



Dynamic pricing with multiple consumers and alternating offers under retailer competition: theory and experiment

Amnon Rapoport¹ · Eyran J. Gisches² · Vincent Mak³ · Rami Zwick¹

Received: 12 December 2023 / Revised: 14 September 2024 / Accepted: 14 September 2024 /
Published online: 1 November 2024
© The Author(s) 2024

Abstract

We introduce and test a stylized model of dynamic pricing under duopolistic competition. In our model, a consumer receives alternating price offers between two retailers over an indefinite number of periods so that the game or “season” terminates with a fixed probability after each period. The two retailers do not know the valuation of the consumer for the good they are competing to sell to the consumer, but they have common knowledge about the probability distribution of the valuation. Our equilibrium analysis suggests that price offers decrease exponentially across periods over the season. Moreover, when there are multiple consumers in the game, as long as their valuations are ex ante independently and identically distributed, the equilibrium predictions are the same regardless of the number of consumers. An experiment on the model showed that subjects acting as retailers often overpriced relative to equilibrium predictions. In addition, the theoretical invariance with respect to the number of consumers did not hold: consumers seemed to be more prone to strategic waiting in the first period of the season when there were multiple consumers (compared with when there was only a single consumer), leading to a decrease in the per-consumer payoff of the retailer who made the price offer in the first period and a corresponding increase in per-consumer payoff of the other retailer. There is also evidence of within-session evolution that led to lower retailer prices that were closer to equilibrium predictions, and higher tendency for consumer strategic waiting, as the session progressed.

Keywords Dynamic pricing · Seller competition · Multiple buyers · Alternating offers

JEL Classification C72 · C78 · C92 · D82

The replication material for the study is available at https://osf.io/g63w5/?view_only=e23272b81a824f4fbd8286a41ede294f.

Professor Amnon Rapoport is deceased.

Extended author information available on the last page of the article

1 Introduction

Revenue management (e.g., Phillips, 2005; Talluri & Ryzin, 2004) can be characterized as a set of propositions for rational retailers of how to sell the right inventory to the right consumer at the right time and the right price (Kimes, 1989). This is mostly achieved by allocation of inventory, by dynamic pricing of the good, or by some combination of both. Subsumed under this very general characterization are multiple models that differ from one another on one or more critical assumptions about the type of consumers who enter the market (e.g., myopic, strategic), the nature of the competition between the retailers (if there are more than one), the information structure (e.g., whether price information is common knowledge among the retailers), the nature of the demand (stochastic vs. deterministic, time-sensitive), capacity constraints, replenishment of the good during the selling season, salvage costs, and, in general, the rules that govern trade in the market. A major contribution of the theoretical literature in revenue management has been the demonstration of how important these assumptions are and what effect they might have on inventory allocation and dynamic pricing. A major contribution of the experimental literature, mostly in economics and marketing, has been to demonstrate systematic deviations of traders' decisions from the theoretical models.

In the present paper, we contribute to both theory and experimentation by introducing and subsequently testing a stylized model of dynamic pricing under retailer competition. Specifically, competition in our model is conducted through the retailers making alternating offers to consumers. Our approach builds on previous studies by Mantin et al. (2007) (MGG), and Mak et al. (2012) (MRG), which also extends a line of experimental research on monopolistic dynamic pricing that includes Rapoport et al. (1995) and Mak et al. (2014).

In the family of models investigated in MGG and MRG as well as here, the consumer's valuation of the good is denoted by v . The two retailers do not know the value of v , but they are assumed to have common prior knowledge about the probability density of v . The two retailers are assumed to maximize their individual payoffs (profits) over a season (horizon) with a finite number of periods, and in each period they can post independently, and with no restrictions, new prices. In each period, the consumer is free to visit one and only one of the two retail stores. A critical element of this family of models concerns the pattern of visits: if in a period the consumer visits a retailer and observes a posted price that exceeds their valuation, then, in the next period, they may either return to the same retailer with probability P or switch to the competing retailer with probability $1 - P$. MGG refer to the special case of $P = 0$, where the consumer deterministically switches retailers from one period to another, as "zigzag competition."

There may be several alternative reasons for this first-order Markov store-choice pattern of visits by the consumer in such finite, multi-period, dynamic pricing settings. As noted in the marketing literature on brand choice and brand loyalty, switching behavior may be attributed to variety seeking (Lillien et al., 1992; Kahn, 1995), as the consumer is assumed to have no prior knowledge

about the posted prices in the market at any given period. The consumer switching probability $1 - P$ may serve as a proxy to the consumer's loyalty, or reflect their geographical proximity to the two stores that may change from one period to another. It also may reflect the consumer's experience at the store, the type of service they encounter, the actual time that elapses from one period to another, and their shopping habits.

MGG considered the case of myopic consumers, who purchase the good if and only if its valuation is higher than the posted price, and then extended their original model to strategic (i.e., forward looking) consumers. MRG considered the latter case in more detail. MGG extended their model of a single consumer to a model with multiple consumers, whose valuations are privately known and independently and identically distributed according to a uniform distribution. Importantly for our study, both MGG and MRG have noted that, under zigzag competition, the setting for multiple consumers trivially coincides with that of one consumer; a major purpose of our study is to test this conclusion experimentally. As discussed in detail below, our experiment showed that the theoretical invariance with respect to the number of consumers did not hold.

The present paper makes two major contributions. The first contribution is an extension of the finite-horizon dynamic pricing models of MGG and MRG to a model where the horizon is indefinite with probabilistic continuation. This is equivalent to an infinite horizon with time discounting in which all players, including retailers and consumers, have the same per period time discount factor for pay-offs (see, e.g., Roth & Murnighan, 1978; Zwick et al., 1992). The motivation of this extension is that for many goods, even so-called seasonal goods, the assumption that the "world ends" precisely after a finite number of periods is frequently made to gain mathematical tractability rather than reflect reality. The finite-horizon assumption might be appropriate for products with which the selling firm only has a limited number of opportunities to change prices during the selling season, such as fresh food that needs to be sold within a day. On the other hand, in cases where the selling season is long with numerous opportunities to change price dynamically, the infinite-horizon approach might be more appropriate.

In terms of experimental research on dynamic pricing, the extension to an infinite horizon seems to create a problem as experimental sessions terminate at finite time. While, strictly speaking, infinite-horizon models with time discounting are not experimentally testable, the way around this problem is to adopt an indefinite horizon approach and setting the continuation probability (i.e., the size of the discount factor) sufficiently low so that the probability of inordinately long sessions becomes negligible. Therefore, in the present study, we implemented the following procedure. If the consumer rejects an offer in a period, and therefore they are still "active" in the market, then the selling season continues with continuation probability δ or ends due to a random process with termination probability $1 - \delta$, and the active consumer earns zero.

The second contribution is a change of focus in the present paper from one consumer to multiple consumers; a major purpose of the present study is to experimentally test the effects of the *number* of consumers on the pricing of the (homogeneous) good and, consequently, on the evolution of this market. We proceed under the

specification, in keeping with MGG and MRG, that if there are multiple consumers in the market, then any consumer's valuation, v , is ex ante independently and identically distributed according to a uniform distribution. We also adopt the complete information specification that, as in MGG and MRG, the price history is commonly known. With these specifications, the equilibrium pricing and purchase decisions of the sellers and consumers should not be affected by the number of consumers. Such equivalence is essentially an equivalence between selling to one buyer with a prior valuation distribution (such as in some business-to-business settings) and selling to a market of multiple consumers with valuations that are independently and identically distributed according to the same distribution (such as for many fast-moving consumer goods). However, behaviorally, it is important to see whether and how human sellers' pricing decisions, as well as human consumers' purchase decisions in response to the prices, could differ according to the size of the market. There are several reasons why the descriptive power of the analytical prediction of equivalence is questionable:

First, some of the consumers may appear to be more myopic than others—they would be more likely to purchase whenever their valuation is higher than the currently offered price, and less likely to attempt to strategically wait for a prospectively lower price with higher earnings in future periods. Retailers who are aware of the presence of such heterogeneity in consumer strategic behavior might adjust their pricing decisions accordingly. The presence of myopic consumer decisions in dynamic pricing contexts is a well-known experimental finding. For example, in an experimental study on dynamic pricing, Mak et al. (2014) reported that a substantial percentage of their buyer subjects behaved myopically although they were not instructed to do so. Osadchiy and Bendoly (2015) also reported that while many of their subjects exhibited strategic behavior, some were consistently purchasing myopically. Further theoretical and experimental developments include Baucells et al. (2017), Kremer et al. (2017), and Chen and Zhao (2020), among others. In our case, if different numbers of consumers lead to differences in consumers' tendency to be myopic or to attempt strategic waiting, the retailers might respond to such differences with different pricing strategies.

Second, the number of consumers can influence the length of the season. This is because the season ends when either a random termination event occurs, or the market clears. With a single consumer, the market clears immediately upon purchase, potentially cutting the season short even if a random termination wouldn't have happened. This has implications for learning in terms of observations of ex post suboptimal decisions, i.e., decisions that are suboptimal vis a vis the ex post observations of the full set of posted prices by retailers in a season after the season is over. In the single-consumer case, the only observable ex post suboptimal decision is ex post "too late" (passing on a lower price and buying later at a higher price, or the season ending without a purchase despite encountering a price that is below valuation). In the multiple-consumer case, ex post "too early" purchases (buying at a higher price than the season's lowest that appears afterwards) become observable alongside ex post "too late" instances. This difference in ex post suboptimal observations might influence consumer behavior in the multiple-consumer case compared to the single-consumer case. Specifically, the ability to observe ex post "too early"

purchases might lead to more strategic waiting behavior by consumers in the multiple-consumer case.

Lastly, since different consumers might purchase in different periods at different prices, there can be social comparison effects that are not present when there is only one consumer. These can be due to, for example, fairness concerns, as well as socially induced regret (over the possibility that another consumer might be able to purchase at a better price) and could affect consumer purchase decisions as well as the pricing decisions of retailers who anticipate such effects. However, since consumers' valuations vary and are only privately known, horizontal comparison between consumers' surplus is not possible. Even if another consumer purchases at a lower price than a focal consumer during the same season (most likely at a later period), there is no guarantee that the other consumer's surplus is higher than the focal consumer's, thus potentially reducing the impact of fairness considerations among consumers. The more pronounced comparison is likely to be the consumer's own ex post comparison of their decisions during a season against the prices in that season, which could induce regret vis-a-vis the realization of ex post suboptimality (as discussed in the previous paragraph).

This paper presents our theoretical and experimental research contributions. Section 2 formally states our model of zigzag competition as employed in our experiment. Our equilibrium analysis suggests that price offers decrease exponentially across periods over the season. Moreover, when there are multiple consumers in the game, as long as their valuations are ex ante independently and identically distributed, the equilibrium predictions are the same regardless of the number of consumers.

Section 3 describes our experiment, which was designed to test the implications of the model and account for the dynamics of play and behavioral deviations across multiple iterations or seasons of the stage game. In accordance with common experimental procedures, we simplified the game presented to the experimental subjects in two different ways. First, for all experimental conditions, we impose a deterministic switching pattern on all the consumers to create a zigzag competition setting. As a consequence, the subjects (retailers) in our experiment had a simpler task as they had to consider only one rather than two sources of uncertainty. Secondly, we have restricted the experimental investigation to a single (relatively high) value of the continuation probability ($\delta=0.75$). To investigate evolution of behavior from experience, we iterated the basic dynamic pricing game for multiple seasons, and had the valuations of each of the consumers drawn randomly and independently on each season. Player roles (i.e., Retailer 1, Retailer 2, Consumer) remained fixed across all seasons to minimize confusion due to constant changes of role and guarantee sufficient data for each player.

Section 4 reports the main findings of the experiment. Briefly, our experimental study showed that retailer subjects often overpriced relative to equilibrium predictions. In addition, the theoretical invariance with respect to the number of consumers was violated: consumers seemed to be more prone to strategic waiting in the first period of the season when there were multiple consumers (compared with when there was only one consumer), leading to a decrease in the per-consumer payoff of the retailer who made the price offer in the first period and

a corresponding increase in per-consumer payoff of the other retailer. There is also evidence of within-session evolution that led to lower retailer prices that were closer to equilibrium predictions, and higher tendency for consumer strategic waiting, as the session progressed. Section 5 concludes by summarizing our results and outlining extensions of the model for future research.

2 The model

2.1 A dynamic pricing model

We model the selling season as a non-cooperative game among three players including two retailers called Retailer 1 and Retailer 2, and N Consumers. Over the duration of the game, which consists of an indefinite number of periods, each Consumer demands one unit of an indivisible good. Each Retailer has N units of the good, which is of zero value for them. Each Consumer's valuation of the good, which is denoted by v , is privately known by the Consumer but not by the Retailers. Rather, the two Retailers hold a common prior knowledge that v is independently and identically distributed among the Consumers following the uniform distribution over the interval $[0, \bar{v}]$.

In period 1 all of the Consumers visit Retailer 1, who offers to sell them the good at the price p_1 . Each Consumer independently and simultaneously decides either to accept the offer or reject it. If all Consumers accept the offer, then the season terminates, each Consumer earns $v - p_1$ (v being the Consumer's own valuation), Retailer 1 earns Np_1 (as the marginal cost of the good is zero for both Retailers), and Retailer 2 earns zero. If some Consumers, say n of them, accept the offer, but the remaining $N - n$ Consumers reject the offer, then each accepting Consumer earns $v - p_1$ and leaves the game, Retailer 1 earns np_1 , and Retailer 2 earns zero, while the season continues to the next period with continuation probability δ and terminates with probability $1 - \delta$.

Period 2—assuming that it takes place—is structured in the same way with the exception that Retailer 2 makes the price offer p_2 . The game proceeds accordingly with the Consumers zigzagging between the two Retailers until last remaining Consumers in the game accept a price p_t at period t or the season terminates with no trade, whichever comes first. In any period t :

- (a) if t is odd, each Consumer who accepts the offer in that period earns a payoff of $v - p_t$, Retailer 1 earn p_t per accepting Consumer, and Retailer 2 earns 0;
- (b) if t is even, each Consumer who accepts the offer in that period earns a payoff of $v - p_t$, Retailer 1 earns 0, and Retailer 1 earn p_t per accepting Consumer.

Every price offer is commonly known once it is posted to the Consumers. Put differently, each Retailer knows the entire history of the offers when it is their turn to post a price and how many Consumers are still active in the market. It

is also commonly known that there is no scarcity: each Retailer holds at least N units of the good.

As mentioned in Sect. 1, this game is formally equivalent to the one in which the Consumers zigzag infinitely many times between the two Retailers and each of the three players has the same per period time discount factor for payoffs.

2.1.1 Equilibrium properties

We focus on pure-strategy rational expectations equilibria in which the Retailer’s price offers are deterministic conditioned on the history of trade, regardless of whether that history is in or out of equilibrium. Proofs of the following results are relegated to Online Appendix 1.

Lemma 1 *In any strategy pricing equilibrium, a Consumer with valuation v_1 does not buy later than a Consumer with valuation v_2 , if $v_1 > v_2$.*

Proposition 1 *Let α be the unique solution in $(0,1)$ of the equation $\delta^2\alpha^4 - 2\alpha + 1 = 0$. The dynamic pricing game has a unique pure-strategy pricing equilibrium, in which:*

- (1) *the equilibrium price in period t is $p_t^* = v_t^* r = \bar{v}\alpha^t(1 - \delta)/(1 - \delta\alpha)$, where $r = [\alpha(1 - \delta)]/(1 - \delta\alpha)$,*
- (2) *a Consumer whose valuation is contained in the interval $(v_{t+1}^*, v_t^*) = (\bar{v}\alpha^t, \bar{v}\alpha^{t-1})$ purchases the good in period t ;*
- (3) *a Consumer whose valuation is exactly $\bar{v}\alpha^{t-1}$ purchases the good in period t or $\max\{1, t - 1\}$;*
- (4) *the ex ante probability of a transaction happening in period t is $\delta^{t-1}(v_t^* - v_{t+1}^*)/\bar{v} = \delta^{t-1}\alpha^{t-1}(1 - \alpha)$;*
- (5) *the ex ante payoff of Retailer i is $\bar{v}\alpha(1 - \alpha)(1 - \delta)(\delta\alpha^2)^{i-1}/[(1 - \delta\alpha)(1 - \delta^2\alpha^4)]$;*
- (6) *the ex ante payoff of the Consumer is $\bar{v}(1 - \alpha)^2(1 + \delta\alpha)/[2(1 - \delta\alpha)(1 - \delta\alpha^2)]$,*

According to Proposition 1 (statement (1)), prices decay exponentially across periods, a feature that is shared with the monopolistic dynamic pricing model analyzed in Rapoport et al. (1995). The Consumer’s behavior and the posted prices exhibit “skimming” characteristics. The segment of the distribution of v that is skimmed decreases in size exponentially across the periods—by a factor of α from period to period—and so are the upper/lower bounds of the skimmed segment. The equilibrium price in any given period is always r times the upper bound of the posterior distribution of v in that period. Lastly, \bar{v} , which is the upper bound of the prior distribution for v , only appears as a scaling factor. All of these results hold because of the symmetry imparted by the random termination process (or, alternatively, because all the players in the dynamic pricing game have the same per period time discount factor and the game has infinitely many periods), as well as the fact that v is uniformly distributed with a lower bound of 0. Table 1 illustrates the equilibrium solution and the out-of-equilibrium properties (to be discussed next) for the parameter values implemented in our experiment, namely, $\delta = 0.75$ and $\bar{v} = 1000$.

Table 1 Equilibrium solution and the out-of-equilibrium properties for the parameter values implemented in our experiment

1. Equilibrium predictions for $\bar{v} = 1000, \delta = 0.75$ (so that $\alpha = 0.521$ and $r = 0.214$ following the notations in Proposition 1).

Period	Eq. price	Eq. price	Eq. valuation of Consumers who purchase the good	
	Retailer 1	Retailer 2	min	max
1	213.79		521.00	1000.00
2		111.38	271.44	521.00
3	58.03		141.42	271.44
4		30.23	73.68	141.42
5	15.75		38.39	73.68
6		8.21	20.00	38.39
7	4.28		10.42	20.00
8		2.23	5.43	10.42
9	1.16		2.83	5.43
10		0.60	1.47	2.83
Ex ante expected payoff	106.83	21.75	328.82	

2. Equilibrium and out-of-equilibrium values for computing best responses for the dynamic pricing game with $\bar{v} = 1000, \delta = 0.75$.

Equilibrium price in period t	$213.79 \times (0.521)^{t-1}$
Equilibrium valuation range for a Consumer who purchases the good in period t	$[1000 \times (0.521)^t, 1000 \times (0.521)^{t-1}]$
Equilibrium ex ante probability of a transaction happening in period t	$47.9\% \times (0.391)^{t-1}$
Equilibrium ex ante payoff of Retailer 1	106.83
Equilibrium ex ante payoff of Retailer 2	21.75
Equilibrium ex ante expected payoff of the Consumer	328.82
Formulas for out-of-equilibrium calculation	$v_{s+1}^i = \min\{2.44p_{ss}, v_s^i\}, p_s^{**} = 0.214 \times v_s^i$

Figure 1 exhibits values of α and r as a function of δ . On both curves, the points for $\delta=0.75$ —the value used in our experiment—are marked. Figure 1 shows that the parameter α is not very sensitive to the value of δ ; in fact, it always lies within the interval $[0.5, 0.55]$ over the full possible range of δ . Recall that α is the ratio by which the equilibrium prices decrease from period to period, and also the ratio by which the support of the updated distribution of the Consumer’s valuation is skimmed from period to period (see Proposition 1). Figure 1 therefore suggests that the equilibrium price always undergoes an approximately 50% discount per period, while the Consumer valuation range is at the same time skimmed by approximately 50%.




Fig. 1 Plots of the parameters $\alpha = v_{t+1}^*/v_t^*$ and $r = p_t^*/v_t^*$ as a function of the probability of continuation δ for the dynamic pricing game

On the other hand, Fig. 1 shows that r is relatively sensitive to δ . Recall that r is the ratio between the equilibrium price in a period and the high end of the valuation distribution in the same period; for example, it is the ratio of the period 1 equilibrium price relative to \bar{v} . Its steady decrease as δ increases, while α remains largely around 0.5, suggests that high discount factor leads to strong competition that drives down prices, while forward looking Consumers' strategic waiting leaves the period-to-period skimming largely unaffected. As $\delta \rightarrow 0^+$, then r tends to the single-period monopolistic selling limit of $1/2$. As $\delta \rightarrow 1^-$, r tends to the Coase conjecture limit of 0 (see e.g., Coase, 1972; McAfee & Wiseman, 2007).

2.2 Out-of-equilibrium properties

The equilibrium solution for the model in Proposition 1 is obviously useful for data analysis if both the Retailers and the Consumer follow the equilibrium path on each period of the selling season. Clearly, it is unrealistic to expect them to do so. If they depart from the equilibrium path, then the expressions stated in the propositions are no longer valid. Best response predictions have to be derived for each of the models to account for the decision in period $t + 1$ conditional upon the *posted* (rather than *predicted*) price offers up to this period.

Consider a sequence of *posted* price offers up to period t and label them p_1, p_2, \dots, p_t . Then, generate the sequence $v_1', v_2', \dots, v_t', v_{t+1}'$ as follows:
 $v_1' = \bar{v}$.

$$v_{s+1}' = \min\{(1 - \delta\alpha)p_s J(1 - \delta), v_s'\}.$$

It can be deduced from the proof of Proposition 1 (see Online Appendix 1) that the valuation of the best responding Consumer who accepts an offer in period s must be contained within $[v_{s+1}', v_s']$. If a best responding Consumer's valuation is higher (lower) than the upper bound of this interval, then they accept (reject) the offer in period s . If their valuation is exactly the lower bound of this interval, then they may or may not accept the price offer; since this is a zero measure event, it is still legitimate to assume that the posterior distribution of the valuation v in period $s+1$ is a uniform distribution on the interval $[0, v_{s+1}']$. This means that the posterior or updated distribution of v in any period which is in or out of equilibrium is, indeed, always a uniform distribution over a support with a lower bound zero.

The best response of a Retailer whose turn is to post an offer in period s , given the posterior distribution at that period, is given by $p_s^{**} = rv_s'$, where r is computed from Proposition 1. Because all the offers are assumed to be known, any Retailer may work out the v_s' in period s , given the previous history of play, whether that is in or out of equilibrium.

Finally, it is clear from the setup of the model and our derivations that, when there are multiple consumers in the game, as long as their valuations are ex ante independently and identically distributed, the equilibrium and out-of-equilibrium predictions are the same regardless of the number of consumers. Players' strategies, including pricing, should also not be affected by how many Consumers there are (i.e., have not yet purchased) in the current period. This feature of our analysis was tested in the experiment reported in the next section.

3 Experimental method

3.1 Subjects

Two hundred and fifty-nine subjects participated in the experiment. The subjects were undergraduate students at the University of Arizona and the University of California, Riverside, who volunteered to take part in a group decision making experiment for payoff contingent on performance.

3.2 Design

The experiment had a three-condition between-subjects design. Subjects in all conditions played the dynamic pricing game described in Sect. 2 with different number of Consumers: one Consumer in Condition 1B, three Consumers in Condition 3B, and five Consumers in Condition 5B, respectively. Subjects participated in experimental session cohorts that, depending on the experimental condition, ranged from 12 to 15 in size. Within a session, they were randomly assigned to groups of three, five, or seven players according to whether the experimental condition was Condition 1B, 3B, or 5B for that session. Condition 1B included 24 groups of three players each

Table 2 Number of sessions and number of groups and subjects in each session in the respective conditions

	<i>Number of sessions/groups/subjects</i>	<i>Total no. of subjects (groups)</i>
Condition 1B	4 sessions \times 5 groups of 3 subjects each	72 subjects (24 groups)
	1 session \times 4 groups of 3 subjects each	
Condition 3B	5 sessions \times 3 groups of 5 subjects each	75 subjects (15 groups)
Condition 5B	8 sessions \times 2 groups of 7 subjects each	112 subjects (16 groups)

(two Retailers, one Consumer) across all of its sessions; correspondingly, Group 3B included 15 groups of five players each (two Retailers, three Consumers); and Condition 5B included 16 groups of seven players each (two Retailers, five Consumers). In each of the three conditions, the game was iterated for 50 seasons.¹ Player roles (i.e., Retailer 1, Retailer 2, Consumer) were held fixed for every subject across all the seasons. However, the composition of each group was determined randomly by re-matching subjects in the same experimental session at the beginning of each season; these measures were implemented in order to minimize sequential effects over the seasons. Each experimental session consisted of two to five groups as detailed in Table 2. The valuation of each Consumer was drawn randomly and from the set $\{1, 2, \dots, 1000\}$ at the beginning of each season.

3.3 Procedure

All the sessions were conducted in a large computer-controlled laboratory with networked terminals that were separated from one another by partitions. Communication between the subjects was strictly forbidden. Once they entered the laboratory, subjects were assigned randomly to their cubicles and proceeded to read the subject instructions at their own pace. The instructions for Condition 5B are presented in Online Appendix 2, where the terms Seller 1, Seller 2, and Buyer stand, respectively, for Retailer 1, Retailer 2, and Consumer in the model and the general terminology in this paper.

The instructions described the trading mechanism, explained the valuation of the good for the Consumer, and then informed the subject about the calculation of the payoff (profit) for each season, the probability of terminating each period randomly, the Retailer's computer screen, the Consumer's computer screen, and the determination of the final payoff for the session. All queries were answered privately by the experimenter.

Each period t was structured as follows. The Retailers were presented with their Decision Screen (see Online Appendix 2) that included information about the player role, period number, season number, and the pricing log of both Retailers' prices

¹ Exceptionally, one session in Condition 5B had 49 seasons due to a technical mishap.

Table 3 Number of seasons of different lengths (number of periods)

Length (period)	1	2	3	4	5	6	7–16	Total
<i>Condition 1B</i>								
Seasons 1–25	352	158	36	39	8	6	1	600
Seasons 26–50	384	140	36	21	11	2	6	600
<i>Condition 3B</i>								
Seasons 1–25	143	114	54	34	15	5	10	375
Seasons 26–50	117	148	61	20	18	4	7	375
<i>Condition 5B</i>								
Seasons 1–25	123	117	56	46	29	9	20	400
Seasons 26–50	119	122	89	35	21	4	8	398

in the season and how many active Consumers remain in the market. If t was odd (even), Retailer 1 (2) was instructed to post their price for the period.

The Consumer was presented with their Buyer Screen that displayed the Buyer's valuation of the good for the season, the period number t , the posted price for period t , the Buyer's potential payoff were they to accept the present offer ($v_t - p_t$), and the pricing log (history) of both Retailers' prices in the season.

At the end of each season, each group member was presented with a summary screen with information about all the posted prices in the season, the member's own transaction (if there was one), and the member's payoff for the season. Once all the groups in the session completed the season, the experiment proceeded to the next season. Thus, all the groups in a given session started each season at the same time.

At the end of the session, after completing 50 seasons, the total payoffs from m randomly selected seasons were converted to US dollars at a rate of 80 experimental profit points = \$1.0 to become the subjects' earnings, plus a \$5 show-up fee which they received privately. The number m was the same across player roles in Condition 1B but varied between roles in Conditions 3B and 5B to mitigate disparity in dollar payments. Specifically, in Condition 1B, $m = 10$ for all roles; in Condition 3B, $m = 5$ for Retailer 1 and the Consumer while $m = 20$ for Retailer 2; in Condition 5B, $m = 3$ for Retailer 1, $m = 12$ for Retailer 2, and $m = 5$ for the Consumer; note that each subject was only informed about the m for the role they were assigned. Upon receiving payments, the subjects were dismissed from the laboratory. All the sessions were completed in less than two hours.

4 Results

Since subjects in the same experimental session were randomly re-matched from season to season during the session (while keeping to the same player role), the independent unit of data analysis should be the session. Hence, the means reported in this section, including those in the tables and figures, are all calculated with session as the unit of analysis unless otherwise stated. On the other hand, when performing statistical tests, we typically use regression (linear or logistic depending on

whether the dependent variable is interval/ratio or binary) clustering variances by session subjects.² The same applies to testing deviation from equilibrium analysis benchmarks—we would typically create a dependent variable that is the difference between the experimental data point and the relevant equilibrium analysis benchmark, and then run an intercept-only regression model without covariates, again clustering variances by session subjects; statistically significant deviation from the benchmark can be tested by looking at whether the intercept term from such a regression analysis is significantly different from zero or not.

Less than 2% of the seasons ended with more than six periods (see Table 3). Hence, where a statistical analysis pertains to period-by-period variables such as the price in each period, we focus only on periods 1 to 6. Lastly, since there is evidence of within-session evolution (see below), we typically conduct statistical analysis with a distinction between whether a data point came from the earlier half of a session (Seasons 1 to 25) or the latter half (Seasons 26 to 50).

In the next subsections, we first report basic results regarding the experimental decisions. We next report a set of analysis comparing the Retailer pricing and Consumer purchase decisions with best responses prescribed by equilibrium analysis. Afterwards, we report analysis that more directly relates to between-condition differences and within-session evolution.

4.1 Basic results

We first note that the average length of each season in the experiment, in terms of periods, is 1.66 for Condition 1B, 2.26 for Condition 3B, and 2.57 for Condition 5B (see also Table 3). These are all significantly different from each other at $p < 0.01$ by regression clustering variances by session.³ The percentage of seasons that ended with transactions being made with all the Consumers is 74.25% (891 out of 1200) in Condition 1B, 47.2% (354 out of 750) in Condition 3B, and 42.86% (342 out of 798) in Condition 5B. The mean number of transactions per season is 0.74 (74% of the number of Consumers) for Condition 1B, 2.00 (67% of the number of Consumers) for Condition 3B, and 3.60 (72% of the number of Consumers) for Condition 5B.

Another basic result from the experiment is that Retailers often overpriced relative to equilibrium predictions (see, e.g., similar results in MRG and Rapoport et al., 1995). This is borne out in Fig. 2, which plots the mean prices by condition and earlier/latter half of the session; the equilibrium predictions are also plotted for comparison. Note that the mean prices did not always decrease as the period increases (a steady decrease would have conformed to a basic equilibrium prediction). Statistical tests show that the mean prices over periods 1 to 6 are often significantly different from equilibrium predictions at $p < 0.05$ according to intercept-only regression

² We use the `proc genmod` routine in the SAS software with a GEE approach for this part of the analysis.

³ For example, when comparing between Conditions 1B and 3B, we regress the number of periods in each season from either of the conditions on a single regressor, namely an indicator function that is 1 when the data point came from Condition 3B, and 0 otherwise. The statistical significance of the estimated coefficient of the regressor informs the conclusion discussed here.

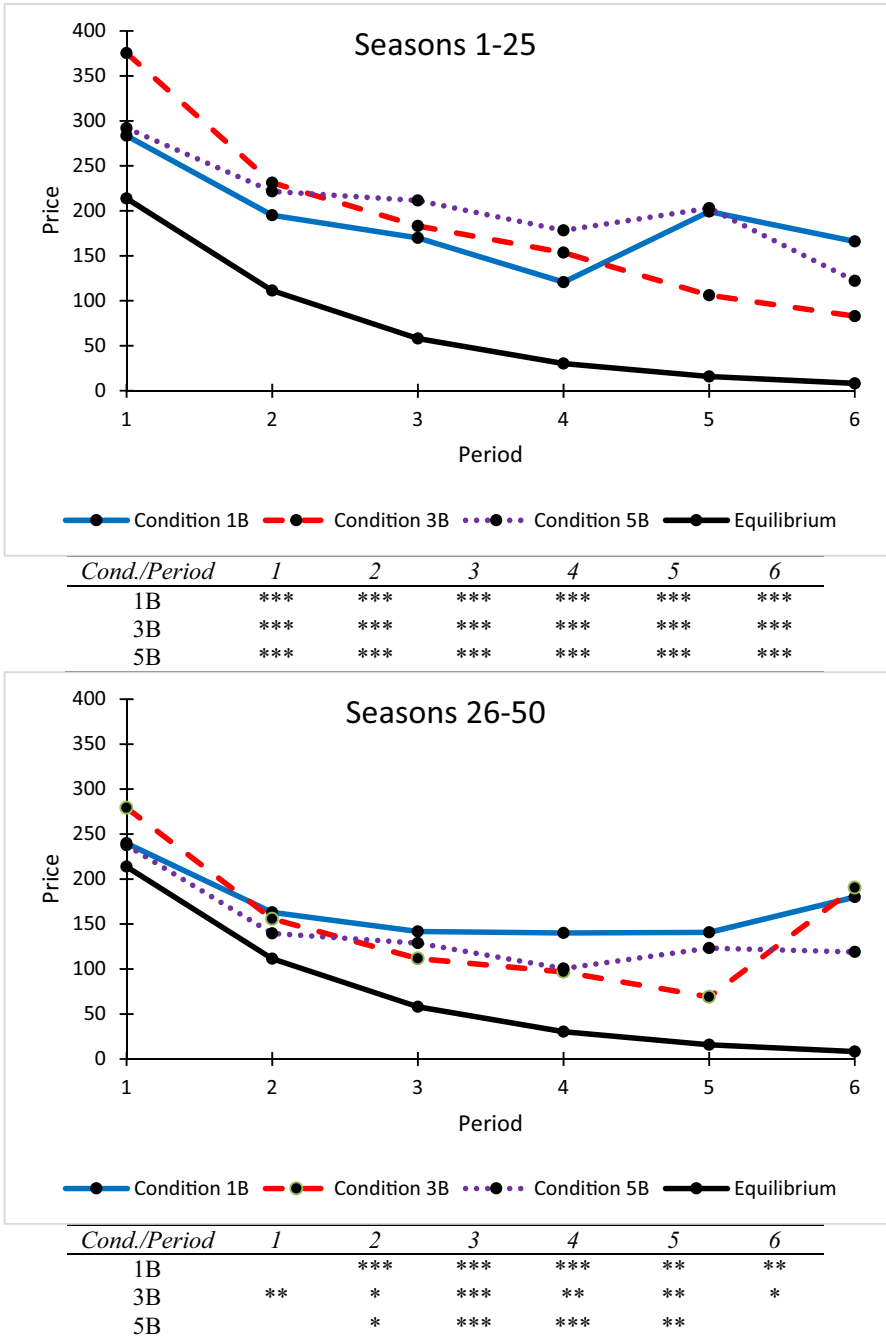


Fig. 2 Mean experimental and equilibrium prices up to period 6. The unit of analysis is the session. The accompanying tables indicate, with one or more asterisks, the mean prices by period that are significantly different from an equilibrium prices according to an intercept-only regression analysis clustering variances by subjects ($*p < 0.05$, $**p < 0.01$, $***p < 0.001$)

Table 4 Experimental and equilibrium percentages of Consumer journeys (see Sect. 4.1) by period of purchase (absolute frequencies in parentheses)

Period	1	2	3	4	5	6	7–9	Total
<i>Condition 1B</i>								
Experiment	46.33 (556)	19.33 (232)	4.17 (50)	3.25 (39)	0.67 (8)	0.33 (4)	0.17 (2)	74.25 (891) ^a
Eq.	47.90	18.72	7.31	2.86	1.12	0.44	0.27	78.61
<i>Condition 3B</i>								
Experiment	30.44 (685)	24.18 (544)	6.40 (144)	3.11 (70)	1.56 (35)	0.44 (10)	0.48 (11)	66.62 (1499) ^b
Eq.	47.90	18.72	7.31	2.86	1.12	0.44	0.27	78.61
<i>Condition 5B</i>								
Experiment	39.62 (1581)	20.75 (828)	6.69 (267)	2.86 (114)	1.25 (50)	0.40 (16)	0.43 (17)	72.01 (2873) ^c
Eq.	47.90	18.72	7.31	2.86	1.12	0.44	0.27	78.61

Note that no transaction took place after period 9 in any condition in the experiment

^aOut of a total count of 1200

^bOut of a total count of 2250

^cOut of a total count of 3990

analysis clustering variances by subjects. The notable exceptions are the period 1 prices in the latter half of the session in Conditions 1B and 5B, suggesting within-session evolution that drove down prices towards the competitive equilibrium level.

We next look at Consumer purchases. Define a Consumer journey as the sequence of price offers received and purchase decisions made during a season by a Consumer player; in Condition 1B, each season involved one Consumer journey, while in Condition 3B (5B), each season involved three (five) Consumer journeys. Table 4 displays the experimental and equilibrium percentages of Consumer journeys that resulted in purchases in periods 1–9. While equilibrium predictions suggest that transactions should occur in 78.6% of the Consumer journeys in the experiment over periods 1–10, the actual percentages in the experiment are markedly lower in Conditions 3B and 5B. In other words, the markets in the experiment cleared much less efficiently than suggested by theoretical analysis.

To have an appreciation of the relative levels of strategic behavior among Consumers in the different conditions, for each Consumer player in each session, we look at instances when they were in the game (i.e. had not yet purchased in the current season) while their valuation was higher or equal to the current price; we then work out the percentage proportion of those instances throughout the session in which they did *not* buy, to form a measurement of attempted strategic waiting. We find that, across Consumer players in Condition 1B, the mean of this proportion is

Table 5 Mean payoffs by condition and earlier/latter half of the session (s.d. in parentheses), together with the relevant equilibrium level

	Condition 1B Seasons		Condition 3B Seasons		Condition 5B Seasons		Eq.
	1–25	26–50	1–25	26–50	1–25	26–50	
Retailer 1 normalized	92.83** (18.87)	99.00 (23.44)	110.98 (38.08)	69.94*** (15.36)	86.76* (30.71)	72.52*** (30.41)	106.83
Retailer 2 normalized	36.48*** (15.16)	26.05 (9.68)	37.73*** (8.21)	38.40*** (8.36)	36.31*** (12.53)	33.48*** (14.01)	21.75
Consumer	285.84*** (31.31)	318.35 (17.76)	219.76*** (38.98)	290.50*** (44.04)	277.92*** (55.67)	316.75 (42.49)	328.82

The unit of analysis is the session. Note that the Retailers' mean payoffs are normalized by the number of Consumers in the experimental game in the corresponding condition, so that, for example, for Condition 3B and Retailer 1, the table shows the mean of the Retailer's payoff divided by three. Where the mean entry is significantly different from equilibrium according to an intercept-only regression analysis clustering variances by subjects, the entry is indicated by one or more asterisks (* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$)

33.00%, compared with 43.73% in Condition 3B and 41.08% in Condition 5B.⁴ This suggests that Consumer players in the latter two conditions had a higher tendency to attempt strategic waiting by not making a purchase even when the current price was lower than their valuation.

Finally, the mean payoffs also suggest significant deviations from equilibrium predictions, as shown in Table 5. Note that, for more meaningful comparison between conditions, the payoff variables for the Retailers in Conditions 3B and 5B are normalized by the numbers of Consumers in those conditions (three and five, respectively). The table also indicates within-session evolution, which will be explored further in Sect. 4.3. With that in mind, we focus on the latter half of the session, in which Retailer 1 (Retailer 2) earned significantly less (more) than in equilibrium in Conditions 3B and 5B, but not 1B; the Consumer, meanwhile, earned significantly less than in equilibrium in Condition 3B in the latter half of the session, but not so in the other two conditions.

4.2 Analysis of deviations from best responses

To further understand these findings, we next report Fig. 3, which displays the mean deviations of Retailer prices relative to best response prescribed by equilibrium analysis. Note that, for period 1 pricing in the figure, we display the equilibrium price as the "best response". Figure 3 shows that the Retailers largely overpriced with respect to best response in the earlier half of the session; however, as the session

⁴ Note that no Consumer player exhibited completely myopic behavior, with which this measure would be 0% for the player. Moreover, this measure is less than 20% for only 14 out of the 149 Consumer players across conditions.

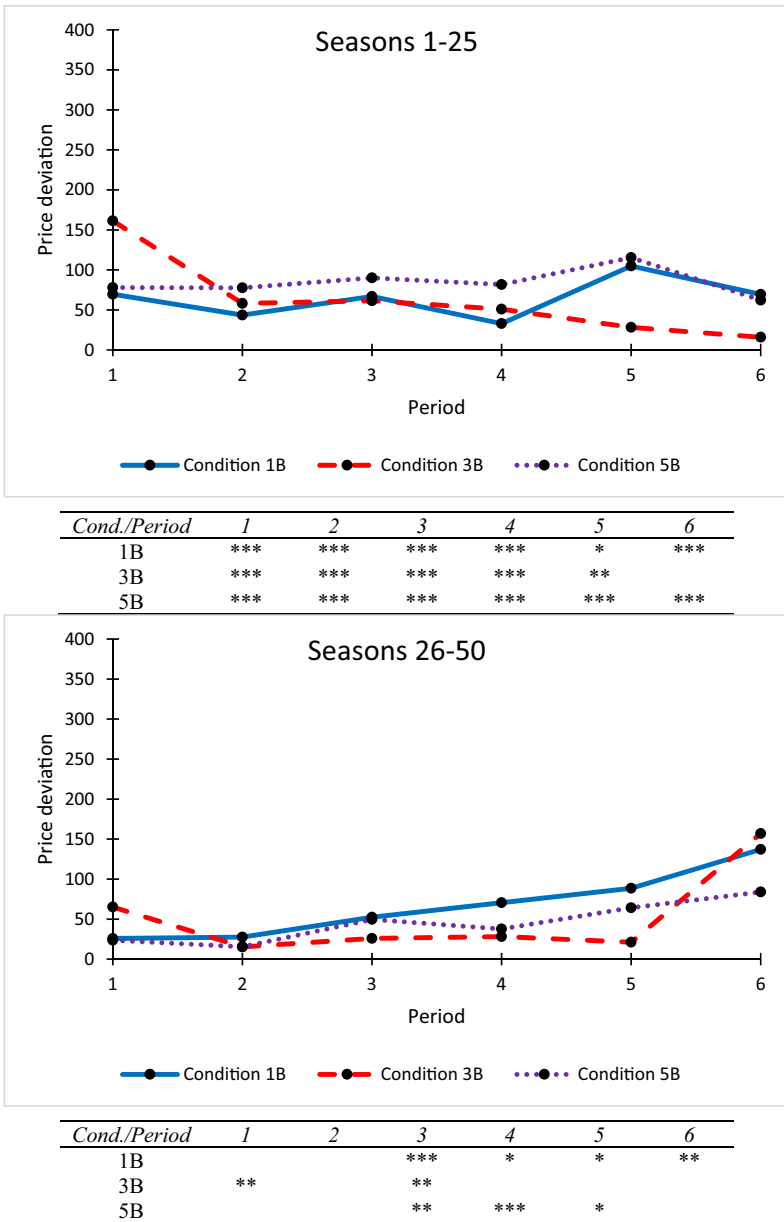


Fig. 3 Mean deviation of experimental prices from best response up to period 6. The unit of analysis is the session. The accompanying tables indicate, with one or more asterisks, the mean deviations by period that are significantly different from zero according to an intercept-only regression analysis clustering variances by subjects ($*p < 0.05$, $**p < 0.01$, $***p < 0.001$)

progressed, Retailers seemed to have been driven by competitive pressure to adhere more to equilibrium analysis prescriptions. This is especially the case with Retailer 2 in period 2 in response to Retailer 1's period 1 price.

We next compare the Consumers' purchase decisions against best response based on equilibrium analysis derived from the equilibrium and out-of-equilibrium results in Table 1. Figure 4 displays the mean deviations of Consumer purchase decisions from best response. We code decisions such that a purchase decision is coded as 1 and a no-purchase decision is coded as 0, so that a purchase decision in the data when equilibrium analysis prescribes no purchase leads to a deviation of 1, while a no-purchase decision in the data when equilibrium analysis prescribes purchase leads to a deviation of -1 ; if a decision in the data is the same what is prescribed by equilibrium analysis, the deviation is 0. Figure 4 shows that purchase decisions in period 1, compared with other periods, seem to adhere more to strategic benchmarks in the form of equilibrium analysis prescriptions, especially in the latter half of the session. But otherwise, Consumers often made purchases when equilibrium analysis suggested that they should not purchase. This type of behavior finds an analogy in previous research on sequential search behavior (e.g., Mak et al., 2019; Zwick et al., 2003), where subjects often stopped to purchase when the optimal strategy was to continue searching.

4.3 Difference between conditions and within-session evolution

Our final set of analysis involves a direct comparison between conditions and examination of within-session evolution. The analysis involves regressions with results that are displayed in Table 6.

The first panel of Table 6, regarding the payoffs, indicates that the theoretical invariance with respect to the number of consumers was violated. For example, focusing on the latter half of the session (i.e., when $\Delta_{1st_half} = 0$ following the terminology for Table 6), Retailer 1's payoffs in Condition 1B was less than its payoffs in the other two conditions, even after normalization by the number of Consumers. Retailer 2, on the other hand, earned a higher payoff in Condition 3B than in Condition 1B.

However, the second panel of Table 6 shows that the Retailers' prices were not significantly different across conditions in the latter part of the session, suggesting that the violation of invariance might not be attributable to prices.

Rather, as evidenced in the third panel of Table 6, Consumers seemed to be more prone to attempt strategic waiting in period 1 in Conditions 3B and 5B than in Condition 1B, at least in the latter half of the session. This is consistent with earlier results as in Fig. 4. This could help to explain Retailer 1's lower normalized payoffs in Conditions 3B and 5B relative to Condition 1B. The panel also shows a lack of corresponding differences in strategic waiting tendency across conditions in period 2. That is, while a higher number of Consumers might prompt Consumers to attempt strategic waiting in period 1, this tendency did not persist into period 2—and beyond, as our analysis for other periods also suggests.

Another analysis approach, summarized in the fourth panel of Table 6, provides corroborating evidence. If Consumers attempted more strategic waiting in


Conditions 3B and 5B than in Condition 1B, there should be correspondingly more instances per Consumer journey where the Consumer's valuation was higher than or equal to the current price. This is because the Consumer became more likely to attempt strategic waiting in such instances and continue the journey without making a purchase. The middle column results of Table 6's fourth panel indeed shows that there was a larger number of instances of such—which we denote as $n_{v \geq p}$ —in Condition 3B (at least in the second half of the session) and Condition 5B, compared with Condition 1B. It is also interesting to see whether, conditioned on $n_{v \geq p}$ being positive, the last instance where “valuation \geq price” was a myopic purchase, which is defined as making a purchase when the best response should be the opposite. The logistic regression results summarized in the right column in Table 6's fourth panel essentially show that, across condition, myopic purchase became more likely the larger number of times the Consumer had attempted strategic waiting in the same Consumer journey. This seeming inconsistency could be due to a simple heuristic that the Consumer only allowed themselves to “miss making a profit” (i.e., waiting despite valuation \geq price) a limited number of times within a Consumer journey; more research might be warranted to find out the in-depth behavioral reasons behind.

Finally, Table 6's fifth panel provides further evidence on Consumer purchase that is consistent with the above findings. The analysis shows that, in the latter half of the session, the higher a Consumers' valuation, the more likely they were to make a purchase in a session, and conditioned on a purchase being made, the lower the number of periods at which the purchase occurred. Moreover, the likelihood to purchase was lower in Condition 3B than in Condition 1B controlling for the valuation, and the purchase period in Conditions 3B and 5B relative to 1B was higher controlling for the valuation, suggesting a higher tendency to attempt strategic waiting in those conditions.

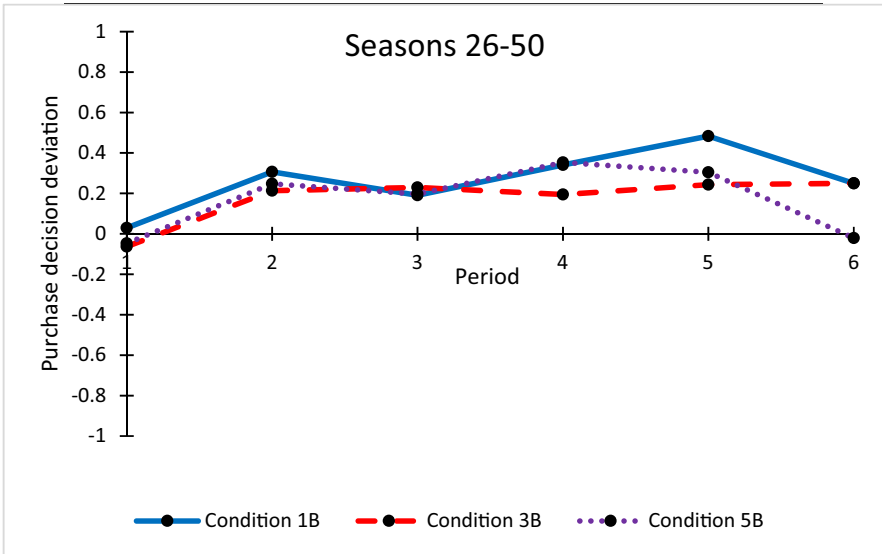
Table 6's panels also illustrate clearly that there was within-session evolution. For example, consistent with Fig. 3, Table 6's panel 2 shows that period 1 prices were overall driven down in the latter half of the session. In parallel, Table 6's panels 3 and 5 (and panel 4 also with partial evidence for Condition 3B), show that there was an increased tendency for strategic waiting among Consumers as the session progressed.⁵

Table 7 helps to drive home the insights from the Table 6 regression results on Consumer decisions. Table 7 shows the frequency distributions of Consumer

⁵ In addition, we examine, for Conditions 3B and 5B only, whether the Retailers' pricing decision in a period t ($t > 1$) depended on C_t , the number of Consumers who were still in the game (i.e., had not yet purchased) at the beginning of period t . Specifically, in parallel with the approach in Table 6's second panel, for the data in these two conditions, we regress p_t on p_{t-1} , C_t , $C_t \times \Delta_{1st_half}$, and $C_t \times \Delta_{1st_half} \times \Delta_{5B}$, for $t = 2, 3, 4, 5, 6$ separately. Focusing on the statistically significant ($p < 0.05$) coefficients related to C_t , we find that p_2 has a negative dependence (estimate = -14.96, s.e. = 4.96, $p < 0.01$) on C_2 and a positive dependence (estimate = 14.10, s.e. = 5.18, $p < 0.01$) on $C_2 \times \Delta_{1st_half} \times \Delta_{5B}$, while p_5 has a positive dependence (estimate = 31.15, s.e. = 9.55, $p < 0.01$) on $C_5 \times \Delta_{1st_half} \times \Delta_{5B}$. All other coefficients related to C_t are non-significant. Focusing on the latter half of the session, we conclude that Retailer 2 in Conditions 3B and 5B had a general tendency to price lower in period 2 the larger the number of Consumers that were still in the game, after controlling for the period 1 price.



<i>Cond./Period</i>	1	2	3	4	5	6
1B		***	***	***		**
3B	***	***	***	***	***	
5B		***	***	***	***	**



<i>Cond./Period</i>	1	2	3	4	5	6
1B		***	***	***	**	
3B	*	***	***	***	*	**
5B	*	***	***	***	***	

Fig. 4 Mean deviation of experimental Consumer purchase decisions from best response according to equilibrium analysis. The unit of analysis is the session. The accompanying tables serve the same function with the same notation as those for Fig. 3

Table 6 Regression analysis results. Variances are clustered by subjects. Standard errors in parentheses. The QIC goodness-of-fit statistic is also listed. Where the estimated coefficient is significantly different from zero, the coefficient is indicated by one or more asterisks (* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$). $\Delta_{3B} = 1$ when the data point is from Condition 3B and $\Delta_{3B} = 0$ otherwise; $\Delta_{5B} = 1$ when the data point is from Condition 5B and $\Delta_{5B} = 0$ otherwise. In addition, $\Delta_{1st_half} = 1$ when the data point is from Seasons 1–25 and $\Delta_{1st_half} = 0$ otherwise. Further specific details for each panel are provided in the respective headings

1. Dependent variables: payoffs (normalized for the Retailers; see Table 5) (linear regression).

^a $p = 0.060$; ^b $p = 0.058$.

	<i>Retailer 1 payoff</i> <i>normalized</i>	<i>Retailer 2 payoff</i> <i>normalized</i>	<i>Consumer payoff</i>
Intercept	99.14*** (7.16)	26.42*** (2.77)	318.70*** (10.99)
Δ_{3B}	-29.20*** (8.26)	11.97** (3.89)	-28.19 ^b (14.86)
Δ_{5B}	-26.62* (10.42)	7.12 (4.47)	-2.11 (13.15)
Δ_{1st_half}	-6.18 (6.07)	9.99** (3.69)	-34.88* (14.52)
$\Delta_{1st_half} \times \Delta_{3B}$	47.21*** (10.19)	-10.66 ^a (5.66)	-35.86* (17.84)
$\Delta_{1st_half} \times \Delta_{5B}$	20.42* (8.95)	-7.23 (5.01)	-3.79 (16.53)
QIC	2776	2758	7450
<i>N</i>	2748	2748	7440

2. Dependent variables: prices in periods 1 and 2 (linear regression). Regressions for prices in periods 3 to 6, using similar models as that for the period 2 price, reveal statistically significant (at $p < 0.05$) positive coefficients for p_{t-1} but no statistically significant coefficients otherwise.

	<i>Period 1 price (p_1)</i>	<i>Period 2 price (p_2)</i>
Intercept	239.59*** (20.16)	7.50 (16.67)
p_{t-1}	–	0.52*** (0.04)
Δ_{3B}	39.56 (29.53)	-0.05 (21.83)
Δ_{5B}	-1.68 (31.31)	2.20 (17.68)
Δ_{1st_half}	44.32*** (11.83)	19.86 (14.75)
$\Delta_{1st_half} \times \Delta_{3B}$	51.87** (18.42)	4.50 (16.07)
$\Delta_{1st_half} \times \Delta_{5B}$	9.78 (19.63)	28.22 (16.40)
QIC	2848	1558
<i>N</i>	2748	1510

journeys distinguished by whether the Consumer made one of four types of ex post suboptimal decisions in the season (i.e., decisions that were suboptimal vis a vis the ex post observations of the full set of posted prices by Retailers in a season after the season was over):

Table 6 (continued)

3. Dependent variables: Consumers' binary purchase decisions in periods 1 and 2 (logistic regression with 1 = purchase, 0 = otherwise). p_t is the price in period t for the regression analysis of period t decisions, $t = 1, 2$. Regressions on decisions in periods 3 to 6 using similar models reveal statistically significant ($p < 0.05$) positive (negative) coefficients for p_t and v for periods 3 to 5, but otherwise no statistically significant coefficients in general.¹

	<i>Period 1 decision</i>	<i>Period 2 decision</i>
Intercept	-0.96*** (0.14)	-0.30** (0.098)
v	0.0092*** (0.0010)	0.0058*** (0.0005)
p_t	-0.0059*** (0.0003)	-0.0091*** (0.0014)
$p_t \times \Delta_{3B}$	-0.0037** (0.0011)	-0.0020 (0.0013)
$p_t \times \Delta_{5B}$	-0.0028** (0.0011)	-0.0025 (0.0013)
$p_t \times \Delta_{1st_half}$	-0.0011* (0.0005)	-0.0008 (0.0010)
$p_t \times \Delta_{1st_half} \times \Delta_{3B}$	0.0054*** (0.0009)	0.0022 (0.0015)
$p_t \times \Delta_{1st_half} \times \Delta_{5B}$	0.0029*** (0.0008)	0.0013 (0.0013)
QIC	5988	3206
N	7440	3358

¹ The exceptions are limited to the regression for the decisions in period 5, with statistically significant ($p < 0.05$) coefficients for $p_5 \times \Delta_{1st_half}$ (negative), $p_5 \times \Delta_{1st_half} \times \Delta_{3B}$ (positive), and $p_5 \times \Delta_{1st_half} \times \Delta_{5B}$ (positive). In the context of the regression results, these coefficients might indicate transient differences across conditions early in the session.

Table 6 (continued)

4. Dependent variables: (1) number of instances, $n_{v \geq p}$, in a Consumer journey in which the Consumer's valuation was higher than or equal to the current price (linear regression); (2) conditioned on the number of instances being positive (i.e., $n_{v \geq p} > 0$), whether the last instance in the Consumer journey was a myopic purchase (logistic regression). v is the valuation in the season.

	<i>Number of instances of valuation \geq price in a Consumer journey ($n_{v \geq p}$)</i>	<i>Whether the last instance of valuation \geq price in a Consumer journey was a myopic purchase, conditioned on $n_{v \geq p} > 0$</i>
Intercept	0.80*** (0.045)	-1.04*** (0.18)
$n_{v \geq p}$	–	0.58*** (0.07)
v	0.0006*** (0.0000)	-0.0013*** (0.0002)
Δ_{3B}	0.19*** (0.05)	-0.039 (0.14)
Δ_{5B}	0.17*** (0.05)	-0.0004 (0.13)
Δ_{1st_half}	0.034 (0.044)	0.12 (0.13)
$\Delta_{1st_half} \times \Delta_{3B}$	-0.22*** (0.06)	0.68*** (0.19)
$\Delta_{1st_half} \times \Delta_{5B}$	-0.07 (0.06)	0.23 (0.15)
QIC	7465	7594
N	7440	6310

5. Dependent variables: (1) whether a Consumer made a purchase or not in a season (logistic regression with 1 = purchase, 0 = otherwise); (2) conditioned on making a purchase in a season, the period in which the purchase was made (linear regression). v is the valuation in the season.

	<i>Whether a purchase was made</i>	<i>Period of purchase conditioned on a purchase being made</i>
Intercept	-0.20 (0.12)	2.71*** (0.072)
v	0.0035*** (0.0001)	-0.0021*** (0.0001)
Δ_{3B}	-0.54*** (0.15)	0.32*** (0.073)
Δ_{5B}	-0.22 (0.14)	0.13* (0.067)
Δ_{1st_half}	-0.52*** (0.16)	0.14* (0.056)
$\Delta_{1st_half} \times \Delta_{3B}$	0.22 (0.19)	-0.061 (0.076)
$\Delta_{1st_half} \times \Delta_{5B}$	0.17 (0.18)	-0.010 (0.068)
QIC	7709	5282
N	7440	5263

The exceptions are limited to the regression for the decisions in period 5, with statistically significant ($p < 0.05$) coefficients for $p_5 \times \Delta_{1st_half}$ (negative), $p_5 \times \Delta_{1st_half} \times \Delta_{3B}$ (positive), and $p_5 \times \Delta_{1st_half} \times \Delta_{5B}$ (positive). In the context of the regression results, these coefficients might indicate transient differences across conditions early in the session

^a $p = 0.060$; ^b $p = 0.058$

Table 7 Percentages of Consumer journeys distinguished by the ex post (sub)optimality of the purchase timing (absolute frequencies in parentheses) (see Sect. 4.3)

	Ex post suboptimal				Ex post optimal	
	Purchase			No-purchase	Purchase	No-purchase
	Payoff < 0	Payoff ≥ 0, “too early”	Payoff ≥ 0, “too late”			
Condition 1B	0.08 (1)	NA	0 (0)	12.42 (149)	74.17 (890)	13.33 (160)
Condition 3B	0.22 (5)	26.40 (594)	0.18 (4)	16.58 (373)	39.82 (896)	16.80 (378)
Condition 5B	0.58 (23)	35.94 (1434)	0.03 (1)	14.26 (569)	35.46 (1415)	13.73 (548)

The total count in each condition is: Condition 1B: 1200; Condition 3B: 2250; Condition 5b: 3990

- (1) Purchase with negative payoff;
- (2) Purchase with non-negative payoff but ex post “too early,” which denotes purchasing in a period after which there is at least one period in the same season when the price is lower;
- (3) Purchase with non-negative payoff but ex post “too late,” which denotes purchasing at a price that is higher than previous prices in the same season⁶;
- (4) No-purchase (that is ex post suboptimal), which denotes not purchasing before the season terminates although at least one price during the season is below the Consumer’s valuation.

In addition, Table 7 shows the frequency distributions of Consumer journeys in which the Consumer made one of two types of ex post optimal decision:

- (1) Purchase at a price that turns out to be the lowest in the season and is not higher than the Consumer’s valuation;
- (2) No-purchase (that is ex post optimal), which denotes not purchasing in a season in which all the prices are higher than the Consumer’s valuation.

Table 7 highlights how Conditions 3B and 5B made possible a substantial number of observations of ex post “too early” purchases that were not possible by design in Condition 1B (as there was only one Consumer in the market in Condition 1B). At the same time, there were lower proportions of ex post optimal purchases in Conditions 3B and 5B, as the season could continue often at lower prices after a Consumer had purchased. All these instances could motivate Consumers to attempt more strategic waiting. The difference in observations between ex post “too early” and ex post “too late”, in Conditions 3B and 5B, was such that there were about twice as many observations of ex post “too early” as ex post “too late”. The observations might have motivated Consumer players to behave more patiently and to

⁶ Since prices often declined during the season, there were few instances of such suboptimal behavior.

attempt to strategically delay purchase compared to the Condition 1B, where only ex post “too late” observations were possible.

5 Conclusions

In this research, we introduce and test a stylized model of dynamic pricing under duopolistic competition. Our model extends previous work on zigzag competition (MGG and MRG, as discussed earlier) in which a consumer receives alternating price offers between two retailers over an indefinite horizon with random continuation probability, or equivalently, infinite horizon with identical per period time discount factor across players. Our equilibrium analysis suggests that price offers decrease exponentially across periods over the season. Moreover, the theoretical predictions are invariant with respect to the number of consumers in the game, as long as their valuations are ex ante independently and identically distributed.

Our experiment demonstrated significant deviations from equilibrium predictions that exhibited consistent and different characteristics from the related previous results from MRG’s finite-horizon experimental study. In our experiment, Retailers often overpriced relative to equilibrium benchmarks, a result that is similar to MRG’s results. However, overpricing in MRG persisted throughout the experimental session while it became mitigated in our indefinite-horizon case.

Moreover, the theoretical invariance with respect to the number of consumers was violated (a research objective that MRG did not study). As the number of Consumers increased, Consumers in the experiment became more prone to strategic waiting and exhibited less myopic behavior. This led to a decrease in the per-Consumer payoff of Retailer 1 in Conditions 3B and 5B compared with 1B. However, the increase in strategic waiting did not seem to extend beyond period 1, and a corresponding increase in per-consumer payoff of Retailer 2. Moreover, despite such *relative* increase in strategic waiting in period 1 across conditions, Consumers were generally not sufficiently strategic when compared to theoretical benchmarks. Hence, similar to the Strategic conditions in MRG, we find that the second Retailer generally earned more than predicted by equilibrium analysis, while the first Retailer and the Consumers generally did worse than predicted by equilibrium analysis.

Such findings are in line with the second possible reason discussed in the Introduction section of this paper regarding why the theoretical invariance with respect to the number of consumers could be violated behaviorally. That is, as the number of Consumers increased, players expected that the market would take longer to clear and the season would last longer, resulting in an increased tendency for strategic behavior. Our experimental findings also shed insights into the validity of the other reasons discussed. For example, the first reason discussed in the Introduction is the possibility that Consumers were heterogeneous in terms of how myopic they were. As the number of Consumer increased across conditions, the probability that at least one Consumer was myopic increased, and Retailers might adjust their pricing strategies accordingly. If this adjustment approach was indeed practiced by the Retailers, we might expect to see increases in prices as the number of Consumers increased—which, however, is not supported by our

data analysis (e.g., the second panel of Table 6). Another reason suggested in the Introduction is related to the possibility that Consumers might be affected by social comparison effects when there were multiple of them, such as fairness concerns and socially induced regret (when, e.g., another Consumer purchased at a better price). Retailers anticipating these concerns might adjust their pricing strategies such that prices across periods and Retailers might be more equal as the number of Consumers increased. However, we do not find evidence in support this possibility. For example, the second panel of Table 6 suggests that the difference between prices in periods 1 and 2 did not decrease as the number of Consumers increased.

Lastly, our analysis suggests evidence of within-session evolution, in that the Retailers' overpricing tendency decreased over the experimental session. This is consistent with competitive pressure driving down prices that got closer to equilibrium predictions. Correspondingly, there was an increased tendency for strategic waiting among Consumers as the session progressed.

Our work can be extended in a number of directions. Further experimentation could be carried out that elicits subject beliefs and/or decision rationale in order to estimate more detailed behavioral models for decisions, test further behavioral predictions, and, in general, find out more about the behavioral factors behind the deviations from theoretical predictions. Another possible direction pertains to the information structure, a direction that MRG also explored. In particular, MRG replaced their main modelling assumption that once a price is posted it is observed by the consumer and the two retailers, by the assumption that Retailer 1 never observes the price offers made by Retailer 2 and vice versa. The effect of this change from complete to incomplete information is quite dramatic, as MRG have proved that if information about prices is restricted in this sense, trade occurs at only zero price in equilibrium. In effect, competition between the two retailers is highly intensified when they cannot monitor each other's prices, and the result becomes the classical Bertrand outcome. An implication of this result is that both retailers would be willing to pay cash in order to gain full transparency of the market. In the present model, because the horizon does not have a fixed finite number of periods as in MRG, it can be expected that the equilibrium analysis will not result in a state of Bertrand outcome with all profits competed away. In yet another direction, the zig-zag competition assumption can be relaxed so that the consumer(s) could switch between retailers with a probability that is between 0 and 1; a first step could be to model this probability as exogenous and constant as in a simple first order Markov process (see MGG). Lastly, to study and isolate specific responses, Retailer or Consumer subjects may be replaced with computerized robots playing different strategies (equilibrium or other). Further theoretical and experimental explorations might reveal more insights into competitive dynamic pricing behavior.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10683-024-09848-8>.

Acknowledgements Eyran J. Gisches, Vincent Mak, and Rami Zwick would like to honor Amnon Rapoport with this paper and express their gratitude to Amnon's advice, collaboration, and friendship over the years. Gisches, Mak, and Zwick are responsible for any errors in the paper.

Funding This study was supported by University of California, Riverside.

Declarations

Competing interests The authors declare that they have no conflicts of interest with respect to their authorship or the publication of this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Baucells, M., Osadchiy, N., & Ovchinnikov, A. (2017). Behavioral anomalies in consumer wait-or-buy decisions and their implications for markdown management. *Operations Research*, *65*(2), 357–378.
- Chen, J., & Zhao, Y. (2020). High price or low price? An experimental study on a markdown pricing policy. *European Journal of Operational Research*, *284*(1), 240–254.
- Coase, R. H. (1972). Durability and monopoly. *Journal of Law and Economics*, *15*, 143–149.
- Kahn, B. E. (1995). Consumer variety-seeking among goods and services: An integrative review. *Journal of Retailing and Consumer Services*, *2*(3), 139–148.
- Kimes, S. E. (1989). Yield management: A tool for capacity-considered service firms. *Journal of Operations Management*, *8*(4), 348–363.
- Kremer, M., Mantin, B., & Ovchinnikov, A. (2017). Dynamic pricing in the presence of myopic and strategic consumers: Theory and experiment. *Production and Operations Management*, *26*(1), 116–133.
- Lilien, G. L., Kotler, P., & Moorthy, K. S. (1992). *Marketing models* (Vol. 803). Prentice-Hall.
- Mak, V., Rapoport, A., & Gisches, E. J. (2012). Competitive dynamic pricing with alternating offers: Theory and experiment. *Games and Economic Behavior*, *75*, 250–264.
- Mak, V., Rapoport, A., Gisches, E. J., & Han, J. (2014). Purchasing scarce products under dynamic pricing: An experimental investigation. *Manufacturing & Service Operations Management*, *16*, 425–438.
- Mak, V., Seale, D. A., Rapoport, A., & Gisches, E. J. (2019). Voting rules in sequential search by committees: Theory and experiments. *Management Science*, *65*(9), 4349–4364.
- Mantin, B., Granot, D., & Granot, F. (2007). Dynamic pricing under first order Markovian competition. *Naval Research Logistics*, *58*(6), 608–617.
- McAfee, R. P., & Wiseman, T. (2007). Capacity choice counters the Coase conjecture. *Review of Economic Studies*, *75*, 317–332.
- Osadchiy, N., & Bendoly, E. (2015). Are consumers really strategic? Implications from an experimental study. Available at SSRN: <https://ssrn.com/abstract=259184>.
- Phillips, R. L. (2005). *Pricing and revenue management*. Stanford University Press.
- Rapoport, A., Erev, I., & Zwick, R. (1995). An experimental study of buyer-seller negotiation with one-sided incomplete information and time discounting. *Management Science*, *41*(3), 377–394.
- Roth, A. E., & Murnighan, J. K. (1978). Equilibrium behavior and repeated play of the prisoner's dilemma. *Journal of Mathematical Psychology*, *17*(2), 189–198.
- Talluri, K. T., & van Ryzin, G. (2004). *The theory and practice of revenue management*. Kluwer Academic Publishers.
- Zwick, R., Rapoport, A., & Howard, J. C. (1992). Two-person sequential bargaining behavior with exogenous breakdown. *Theory and Decision*, *32*(3), 241–268.
- Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. V. (2003). Consumer sequential search: Not enough or too much? *Marketing Science*, *22*(4), 503–519.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Amnon Rapoport¹ · Eyran J. Gisches²  · Vincent Mak³  · Rami Zwick¹ 

✉ Rami Zwick
rami.zwick@ucr.edu

Eyran J. Gisches
eyran@email.arizona.edu

Vincent Mak
v.mak@jbs.cam.ac.uk

¹ School of Business, University of California, Riverside, Riverside, CA 92521, USA

² Department of Management Information Systems, Eller School of Management, University of Arizona, Tucson, AZ 87521-0108, USA

³ Cambridge Judge Business School, University of Cambridge, Trumpington Street, Cambridge CB2 1AG, UK