. (6) is equivalent to $\theta_i \leq g(m, m', \mathbf{s}, \alpha)$. Therefore, $g(\cdot)$ represents the threshold of local demand state above (below) which ordering m batches is better than ordering m' batches when m is larger (smaller) than m'.

Case 2. If $B(m, m', \mathbf{s}, \alpha) = 0$, then

$$g(m, m', \mathbf{s}, \alpha) = \begin{cases} \underline{\Theta} & \text{if } m > m' \text{ and } \Delta(m, m', \mathbf{s}, \alpha) > 0, \\ \bar{\Theta} & \text{if } m < m' \text{ and } \Delta(m, m', \mathbf{s}, \alpha) > 0, \\ \bar{\Theta} & \text{if } m > m' \text{ and } \Delta(m, m', \mathbf{s}, \alpha) < 0, \\ \underline{\Theta} & \text{if } m < m' \text{ and } \Delta(m, m', \mathbf{s}, \alpha) < 0, \\ \bar{\Theta} & \text{if } m > m' \text{ and } \Delta(m, m', \mathbf{s}, \alpha) = 0, \\ \bar{\Theta} & \text{if } m < m' \text{ and } \Delta(m, m', \mathbf{s}, \alpha) = 0. \end{cases} \tag{8}$$

When $\Delta(m, m', \mathbf{s}, \alpha) > 0$, it is easy to see that ordering m batches is always better than ordering m' batches over the entire support of θ_i. Therefore, when $m > m'$(or $m < m'$), we set $g(m, m', \mathbf{s}, \alpha) =$

$\underline{\Theta}$ (or $g(m, m', \mathbf{s}, \alpha) = \bar{\Theta}$) so that above (or below) this threshold, ordering m batches always dominates ordering m' batches. Conversely, if $\Delta(m, m', \mathbf{s}, \alpha) < 0$, ordering m' batches is always better than ordering m batches over the entire support of θ_i and we set $g(m, m', \mathbf{s}, \alpha)$ accordingly. If however $\Delta(m, m', \mathbf{s}, \alpha) = 0$, the two order quantities give identical expected payoff. Here we need to apply the assumption that retailer i chooses the lower order quantity.

Combining the thresholds with the constraint $m \leq \bar{m}$, we reach the conditions in Lemma 1. ■

The threshold policy says that there exists a range of θ_i over which it is optimal for retailer i to order m batches. Notice that the thresholds are state dependent. Let

$$\begin{aligned} \underline{\boldsymbol{\theta}}_i(\mathbf{s}, \alpha) &= \{\underline{\theta}_{i0}(\mathbf{s}, \alpha), ..., \underline{\theta}_{iM}(\mathbf{s}, \alpha)\}, \\ \bar{\boldsymbol{\theta}}_i(\mathbf{s}, \alpha) &= \{\bar{\theta}_{i0}(\mathbf{s}, \alpha), ..., \bar{\theta}_{iM}(\mathbf{s}, \alpha)\}. \end{aligned}$$

These two $M + 1$ dimensional vectors completely determine retailer i's optimal order distribution. We now characterize the properties of the thresholds.

Lemma 2 *For any $m < m'$, $m, m' \in \{0, ..., M\}$,*

$$\bar{\theta}_{im}(\mathbf{s}, \alpha) \leq \underline{\theta}_{im'}(\mathbf{s}, \alpha).$$

Proof. By the definition in Lemma 1, $\bar{\theta}_{im}(\mathbf{s}, \alpha) \leq \max\left\{\frac{bm'\delta - \alpha}{\varepsilon}, g(m, m', \mathbf{s}, \alpha)\right\}$ and $\underline{\theta}_{im'}(\mathbf{s}, \alpha) \geq \max\left\{\frac{bm'\delta - \alpha}{\varepsilon}, g(m', m, \mathbf{s}, \alpha)\right\}$. Notice that $g(m, m', \mathbf{s}, \alpha) = g(m', m, \mathbf{s}, \alpha)$ by the definition in Eq. (7) and (8). Therefore, $\bar{\theta}_{im}(\mathbf{s}, \alpha) \leq \underline{\theta}_{im'}(\mathbf{s}, \alpha)$. ■

This monotonicity property of the thresholds implies that when the local demand state is higher, the optimal order quantity increases. The property also ensures that there is no overlapping θ_i regions whenever ordering m and m' batches have positive probability of being optimal (i.e., when $\underline{\Theta} \leq \underline{\theta}_{im} < \bar{\theta}_{im} \leq \bar{\Theta}$ and $\underline{\Theta} \leq \underline{\theta}_{im'} < \bar{\theta}_{im'} \leq \bar{\Theta}$).

Lemma 3 *For any $m < m'$, $m, m' \in \{0, ..., M\}$, if $\underline{\theta}_{im}(\mathbf{s}, \alpha) \leq \bar{\theta}_{im}(\mathbf{s}, \alpha)$ and $\underline{\theta}_{im'}(\mathbf{s}, \alpha) \leq \bar{\theta}_{im'}(\mathbf{s}, \alpha)$, then $\bar{\theta}_{im}(\mathbf{s}, \alpha) = \underline{\theta}_{im'}(\mathbf{s}, \alpha)$ when either of the following conditions is satisfied:*

i) $m' = m + 1$;

ii) $m' > m + 1$ and $\underline{\theta}_{i,m''}(\mathbf{s}, \alpha) > \bar{\theta}_{i,m''}(\mathbf{s}, \alpha)$ for all $m < m'' < m'$.

Proof. By Lemma 2, we have $\bar{\theta}_{im}(\mathbf{s}, \alpha) \leq \underline{\theta}_{im'}(\mathbf{s}, \alpha)$. Thus we only need to show $\bar{\theta}_{im}(\mathbf{s}, \alpha) < \underline{\theta}_{im'}(\mathbf{s}, \alpha)$ does not hold. Suppose not, there exists a $\tilde{\theta} \in (\bar{\theta}_{im}(\mathbf{s}, \alpha), \underline{\theta}_{im'}(\mathbf{s}, \alpha))$. For any such $\tilde{\theta}$, the following two statements hold: 1) ordering m batches is better than ordering z batches for any $z = 0, ..., m-1$ (because $\tilde{\theta} > \bar{\theta}_{im}(\mathbf{s}, \alpha) \geq \underline{\theta}_{im}(\mathbf{s}, \alpha)$); 2) ordering m' batches is better than ordering z batches for any $z = m+1, ..., M$ (because $\tilde{\theta} < \underline{\theta}_{im'}(\mathbf{s}, \alpha) \leq \bar{\theta}_{im'}(\mathbf{s}, \alpha)$). Moreover, because i) $m' = m+1$ or ii) $\underline{\theta}_{im''}(\mathbf{s}, \alpha) > \bar{\theta}_{im''}(\mathbf{s}, \alpha)$ for all $m < m'' < m'$, the optimal order quantity for any $\tilde{\theta} \in (\bar{\theta}_{im}(\mathbf{s}, \alpha), \underline{\theta}_{im'}(\mathbf{s}, \alpha))$ should be from set $\{m\delta, m'\delta\}$. This contradicts the fact that $\tilde{\theta} \notin \left[\underline{\theta}_{im}(\mathbf{s}, \alpha), \bar{\theta}_{im}(\mathbf{s}, \alpha)\right]$ and $\tilde{\theta} \notin \left[\underline{\theta}_{im'}(\mathbf{s}, \alpha), \bar{\theta}_{im'}(\mathbf{s}, \alpha)\right]$, which are the necessary conditions for $m\delta$ and $m'\delta$ to be optimal, respectively, according to Lemma 1. ■

This lemma ensures that the optimal regions of θ_i for the order quantities are connected and ordered increasingly. With the properties stated in Lemma 2 and 3, we can eliminate duplicated thresholds and simplify the threshold policy as: for any $m \in \{0, ..., M\}$, it is optimal for retailer i to order $m\delta$ if and only if $\eta_{im}(\mathbf{s}, \alpha) \leq \theta_i \leq \eta_{i,m+1}(\mathbf{s}, \alpha)$, where $\eta_{im}(\mathbf{s}, \alpha)$ can be found iteratively as follows:

$$\eta_{im}(\mathbf{s}, \alpha) = \begin{cases} \underline{\Theta} & \text{if } m = 0 \\ \min\left\{\bar{\Theta},\ \ \max\left\{\eta_{i,m-1}(\mathbf{s}, \alpha), \bar{\theta}_{i,m-1}(\mathbf{s}, \alpha)\right\}\right\} & \text{if } m = 1, ..., M \\ \bar{\Theta} & \text{if } m = M+1 \end{cases}$$

Note that the thresholds derived from Lemma 1 might be outside $[\underline{\Theta}, \bar{\Theta}]$, it is thus necessary to set $\eta_{im}(\mathbf{s}, \alpha) \in [\underline{\Theta}, \bar{\Theta}]$. In addition, $\underline{\theta}_{im}(\mathbf{s}, \alpha) > \bar{\theta}_{im}(\mathbf{s}, \alpha)$ may arise and it is nonetheless without loss of information to set $\eta_{im}(\mathbf{s}, \alpha) = \eta_{i,m+1}(\mathbf{s}, \alpha)$ because $\sigma_{im}(\mathbf{s}, \alpha) = 0$. By definition, $\eta_{im}(\mathbf{s}, \alpha) \in [\underline{\Theta}, \bar{\Theta}]$ and $\eta_{im}(\mathbf{s}, \alpha) \leq \eta_{i,m+1}(\mathbf{s}, \alpha)$, $m = 0, ..., M$. This completes the proof for Proposition 1. ■

Proposition 2 Proof. Because no intertemporal dependence exists that affects the retailers' order quantities under fixed allocation, it is optimal for the retailers to maximize their per-period profits. With no capacity constraint, retailer i's per-period profit function is concave and is maximized at $q^M(\alpha, \theta)$. With capacity constraint, it is optimal for retailer i to sell as much as possible up to $q^M(\alpha, \theta)$. Hence, his optimal order quantity is the minimum of $q^M(\alpha, \theta)$ and his maximum allocation. Retailer i's maximum allocation equals the sum of his guaranteed allocation K/N and the maximum possible residual capacity unclaimed by all other retailers, which occurs when they experience the lowest local demand at $\underline{\Theta}$. ■

N retailers. For the case of N retailers, we modify the proof for Proposition 1. All three lemmas needed to prove the proposition still hold. The proof for Lemma 2 and 3 are the same as the case of two retailers. For Lemma 1, the proof is modified as follows. Consider the optimization problem of retailer i:

$$
\begin{aligned}
V_i(\mathbf{s},\alpha,\theta_i) &= \max_{x_i(s,\alpha,\theta_i)\in\{0\delta,...,\bar{m}\delta\}} \sum_{\mathbf{x}_{-i}\in\{0\delta,...,M\delta\}^{N-1}} \Pr(\mathbf{x}_{-i}) \\
&\quad \cdot\left\{\pi_i(\mathbf{s},\alpha,\theta_i,x_i,\mathbf{x}_{-i}) + \beta\sum_{\alpha'\in\mathcal{A}} P(\alpha'|\alpha)V_i\left(\mathbf{q}(\mathbf{s},x_i,\mathbf{x}_{-i}),\alpha'\right)\right\}.
\end{aligned}
$$

where $\bar{m}\equiv\max\{m\in\mathbb{Z}|\alpha+\varepsilon\theta_i-bm\delta\geq 0\}$ and

$$
\Pr(\mathbf{x}_{-i}) = \prod_{j\neq i}\prod_{l=0}^{M}\sigma_{jl}(\mathbf{s},\alpha)^{\mathbf{1}[x_j=l\delta]}.
$$

For $m=1,...,M-1$, $x_i^*=m\delta$ if and only if $m\leq\bar{m}$ and

$$
\begin{aligned}
&\sum_{\mathbf{x}_{-i}\in\{0\delta,...,M\delta\}^{N-1}} \Pr(\mathbf{x}_{-i})\left\{\pi_i(\mathbf{s},\alpha,\theta_i,m\delta,\mathbf{x}_{-i})+\beta\sum_{\alpha'\in\mathcal{A}}P(\alpha'|\alpha)V_i(\mathbf{q}(\mathbf{s},m\delta,\mathbf{x}_{-i}),\alpha')\right\}\\
\geq &\max_{x_i\in\{0\delta,...,M\delta\}}\sum_{\mathbf{x}_{-i}\in\{0\delta,...,M\delta\}^{N-1}} \Pr(\mathbf{x}_{-i})\left\{\pi_i(\mathbf{s},\alpha,\theta_i,x_i,\mathbf{x}_{-i})+\beta\sum_{\alpha'\in\mathcal{A}}P(\alpha'|\alpha)V_i(\mathbf{q}(\mathbf{s},x_i,\mathbf{x}_{-i}),\alpha')\right\}.
\end{aligned}
$$

The boundary cases, i.e., $m=0$ and $m=M$ can be easily derived following the same logic and is omitted for brevity. The second inequality is equivalent to for any $m'\neq m$, $m'\in\{0,...,m-1,m+1,...,M\}$,

$$
\theta_i\varepsilon B(m,m',\mathbf{s},\alpha)\geq -\Delta(m,m',\mathbf{s},\alpha).
$$

where

$$
\begin{aligned}
\Delta(m,m',\mathbf{s},\alpha) &= \sum_{\mathbf{x}_{-i}\in\{0\delta,...,M\delta\}^{N-1}} \Pr(\mathbf{x}_{-i})\left\{\pi_i(\mathbf{s},\alpha,0,m\delta,\mathbf{x}_{-i})+\beta\sum_{\alpha'\in\mathcal{A}}P(\alpha'|\alpha)V_i\left(\mathbf{q}(\mathbf{s},m\delta,\mathbf{x}_{-i}),\alpha'\right)\right.\\
&\quad\left.-\pi_i(\mathbf{s},\alpha,0,m'\delta,\mathbf{x}_{-i})-\beta\sum_{\alpha'\in\mathcal{A}}P(\alpha'|\alpha)V_i\left(\mathbf{q}(\mathbf{s},m'\delta,\mathbf{x}_{-i}),\alpha'\right)\right\},\\
B(m,m',\mathbf{s},\alpha) &= \sum_{\mathbf{x}_{-i}\in\{0\delta,...,M\delta\}^{N-1}} \Pr(\mathbf{x}_{-i})(q_i(\mathbf{s},m\delta,\mathbf{x}_{-i})-q_i(\mathbf{s},m'\delta,\mathbf{x}_{-i})).
\end{aligned}
$$

The rest of the proof is the same as the case of two retailers. ■

II. Approximation Logic for Fractional Allocations

We use a two-retailer example to illustrate the logic for approximating fractional number of batches using nearby integer number of batches. The logic can be extended to the N-retailer setting and is omitted for brevity. Let $m_i\equiv q_i/\delta$ and let $\lceil\cdot\rceil$ ($\lfloor\cdot\rfloor$) denote the operator to take the nearest integer larger (smaller) than the argument. It is straightforward to show that fractional number of batches only occur when capacity is lower than the sum of orders. Therefore, when $m_i\neq\lceil m_i\rceil$, we must have $\sum_{i=1}^{N} \lceil m_i \rceil = M + 1$ and $m_1 - \lfloor m_1 \rfloor + m_2 - \lfloor m_2 \rfloor = 1$. The allocations can be approximated using the following random variable:

$$\mathbf{q} = \begin{cases} (\lceil m_1 \rceil \delta, \lfloor m_2 \rfloor \delta) & \text{w.p. } m_1 - \lfloor m_1 \rfloor, \\ (\lfloor m_1 \rfloor \delta, \lceil m_2 \rceil \delta) & \text{w.p. } m_2 - \lfloor m_2 \rfloor. \end{cases}$$

III. Distribution Function for Beta (3,3)

$$F(\theta) = \begin{cases} 0 & \text{if} \quad \theta < -1, \\ \frac{1}{2} + \frac{15}{16}\theta - \frac{5}{8}\theta^3 + \frac{3}{16}\theta^5 & \text{if} \quad -1 \leq \theta < 1, \\ 1 & \text{if} \quad \theta \geq 1. \end{cases}$$

IV. Additional Figures

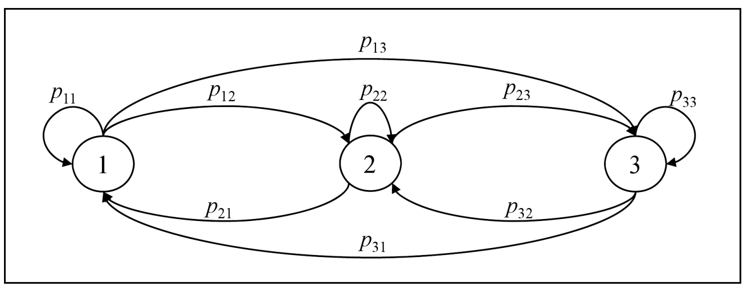

Figure 8: A Three-State Common Demand Process with Transition Probability p_{uv}

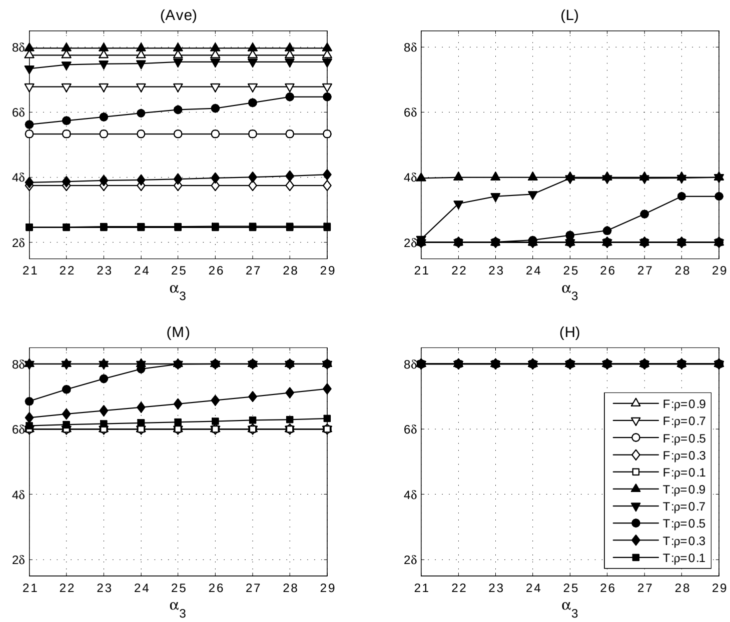

Figure 9: Expected Total Sales as a Function of α_3 Parameterized by ρ Under Fixed (F) and Turn-and-Earn (T) Allocation. $K = 8\delta$, $\mathcal{A} = \{5, 13, \alpha_3\}$, $\varepsilon = 4$.

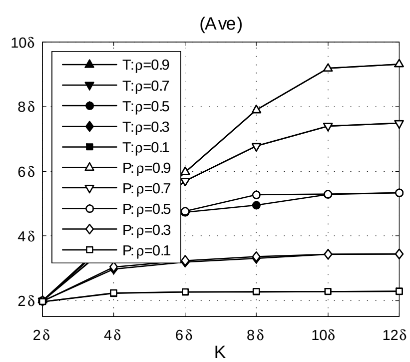

Figure 10: Expected Total Sales as a Function of K Parameterized by ρ Under Turn-and-Earn (T) and Proportional Turn-and-Earn (P) Allocation. $\mathcal{A} = \{5, 13, 21\}$, $\varepsilon = 4$.