Chapter 1

Cheap talk in Operations: Role of intentional vagueness

Gad Allon
Kellogg School of Management,
2001 Sheridan Road Evanston,
IL 60208, g-allon@kellogg.northwestern.edu

Achal Bassamboo
Kellogg School of Management,
2001 Sheridan Road Evanston,
IL 60208, a-bassamboo@kellogg.northwestern.edu

1.1 Introduction

Provision of real-time information by firms to their customers has become prevalent in recent years in both the service and retail sectors. Service providers use announcements to inform customers about anticipated delays, whereas retailers provide the customers with information about the inventory level and the likelihood of a stock-out. Often, this information cannot be credibly verified by the customers. The question of which information should the firm share with its customers is a complex one, and its answer depends among other things on the dynamics of the underlying operations and the customer behavior.

Most of the Operations Management literature addressing this issue analyzed two categories of information provided to the customer: (i) full information - the state of the system, as known to the system manager when the customer arrives, and (ii) no information - where no information is provided, and customers must base their decisions on their expectation regarding the system performance. The main assumption made in the former category of literature is that customers treat the information provided regarding the state of the system as a-priori verified (i.e., credible), and act accordingly in making their decisions. The two main issues with this assumption are: (i) Customers are seldom naive in their attitude towards any information provided by interested parties, and thus take such announcements with a “grain of salt”. Moreover, under the assumption of “naivety,” it makes sense for the firm to deviate from the truth-telling policy. The option that the firm might *lie*, given that the customer always believes the firm, is never explored in the literature. (ii) Further, prior work implicitly assumes that the announcements have a literal meaning in terms of the availability (in retail) or delay (in services) or average waiting time. However, as stated above, many service providers use verbal messages that need to be further processed in order for customers to make the decision. For example, without processing, it is not clear what “high volume of calls” or “almost gone” mean in terms of delay in the system (in services) and availability of the product (in retail) in these commonly used statements. This problem is clearly a consequence of the first issue since, without processing, only announcements with literal meaning are possible. The combination of these two issues contributed to the fact that only simple (i.e., no-information or full-information) announcements were discussed, while in practice we observe a much richer variety of announcements.

This chapter surveys models that address these issues. In particular, the customers in these models treat information provided by the service provider as unverified and non-binding. These models, thus, treat customers as strategic in the way they process information, as well as in making the decisions (that is, in service settings whether to join or balk, and whether to buy or wait in retail), and the firm as strategic in the way it provides the information. The customers and the firm are assumed to be self-interested in making their decisions: the firm in choosing which announcements to make and the customers in interpreting these and making the decisions. Note that, while previous models assumed customers to be strategic in the way they make decisions (being forward-looking) or in the way they form expectations, these models are the first to study settings in which customers are strategic in the way they *interpret* information provided by other parties. That is, customers do not take the messages or the information provided by the firm at their “face value.”

This allows us to characterize the equilibrium language that emerges between the firm and its customers. By doing so, not only do these models relax the assumption that customers are naive in their treatment of the announcements, but also demonstrate that many of the commonly used announcements arise in equilibrium. For example, in services, the spectrum of possible equilibria will range from announcements that are analogous to the verbal type, describing the volume of arriving customers as high or low to the detailed waiting time announcements, both common in service systems. In retail settings, it is shown that an informative language is not possible between a single retail and its customers. These models are among the first to show that the spectrum of announcements that exists in real-world applications can emerge as an equilibrium of a game between the provider and her customers.

This chapter surveys the emerging literature that deals with the strategic nature of the information transmission in a practical operational setting, where unverifiable, non-committal, real-time information is provided by a self-interested firm to selfish customers.

In this literature, the announcements made by the system manager is modelled as "cheap talk," i.e., pre-play communication that carries no cost. Cheap talk consists of costless[1], non-binding, non-verifiable messages that may affect the customer's beliefs. It is important to note that while providing the information does not *directly* affect the payoffs, it has an indirect implication through the customer's reaction and the equilibrium outcomes. The information has no impact on the payoffs of the different players per se, i.e., the payoffs of both sides depend only on the actions taken by the customer and queueing dynamics. This, in turn, means that if the customer does not follow the recommendation made by the firm, he is not penalized, nor is he rewarded when he follows them. However, as it will be shown, the announcements do have an impact on the service provider's profits and the customers' utility, in equilibrium. This is in agreement with both the cheap talk literature (See Crawford and Sobel (1982)) and the operations management literature with strategic customers (See Naor (1969) for a queueing application and Aviv and Pazgal (2007) for a retail application, where the information provided to the customer in the form of full visibility of the state of the system does not alter the customers utility directly; however, it allows him to make a knowledgeable decision and thus affects his utility in a indirect manner).

The focus of these models is dealing with the *strategic* interaction between the customer and the firm in a setting in which their incentives are *misaligned*, when *unverifiable, costless, and non-binding* information is provided to the customer. In

[1]We assume that the cost associated with conveying the message is negligible. In most practical service organizations, while the provider needs to incur fixed costs, for example, by investing in a more sophisticated IT infrastructure to learn the state of the system, the marginal cost of providing the information to the customer is insignificant. There is a voluminous literature starting with Spence (1973) dealing with models where signaling is not costless, and the mere fact that players are willing to incur a cost provides a signal.

all of the instances described in this chapter, the information is always unverifiable and has no contractual bearing. This is in contrast to service-level *guarantees*, such as those made by Dominos Pizza, Ameritrade, and E$*$trade to name a few, where the commitment is both contractually binding and verifiable.

A reading guide. The next section reviews the classical cheap talk model introduced by Crawford and Sobel (1982). We discuss the challenges one faces in developing a framework that echoes the classical cheap talk model for dynamic operational settings. Section 1.3 describes the cheap talk game in a service setting, and Section 1.4 describes the cheap talk game in retail.[2] These sections are almost independent and can be read in any order. Section 1.5 summarizes the finding in the previous section and contrasts the equilibrium language in the queueing with the retail one. We conclude the chapter by surveying related literature and future direction.

1.2 Classical Cheap Talk Game

In this section, we provide an overview of the cheap talk game introduced in Crawford and Sobel (1982). This a game played between a *sender* who has some private information and a *receiver* who takes the action which impacts the payoff of both players. We next define the game and highlight the key findings.

1.2.1 Model

The game proceeds as follows: The Sender observes the state of the world, which we shall denote by Q, which is private information and is uniformly distribution on the unit interval. The Sender then sends a signal (or a message) denoted by $m \in \mathcal{M}$. (Here $\mathcal{M}$ denotes the set of all signals that can be used by the Sender.) The Receiver processes this information and chooses an action y which determines the players payoff. The Sender obtains a utility which depends on: (a) the action taken by the Receiver, y; (b) the state of the world Q and (c) his bias which we denote by b, and is given by $V(y, Q, b) = -(y - (Q + b))^2$. The Receiver, on the other hand, obtains a utility which depends only on: (a) his own action y, and (b) the state of the world, Q, and is given by $U(y, Q) = -(y - Q)^2$. [3]

[2] All the proofs of the results in Sections 1.3 and 1.4 are in Allon et al. (2007) and Allon and Bassamboo (2008), respectively.

[3] We adopt a notation that is different from the one used in Crawford and Sobel (1982). This is done in order to be consistent with the notation developed in the model used in the latter part of the chapter. For instance, Q, which denotes the state of the world, would correspond to the queue length in services, and the quantity-on-hand for retail.

The Markov Perfect Bayesian Nash equilibrium of the above game requires that (a) The Sender's signaling rule yields an expected-utility maximizing action for each of the state of the world Q, fixing the action rule for the Receiver; and (b) The Receiver responds optimally to each possible signal using Bayes' Rule to update his prior, taking into account the Sender's Signaling rule and the message/signal received from the Sender.

1.2.2 Key Results

For this classical cheap talk game, there always exists an equilibrium where *no information* is transmitted from the Sender to the Receiver, irrespective of the parameters of the problem. In fact this is the only equilibrium of the game when the bias b exceeds $1/4$. However, when b is less than $1/4$, informative equilibria exist. All these equilibria share the same structure that they partition the state-space (i.e., the unit interval) into finite number of intervals. On each of these intervals the Sender uses the same message. Further, they show that the number of intervals is bounded from above by an integer which is a function of the bias and is denoted by $N(b)$. The equilibrium where the sender partitions the state-space into exactly $N(b)$ partitions is referred to as the *most informative equilibrium*. Further, it is shown that among all the equilibria, both the Sender and Receiver are better off in expectation under the most-informative equilibrium.

1.2.3 Other Applications of the Classical Cheap Talk Model

A variety of papers study mixed-motive economic interaction involving private information and the impact of cheap talk on the outcomes. Farrell and Gibbons (1989) study cheap talk in bargaining; in political context cheap talk has been studied in multiple papers including Austen-Smith (1990), and Matthews (1989). A recent paper by Ren et al. (2007) studies a cheap talk game where a retailer shares forecast information with a supplier. These models almost exclusively focus on static environments. In operational systems information, transmission which is typically done in real-time, cannot be categorized in the classical model and the dynamic environment is, in general, multidimensional and complex.

1.2.4 Discussion

The framework used in this chapter echoes the cheap-talk model proposed in Crawford and Sobel (1982). Driven by the applications in operations, the models have two novel features: first, the game is played with multiple receivers (customers) whose actions have externalities on other receivers; and second, the stochasticity of the state-of-the-world (i.e., the state of the system) is not exogenously given but is determined endogenously. In particular, the private information in these model (for example, the queue length or the inventory position at any given time in service and retail setting, respectively) is driven by the system dynamics, which in turn depend on the equilibrium strategies regarding the information and actions of both the firm and the customers. As we shall see, this multiplicity of receivers with externalities and the endogenization impacts both the nature of the communication as well as the outcome for the various players. This endogeneity, which is crucial for modeling operational setting with customers interaction, is absent in the previous cheap talk literature.

To highlight the impact of the system dynamics, note that there are two types of uncertainties faced in these models: (i) Uncertainty regarding the state of the system when a customer arrives, which is a private information held by the service provider. This type of uncertainty exists in Crawford and Sobel's model as well. (ii) Uncertainty regarding the evolution of the system: Even after announcements are made and the customer decides on his action, both the service provider and the customers are exposed to uncertainty regarding the future dynamics. Note that the latter type of uncertainty is not modeled in Crawford and Sobel (1982). Hence, the definition of the equilibrium concept would require solving a dynamic optimization problem.

1.3 Service Application

In this section, we will survey an *endogenized* cheap-talk model which studies the equilibrium language emerging in a service setting. This model is motivated by the prevalence of the practice of informing customers regarding anticipated delays. Call centers often use recorded announcements to inform callers of the congestion in the system and encourage them to wait for an available agent. While some of these announcements do not provide much information - such as the common message, "Due to high volume of calls, we are unable to answer your call immediately," some call centers go as far as providing the customer with an estimate of his waiting time or his place in the queue. In many service systems where the real state of the system is invisible to customers, delay announcements will affect customers' behavior and may, in turn, have significant impact on the system performance.

1.3.1 Model

We consider a service provider, modeled as an M/M/1 system. Customers arrive to the system according to a Poisson process with rate λ. Service times are exponentially distributed with mean $1/\mu$. We assume that $\lambda < \mu$. We assume that all customer are ex-ante symmetric: customers obtain a value R if they are served, and incur a waiting cost that is proportional to the time spent in the system, with a unit waiting cost of c. Thus, a customer arriving to the system obtains the following utility:

$$U(y) = \begin{cases} R - cw & \text{if } y = \text{"join,"} \\ 0 & \text{if } y = \text{"balk,"} \end{cases} \tag{1.1}$$

where y is the decision made by this customer and w denotes its sojourn time in the system. Throughout the paper, we shall assume that $R > \frac{c}{\mu}$, this assumption ensures that in the absence of delays, the service is beneficial to the customer, on average. Clearly, if $R < \frac{c}{\mu}$, no customer will join regardless of the system announcements. When a customer arrives, the system manager has private information regarding the number of customers currently waiting in queue, denoted by the random variable Q. Its distribution will depend on the equilibrium strategies of both the provider and the customers, unlike in the classical cheap talk games where the distribution of the state-of-the-world is exogenous.

We assume that if the customer is satisfied (i.e., he obtains non-negative utility from the transaction), the service provider obtains a positive revenue of $\overline{v}$, while if the customer is dissatisfied (i.e. he obtains a negative utility), the service provider incurs a cost of $-\underline{v}$. Thus the profit function captures the fact that the firm makes higher profit when the customer is satisfied versus when he is not.

Formally, depending on the action taken by the a customer, and his actual sojourn time in the system, the firm obtains the following revenues:

$$\pi(y) = \begin{cases} \overline{v} > 0 & \text{if } y = \text{"join" and } R \geq cw, \\ \underline{v} \leq 0 & \text{if } y = \text{"join" and } R < cw, \\ 0 & \text{if } y = \text{"balk."} \end{cases} \tag{1.2}$$

Such profit functions arise naturally in several settings. One such environment is service processes outsourcing. Typically, the outsourcing firm requires the provider, (for example, a call center) to provide an adequate and timely service to the referred customers. The referring firm then pays the call center only for the satisfied customers and penalizes the provider for the dissatisfied ones. Such a structure will also arise in cases where the firm earns certain revenues from satisfied customers but loses goodwill with every dissatisfied one. Further, we would like to point out that this analysis can be generalized for the setting where the firm's profit from a customer is a monotone decreasing function of the customer's waiting time.

We assume that the customer decides whether to join or not based on the information he can infer from the system manager regarding the current state of the system, denoted by I, in order to maximize its expected utility. Therefore the customer will join, if and only if $R \geq c\mathbb{E}(w|I)$, where I is the information provided to this customer.

Note that the customer's and the service provider's incentives are not completely misaligned: both prefer short waiting times, which result in higher utility for the customer and higher profits for the service provider. At the same time, we observe that the incentives are not perfectly aligned and this would lead to equilibria described in the next section. We refer the reader to Farrell and Rabin (1996) for a discussion of settings in which incentives are perfectly misaligned.

1.3.2 Problem formulation

In this section we formally define the game between the service provider and the customers. The equilibrium concept we employ is one of Markov Perfect Bayesian Nash equilibrium, which is simply a Nash equilibrium in the decision rules that relate agents' actions to their information and to the situation in which they find themselves, allowing for the strategies to depend only payoff-relevant histories. Recall that customers are indistinguishable and their strategies are ex-ante symmetric, both in their interpretations of the signals and in their actions. Let $\mathcal{M} = \{m_1, m_2, \ldots\}$ represent the set of feasible signals that the firm can provide to the customer. We can represent the signaling rule by a function $g : \mathbb{Z} \mapsto \mathcal{M}$, where $g(q) = m$ if the firm uses the signal m when the queue lenght is q. Let $y : \mathcal{M} \mapsto 0, 1$ denote the strategy of the customer, where $y(m)$ is the probability that a customer joins when the firm signals m. Consequently, we interpret $y(m) = 1$ as a "join" decision and $y(m) = 0$ as a "balk" decision and we will use this alternative terminology interchangeably. Note that the above signaling and action rules restrict attention to pure strategies. The requirements of a Markov Perfect Bayesian Nash equilibrium in our context are rather intuitive. Given a signaling rule for the system, customers with an action rule that dictates joining the system when the signal is m will not deviate from this rule if their expected conditional utility, given by $\mathbb{E}[R - c\frac{q+1}{\mu}|g(q) = m]$, will be negative by doing so. Given the customer's action rule $y(m)$, the firm will deviate from its signaling rule $g(q)$ if it maximizes its steady-state profit, i.e, if $g(q)$ solves an appropriate Markov Decision Process (see below) with respect to the action rule $y(m)$. The above is formalized in the following definition.

Definition 1.3.1 (**Markov Perfect Bayesian Nash Equilibrium**) *We say that the signaling rule $g(q)$ and the action rule $y(m)$ constitute a Markov Perfect Bayesian Nash Equilibrium (MPBNE), if they satisfy the following conditions:*

1. *Let $N = \inf\{q : y(g(q)) = 0\}$. Let p_q^N be the steady state probability that*

the number of customers in an M/M/1/N is q[4]*. For each* $m \in \mathcal{M}$*, we have*

$$y(m) = \begin{cases} 1 & \frac{\sum_{\{q:g(q)=m\}}\left[R - c\frac{q+1}{\mu}\right]p_q^N}{\sum_{\{q:g(q)=m\}} p_q^N} \geq 0, \\ 0 & \text{otherwise.} \end{cases}$$

2. *With* $f(j) = (\overline{v} - \underline{v})\mathbb{P}\{W(j+1) \leq \frac{R}{c}\} + \underline{v}$*, there exist constants* $J_0, J_1, \ldots,$ *and* γ *that solve the following set of equations:*

$$\begin{aligned} J_0 &= \max_{m \in \mathcal{M}} \left\{ \frac{f(0)y(m) - \gamma}{\lambda} + J_0(1 - y(m)) + J_1 y(m) \right\} \\ &= \frac{f(0)y(g(0)) - \gamma}{\lambda} + J_0(1 - y(g(0))) + J_1 y(g(0)) \\ J_q &= \max_{m \in \mathcal{M}} \left\{ \frac{f(q)y(m) - \gamma}{\lambda + \mu} + \frac{\mu}{\lambda + \mu} J_{q-1} + \frac{\lambda}{\lambda + \mu}(J_q(1 - y(m)) + J_{q+1}y(m)) \right\} \\ &= \left\{ \frac{f(q)y(g(q)) - \gamma}{\lambda + \mu} + \frac{\mu}{\lambda + \mu} J_{q-1} + \frac{\lambda}{\lambda + \mu}(J_q(1 - y(g(q))) + J_{q+1}y(g(q))) \right\} \end{aligned} \tag{1.3}$$

In the above definition of MPBNE, the first condition uses the Bayesian rule for the customer based on the signaling function g to determine whether to join or balk. The second condition states that the composite function $y \circ g$ solves the *admission control type* MDP for the firm. In the optimality equations (1.3), the constant γ represents the long-run average profit made by the firm under optimal policy, and constants $J_0, J_1, \ldots$ represent the *relative cost* for states $0, 1, \ldots$.

1.3.3 Informative equilibria

While the definition of the pure strategy MPBNE in the previous section is complete, it is not directly amenable for further analysis. Thus, the first step towards characterizing the equilibria is to show that any pure strategy MPBNE can be described using a threshold level. The next proposition shows that such a mapping always exists.

Proposition 1.3.1 *Let the pair* $y(m)$ *and* $g(q)$ *be a pure strategy MPBNE such that* N *defined in condition (1) of Definition 1.3.1 is finite. Then there exists a constant* $\overline{q}$ *such that the pair* $(\widetilde{g}(\cdot), \widetilde{y}(\cdot))$ *given by*

$$\widetilde{g}(q) = \begin{cases} m_1 & q \leq \overline{q}, \\ m_0 & \text{otherwise.} \end{cases}, \quad \widetilde{y}(m) = \begin{cases} 1 & m = m_1, \\ 0 & \text{otherwise.} \end{cases} \tag{1.4}$$

forms a MPBNE with the same firm profit and customer utility.

[4]Note that p_q^N can be thought of as the beliefs of the agents on the state of the systems. These beliefs are consistent with the the strategy of the other players.

The above result implies that instead of studying the actions taken by customers and the announcement made by the firm in each state of the system (i.e., queue length), we can focus on the threshold queue length, below which the customer's action will be "join," while above which it will be "balk." Note that the equilibria characterized using the above proposition requires that the constant N in Definition 1.3.1 be finite. There may exist equilibria where the constant N is infinite. We shall discuss these in Section 1.3.4.

While every pure strategy MPBNE with finite N is equivalent to a pure strategy BNE induced by some threshold, the converse is not true, i.e., not all thresholds induce a pure strategy MPBNE. Indeed thresholds below q^* defined by (1.5) below and above a certain level cannot form a pure strategy MPBNE. Thus, given a threshold level, one needs to verify that it indeed induces a pure strategy MPBNE via the functions $\widetilde{g}$ and $\widetilde{y}$. Since we frequently use this notion, we formally define it below.

Definition 1.3.2 *We say that the threshold* $\overline{q}$ *induces a pure strategy MPBNE if the pair* $(\widetilde{g}(\cdot), \widetilde{y}(\cdot))$ *given by (1.4) forms a MPBNE, and this pair is said to be the induced MPBNE by this threshold.*

Before delving into the analysis of the model and the characterization of the equilibrium, we would like to take a step back and develop intuition into the possible regimes and outcomes. In order to do that, and knowing that we can focus on threshold levels, we characterize two important threshold levels: the first, q^*, denotes the threshold value above which a customer *will not* join, given that he has **full information** of the state of the system, and below which he *will join*. The second threshold level, $\widehat{q}$, is motivated by the service provider's point of view, and denotes the threshold level below which the service provider would like the customers to join, and above which she would like them to balk, if she had **full control** of their actions.

Full information. We will define q^* to be the threshold value above which the customer will not obtain positive utility, in expectation, given full queue length information. It is easy to see that

$$q^* = \left[\frac{R\mu}{c}\right], \tag{1.5}$$

where $[\cdot]$ is the bracket function; i.e., q^* is the largest integer not exceeding $R\mu/c$. Note that this threshold pertains to the marginal customer who decides to balk. We will refer to this as the first-best from the customer's perspective, as this maximizes the utility for the individual (selfish) customer. Note that, as shown in Naor (1969), this threshold, which is based on self-optimization (to use Naor (1969)'s terminology), falls short of maximizing the overall expected utility of the customer population.

Full control. From the service provider's point of view, deciding on a threshold level amounts to deciding what should be the finite waiting space in an $M/M/1/k$ queueing system. For each value of k, the expected number of customers joining the queue per unit of time equals $\lambda\frac{1-\rho^k}{1-\rho^{k+1}}$ where $\rho = \frac{\lambda}{\mu}$. Let $\widehat{q}$ denote the optimal waiting space. Thus, $\widehat{q}$ solves the following full control optimization problem:

$$\widehat{q} = \arg\max_k \lambda \frac{1-\rho^k}{1-\rho^{k+1}} \left[\overline{v}\beta(k) + \underline{v}(1-\beta(k))\right], \tag{1.6}$$

where $\beta(k) = \mathbb{P}\left(W_k \leq \frac{R}{c}\right)$, and W_k is the steady-state sojourn time of the customers who join the $M/M/1/k$ queue. The following proposition is given to show that such a threshold exists, and to discuss the properties of the objective function of the full-control optimization problem faced by the service provider.

Proposition 1.3.2 *The function defined by*

$$\Pi(k) := \lambda \frac{1-\rho^k}{1-\rho^{k+1}} \left[\overline{v}\beta(k) + \underline{v}(1-\beta(k))\right],$$

is unimodal in k*, i.e., there exists* $k^* \in \{1, 2, \ldots, \infty\}$ *such that the function* $\Pi(k)$ *is strictly increasing for* $k < k^*$ *and strictly decreasing for* $k \geq k^*$*.*

Using these two quantities, q^* and $\widehat{q}$, which are based on unilateral optimization under full information to the customers and the full control of the service provider respectively, we can identify three regions. These regions are based on the misalignment between the customers and service provider and correspond to different levels of, the so called *bias* in the cheap talk literature. Each of these regions results in a different type of conflict of interest, and thus different equilibria and outcomes for both sides. Figure 1.3.3 depicts the different regions and the equilibrium announcements in each one, which we will next discuss. We will initially outline the key equilibrium in each of three regions, and the intuition behind them. The intuition will be followed by a formal statement in Proposition 1.3.3. The three cases are:

I. Complete alignment: $q^* = \widehat{q}$. In this region, the interests of the two parties are completely aligned, and thus the pure strategy MPBNE is as follows: The firm gives two signals: i) the first for low congestion, which can be denoted as "Low." This signal is announced if the queue length is below q^*. ii) A second signal denoted by "High," which indicates high congestion, and is given when the queue length exceeds q^*. Thus we have $g(q) = $ "Low" if $q < q^*$ and $g(q) = $ "$High$" otherwise; the customer joins the queue when he/she receives the signal "Low" and balks otherwise, i.e., $y($"Low"$) = $ "join", $y($"High"$) = $ "balk".

As stated before, this is the key equilibrium in this region; however, this need not be the unique pure strategy MPBNE. As discussed in Allon et al. (2007) there are multiple equilibria in this model. However, it can be shown that even the more informative equilibria are equivalent to the one described above.

II. Overly patient customers: $q^* > \widehat{q}$. In this region, if customers are endowed with full information, they would like to join the system even when the service provider would like them to balk (if she had full control). Thus, we use the term "overly patient" to emphasize the fact that, in this case, customers are willing to join a more congested system than what the firm would like. Specifically, when the queue length is between $\widehat{q}$ and q^*, the customers would like to join whereas the firm would like them to balk.

We will show that there is no threshold which is immune to defection by both the customers and the firm and consequently that there is no MPBNE in pure strategies. Indeed, for pure strategy MPBNE to exist the firm should be able to signal "High" and customers who receive "High" should balk. The only threshold immune to profitable deviation by the firm is $\widehat{q}$. Given that under any pure strategy MPBNE, the customers respond to "High" by balking, a profitable deviation for the firm from any other candidate threshold is to announce "High" at $\widehat{q}$. The customers, however, know that $\widehat{q} < q^*$ so that $\widehat{q}$ cannot induce an equilibrium: an arriving customer that receives the signal that instructs him to "balk", can deviate from the prescribed equilibrium strategy by joining; the customer will then earn positive utility (since the *only* state in which he can receive such a signal is on the threshold itself, which is, by assumption, below q^*), and thus *detect* (on average) that such a deviation is profitable - hence ruling out the possibility of a pure strategy MPBNE.

III. Impatient customers: $q^* < \widehat{q}$. In this region, the service provider would like the customers to join a more congested system than the one they wish to join. Specifically, when the queue length is between q^* and $\widehat{q}$, the firm would like the customers to join, whereas the customers would like to balk. In order to study this region, we define $F(q)$ to be the customer's expected utility if he finds q customers in the system upon arrival and decides to join the queue; i.e, $F(q) := R - c\frac{q+1}{\mu}$. We define for $\ell < k$,

$$G(\ell, k) = \sum_{q=\ell}^{k-1} p_q^k F(q), \tag{1.7}$$

where $p_q^k := \frac{\rho^q(1-\rho)}{1-\rho^{k+1}}$ is the steady state measure of the $M/M/1/k$ queue. Here, $G(0, k)$ is interpreted as the average utility of a customer joining the $M/M/1/k$ queue.

Then, we have two subcases to consider:

a) $G(0, \widehat{q}) \geq 0$: if the firm announces "Low" when the queue length is below $\widehat{q}$ and "High" otherwise, the customer would like to join when they get the "Low" signal, as their expected utility is positive (since $G(0, \widehat{q}) > 0$). Further, since in equilibrium "High" would be announced only when the queue exactly equals $\widehat{q}$, the customer would balk as they know that $q^* < \widehat{q}$. This is optimal for the firm and also describes our pure strategy MPBNE for this setting. Thus, the firm is capable of achieving its first best profits and operates as if it has full control over the customer decisions.

b) $G(0, \widehat{q}) < 0$: In this case there is no threshold-induced pure strategy MPBNE. For pure strategies to exist the firm should be able to signal "Low" and customers who receive "Low" should join. As in case II, the only threshold immune to profitable deviation of the firm is $\widehat{q}$. However, the customers know that $\widehat{q} > q^*$, thus the threshold $\widehat{q}$ cannot constitute an equilibrium: an arriving customer that receives a signal that instructs him to "join" would obtain negative expected utility and thus can deviate from the prescribed equilibrium strategy by balking and obtaining zero utility. This rules out the possibility of a threshold-induced pure strategy MPBNE.

   The intuition of the above is as follows: if the expected utility of the customers under an $M/M/1/\widehat{q}$ system, as given by $G(0, \widehat{q})$, is positive, they will have no incentive to deviate. Any deviation here will lead to zero utility for the customers. If, on the other hand, their utility is negative, they would be better-off by not joining at all. Consequently, the threshold $\widehat{q}$ can not induce a pure strategy MPBNE. Further, no other threshold is immune to profitable deviation on the firm's part. Thus, in case III(b) there does not exist a pure strategy MPBNE. We emphasize, however, that in case III(a) the customer can be lured, by using intentional vagueness, to join the system even in states in which they obtain negative expected utility as long as their utility averaged over all state in which they join is positive.

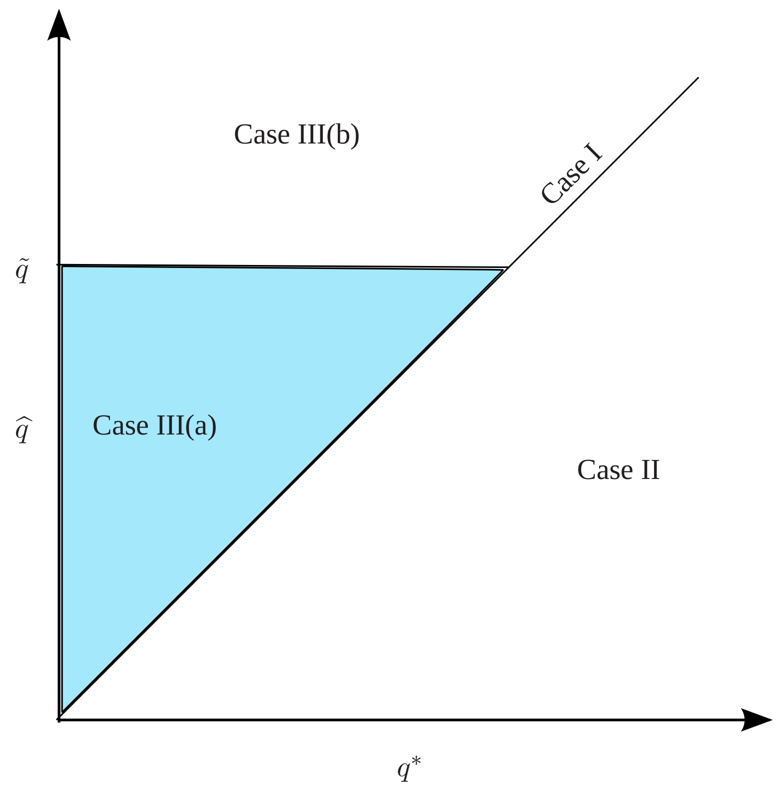

Figure 1.1: Three regions based on full control and full information

We turn now to the formal statement and proof of the equilibria we have discussed thus far. To this end, we let Π_{FI} and Π_{FC} be the firm's profit under full information and full control respectively. Let U_{FI} and U_{FC} denote the expected utility of the customers under full information and full control, respectively. As discussed before, Π_{FC} is the first-best profit for the firm and U_{FI} is the first-best utility for the customer. The next proposition summarizes the above result and also compares the firm's profit and expected customer utility under the different equilibria.

Proposition 1.3.3

I. *If $q^* = \widehat{q}$, then q^* induces a pure strategy MPBNE. Under this equilibrium the firm's profit equals Π_{FC} and the expected utility of the customers is U_{FI}.*

II. *If $q^* > \widehat{q}$, there is no finite q that induces a pure strategy MPBNE.*

III. *If $q^* < \widehat{q}$, then:*

- ***(a)*** *If $G(0, \widehat{q}) > 0$, $\widehat{q}$ induces a pure strategy MPBNE. Under this equilibrium the firm's profit equals Π_{FC} and the expected utility of the customers is U_{FC}.*
- ***(b)*** *If $G(0, \widehat{q}) \le 0$, there is no finite q that induces a pure strategy MPBNE.*

To summarize the findings so far: we have identified three regions, each with a different equilibrium behavior. We observed that a pure strategy MPBNE exists only if the firm's and the customers' incentives are perfectly aligned or if the customers are mildly impatient. We find that in these equilibria, only a two-signal language is required, thus providing analytical support to the common "high congestion/low congestion" announcement observed in practice. Proposition 1.3.3 establishes conditions for the existence of pure-strategy *informative* MPBNE's as a function of the system parameters and characterizes these whenever they exist. It also raises two important questions: are the equilibria outlined above (where they exist) the only equilibria. Further, does the lack of equilibria (for the appropriate regions) suggests that no equilibrium language whatsoever is possible. To discuss these questions, we shall consider these two types of equilibria. First, we show the existence of a *babbling* equilibria, where the firm provides no information. Next, we extend the definition of MPBNE to allow customers to randomize their actions. We characterize the *non-informative* as well as the *informative mixed* strategy MPBNE. Here, the informative mixed strategy MPBNE is again a two-signal language. While other equilibria can be constructed as well, they are equivalent to the two-signal equilibrium.

1.3.4 Non-informative and Other equilibria

The equilibria constructed above are based on a signaling rule with two signals. In practice, however, there are many service providers that share no information whatsoever with the customer, whether it is direct information or one that is implicit in the type of recorded music heard while waiting. Are these systems, where no information is transmitted, in equilibrium? It turns out that such an equilibrium may indeed exist in our setting. When it does exist, it is referred to as a "babbling" equilibrium, to denote that no information is transmitted, and any information provided is treated by the customers as meaningless. In the setting of Crawford and Sobel (1982), such an equilibrium always exists and is sometimes the only possible one. In our model, however, such an equilibrium-in-pure-strategies exists only under certain conditions derived below. In our model, a "babbling equilibrium" exists in pure strategies if, in the absence of information, all customers join (otherwise, given that customers know that all customers balk, they have an incentive to join and earn positive utility). If all customers join, the resulting queueing system is an $M/M/1$ queue (i.e. with infinite waiting space), in which case the average waiting time is $E[W] = 1/(\mu - \lambda)$ and customers join if $R \geq cE[W]$, i.e, if $R \geq \frac{c}{\mu-\lambda}$. In this equilibrium, if indeed all customers join, the system manager can obtain the following profits $\pi^{NI} = \lambda e^{-(\mu-\lambda)\frac{R}{c}}(\overline{v} - \underline{v}) + \lambda \underline{v}$.

Observe that if $R < \frac{c}{\mu-\lambda}$, we cannot have a babbling equilibrium. This underscores one of the the differences between the setting of Crawford and Sobel (1982) and our setting. While the uncertainty in Crawford and Sobel (1982) is independent of the equilibrium dynamics, in our setting there is a clear dependence between the uncertainty (as embedded in the steady-state distribution of the queue) and the resulting equilibrium. This manifests itself in the fact that the babbling equilibrium may not exist. To provide rigorous characterization we have the following result.

Proposition 1.3.4 *There exists a pure strategy babbling equilibrium if and only if $R \geq \frac{c}{\mu-\lambda}$. Further, if $q^* < \widehat{q}$ and $G(0, \widehat{q}) < 0$, i.e. Case III(b) of Proposition 1.3.3, there does not exist a pure strategy babbling equilibrium.*

The following proposition shows that even though a babbling equilibrium may exist, the firm's profit obtained under it is dominated by the firm's profits under the two-signal equilibria described above. Further, the overall customer's expected utility is lower under the babbling equilibrium as compared to that achieved under the two-signal one.

Proposition 1.3.5 *Assume that $R \geq \frac{c}{\mu-\lambda}$ so that the babbling equilibrium is a pure strategy MPBNE. The firm's profits under babbling equilibrium are always dominated by the two-signal equilibrium described in Proposition 1.3.3, if it exists. Further, the customers' expected utility is higher under the equilibrium described in Proposition 1.3.3 than under the babbling equilibrium, if it exists..*

Proposition 1.3.5 emphasizes the value of communication. Even though a non-informative (babbling) equilibria does exist, both the service provider and the customers are always better off when they move to more informative equilibria if such equilibria exist, i.e., to a two-signal equilibria. This communication does not necessarily maximize the customer's overall expected utility but it does improve it. The logic behind Proposition 1.3.5 is as follows: Naor (1969) shows that when customers are self interested and can observe the length of the queue prior to joining, their optimal threshold q^* will be higher than what the social optimum prescribes but it will be finite. In our setting, we observe that for the two-signal equilibrium, the threshold queue length is at least as high as q^*. Further, for the babbling equilibrium, when it exists, the threshold is infinite. Thus, using information improves the customer's overall expected utility when compared to settings where the service provider is giving no information. Note that this improvement is present in the absence of any verification or credibility of the information provided by the service provider.

At this point, we remind the reader that in the region where the customers are very impatient (region III(b)), there is no pure strategy MPBNE neither informative nor non-informative. Without expanding the strategy set for the customer or the firm, it is unclear how the system would behave in this parameter regime. In particular, the customer behavior is unpredictable for the service provider. This issue is alleviated by considering randomization on the part of the customer. We next discuss these results in passing. For more details and formal analysis of these equilibria the reader is referred to Allon et al. (2007).

Mixed strategy non-informative equilibria With the restriction to pure-strategies MPBNE we have shown above that babbling equilibria need not exist. When customers are allowed to use mixed strategies, such equilibria always exist.

The customers randomize among joining and balking, to form a mixed strategy MPBNE as follows: they choose a probability of joining θ that satisfies $R = \frac{c}{\mu - \theta\lambda}$, if $R < \frac{c}{\mu - \lambda}$ and $\theta = 1$ otherwise. Under this equilibria, the arrival process is thinned by the customer randomization such that an arriving customer is indifferent between joining and balking. In particular, the customers do not have any profitable deviation.

Informative cheap talk can be viewed as a mechanism to coordinate incentives of the service provider and the customers when credible information cannot be transmitted. If only babbling equilibrium exists, it might suggest that the non-creditability is hampering any possibility of coordination whatsoever between the players. This is exactly the issue we explore below when we examine whether there is a possibility of improvement in the coordination between the service provider and its customers.

Mixed strategy informative equilibrium Allon et al. (2007) shows that in addition to the babbling equilibria, there may exist more *informative* MPBNEs in mixed strategies. The results in Allon et al. (2007) imply that there are only two possible types of two-signal mixed strategy MPBNE in which randomization is used. The two types can be described as follows: The firm announces "High" and "Low" based on the threshold q_{mix}, (a) in the first type of MPBNE, which we shall refer to as *Join or Randomize* equilibria, the customers who receive "Low" join the system and the customers who receive "High" would join the system with probability $\theta \in (0, 1)$ and balk otherwise; (b) in the second type of MPBNE, which we shall refer to as *Randomize or Balk* equilibria, the customers who receive "Low" join the system with probability θ and balk otherwise, and the customers who receive "High" would balk. Note that both of these types of equilibria are completely defined by two parameters: the threshold q_{mix} used by the firm for signaling and the randomization parameter θ.

Intentional vagueness Allon et al. (2007) shows that unless the firm and the customer are perfectly aligned (that is, $q^* = \widehat{q}$), the equilibrium language always involves *intentional vagueness*. For example, under region (IIIa), the firm uses intentional vagueness to lure customer to join a system they would not join if they had full information. Under mixed strategy equilibria, the firm uses intentional vagueness to ensure that the customer randomizes between joining and balking.

Thus, the firm even though always tells the truth it is almost always an *incomplete* truth.

1.4 Retail Application

In this section we shall apply the above framework to a retail setting. Here, a retailer is trying to sell a product over a time horizon, and provides availability information to the arriving customers who makes a decision whether to buy or wait. For example, the web-retailer **`sierratradingpost.com`** uses the tag "almost gone!" for some of the products, and in its Frequently Asked Questions section explains this tag as follows:

> *If an "almost gone!" label appears next to the item, the sell out risk is very high. We recommend that you place your order immediately.*

Several other web based retailers, such as BarnesandNoble.com and Circuitcity.com, allow customers to search for the availability of specific products for in-store pick-up. Along the same lines, web-based travel agencies such as Expedia.com allow customers to view the availability of airline tickets on specific flights, prior to making the purchasing decision. Similarly, brick-and-mortar stores use different display modes to inform customers about availability. The different displays range from showing ample stock per item to showing only a single item per available product. In all of these examples, the information shared cannot be fully verified by the customers. In the brick-and-mortar examples, a customer does not know if there are more than a single item available even if only one is displayed, and cannot verify whether the stock is indeed low, even if a tag "almost gone!" is attached to an item in web-retailing. In this section, we shall study a formal model and study the emerging equilibrium language between the retailer and its customers. We shall also study the setting when there are multiple decenteralized information channels available to the customers.

1.4.1 Model

Consider a firm that sells a product during a finite length regular season denoted by $[0, \tau]$ followed by a sales season. Here, τ is a stopping time whose distribution is known to both the firm and the customers. Thus, the sales period begins at a random time and both the firm and its customers observe it only once the sales season starts. Further, we assume that the cumulative distribution function of τ is F_τ. We shall make the following assumption with regards to the distribution of the length of the regular season, τ.

Assumption 1.4.1 $\mathbb{E}[\tau - t|\tau > t]$ *is a non-increasing function.*

Simply put, the above assumption requires that in expectation the "sales period" is getting closer as time goes on. The impossibility result described in the paper would hold even under general conditions but to characterize the specific structure of the equilibrium we shall make this assumption. Let $Q = \{Q(t) : t \in [0, \tau]\}$ be the quantity on hand process, i.e., $Q(t)$ denotes the number of products on hand at time $t \in [0, \tau]$. Thus, $Q(0)$ denotes the initial inventory at the beginning of the regular season. Similarly, $Q(\tau)$ denotes the inventory at the end of the regular season and hence the inventory which is being offered at a discounted price during the sale season. Note that the actual evolution of the quantity-on-hand process $Q(t)$ is determined by both the arrival process of the customers and their buying decisions, which depend on the information they have, which includes the information provided by the firm.

Customers arrive according to a Poisson process with rate λ. We denote this arrival process by $N = \{N(t) : t \in [0, \tau]\}$, where $N(t)$ is the number of customers that arrived in the interval $[0, t]$. We assume that the firm sells the product for the price p during the regular season. All units that are left at time τ are discounted and sold at a random price S. We assume that S is a random variable which is independent on all other stochasticity in the system and satisfies $\mathbb{P}(S \leq p) = 1$. Further, we assume that the products during the sale season are sold instantaneously at time τ. Thus, the firm's revenue is $p(Q(0) - Q(\tau)) + SQ(\tau)$. Customers are assumed to be ex-ante symmetric and obtain value v for the purchased product. Here, we assume $v > p$. A customer that arrives at time $t \in [0, \tau]$ makes the decision whether to buy immediately or wait for the sales season. (If $Q(t) = 0$ then there is no decision to be made.) If he buys immediately, he obtains a utility of $v - p$ which we assume to be positive. If he decides to wait until the end of the period for the sale then he obtains the product with probability $A(Q(\tau))$, where $A(x)$ is the probability that any single customer can obtain the product during the sale period if the sale starts with x units on hand[5]. Depending on whether he is able to buy the product during the sales season or not, he obtains $(v - S) - c^W(\tau - t)$ or $-c^W(\tau - t)$, respectively. Here c^W is the waiting cost incurred by the customer, associated with the inconvenience of not obtaining the product immediately. Hence his expected utility is given by $\mathbb{E}[(v - S)A(Q(\tau)) - c^W(\tau - t)|\tau > t]$, where the expectation is over the quantity available at the beginning of sales period, $Q(\tau)$.

We shall refer to $A(Q(\tau))$ as the *availability* of the product during the sale season. The customer has the option to leave the market, and obtain zero utility, but it can be easily seen since $v > p$ that the option of leaving the market is dominated by the "buying now" option. One can envision a more elaborate model for the availability of the product during the sale season. All the structural results from the paper will continue to hold, even if the availability function depends on other factors. However, since the focus of this paper is on the communication, we restrict attention to the above described availability model.

1.4.2 No Information and Full Information Strategies

The main focus of this paper is to characterize the ability (or lack thereof) to communicate unverifiable information to a strategic customer by a retailer. In order to be able to discuss the specific model of communication we will initially discuss the customers behavior under two benchmarks. These correspond to two possible strategies on the firm's part: (i) The strategy of providing no information, and (ii) The strategy of providing the customer full information regarding the availability of the item upon his arrival. The question whether these strategies would emerge in equilibrium is a separate one and would be addressed later in the chapter when we study the game between the retailer and its customers. There, (See Section 1.4.3) we will allow the firm to use different information sharing rules. We will next describe the customers behavior in response to both of these strategies, forming an equilibrium among themselves.

[5]We assume the probability that a customer can obtain the product during the sales period depends on the demand during the period only through the number of sales that occurred. This corresponds, for example, to cases where there are other customers that arrive during the sale period, and do not arrive during the regular season. Cachon and Swinney (2007) describe these customers as "bargain hunters," who frequent the store only during the sale season. The resulting availability for a specific customer in their model is similar to ours.

No information Solution In this setting, we assume that the firm is not providing any information with regards to the inventory position. Note that this is equivalent to the case, where the customers have decided to disregard any information provided by the firm. Since the customers cannot observe the state of the system, they have to rely on the time to make their decisions. Thus, the strategy of the customer shall simply be a function of time. The customer's strategy is represented by $y = \{y(t) : t \in [0, \infty)\}$, where $y(t) \in [0, 1]$ is the probability that a customer arriving at time t buys the product if faced with a decision. (Note that if $t > \tau$ or $Q(t) = 0$ then the customer can not buy the product and there is no decision to be made.) We next define the notion of Markov Perfect Bayseian Nash equilibrium (NE) under-no-information:

Definition 1.4.1 *We say that y forms a* MPBNE under no information*, if the following is satisfied for all $t \in [0, \tau]$:*

$$y(t) \in \arg\max_{\theta \in [0,1]} \theta[(v-p) - (v-S)\mathbb{E}[A(Q_y(\tau))] + c^W(\tau - t)|\tau > t],$$

where $Q_y(\tau)$ is the quantity on hand at time τ if the customers follow strategy y.

The definition requires that the customer buys with probability one if his utility from buying is strictly greater than his utility from waiting, assuming other customers follow their time-dependent strategies y. Similarly, his probability of buying is zero if the utility from buying is strictly dominated by that obtained from waiting. If the utilities from buying and waiting are equal, he randomizes between buying and waiting.

The next result shows that there exists a MPBNE under no information in *pure* strategies, i.e., a MPBNE for which $y(m) \in \{0, 1\}$.

Proposition 1.4.1 *There exists a NE under no information in pure strategies. Specifically, there exists $\hat{\tau}$ such that*

$$y(t) = \begin{cases} 1 & t \le \hat{\tau} \\ 0 & t > \hat{\tau} \end{cases}$$

forms a pure strategy NE.

The above theorem shows that there exists an equilibrium among the customers when the firm does not provide any information. One can view this equilibrium as self-organization of the customers among themselves in the absence of any information. Further, this equilibrium exists in pure strategy, i.e., the arriving customer would buy or wait with probability one, depending on the arrival epoch. Note that under the monotonicity assumption 1.4.1, we have that there exists a threshold $\hat{\tau}$

until which the customer buys and does not buy after that. However, if this assumption is relaxed, then there still exists pure strategy equilibrium in which multiple switch-over points exist, that is, a customer arriving up to time t_1 will purchase the product, a customer arriving between t_1 and t_2 will wait, and a customer arriving after t_2 will buy immediately, again.

Full Information Solution In this setting we assume that the customers have perfect information regarding the quantity on hand, based on which they make their buying/waiting decisions. The customers' strategy in this setting is defined via a mapping $y : \mathbb{Z}_+ \times [0,\infty) \mapsto [0,1]$, where $y(q,t)$ is the probability that a customer arriving at time t buys the product immediately when the quantity on hand is q and $t \leq \tau$. We next define the NE under-full-information.

Definition 1.4.2 *We say that y forms a* NE under-full-information*, if the following is satisfied for all $t \in [0,\infty)$:*

$$y(q,t) \in \arg\max_{\theta\in[0,1]} \theta[(v-p)-(v-\mathbb{E}[S])\mathbb{E}[A(Q_y(\tau))|Q(t)=q,\tau>t]+c^W\mathbb{E}[(\tau-t)|\tau>t]],$$

where $Q_y(\tau)$ is the quantity on hand at time τ if the customers follow strategy y.

To characterize the NE under full information, without loss of generality we can restrict ourselves to threshold induced NE. The reason for this is the fact that for any $q,t \in \mathbb{Z}_+ \times [0,\infty)$, if $y(q,t)=0$ then $y(q',t)=0$ for all $q'>q$. In addition, if two equilibria y and y' differ on a set of Lebesgue measure zero, then the outcomes of the games, in terms of the customers' utility and the firm's profit, are identical. We next define the customer strategy induced by a threshold function $\eta = \{\eta(t) : t \in [0,\infty)\}$.

Definition 1.4.3 *We say that a function η induces the customer strategy y if*

$$y(q,t) = \begin{cases} 1 & q < \eta(t) \\ 0 & \textit{otherwise}. \end{cases}$$

Further, we say that η induces a NE under-full-information if η-induced customer strategy y forms a NE under-full-information.

The next result shows that there is a unique threshold η that induces a NE under full information. To this end, note that since $A(\cdot)$ is non-increasing function, we have that A^{-1}, which denotes the inverse of A, is well defined and is also non-increasing function.

Proposition 1.4.2 *There is a unique NE under full information and it is induced by $\eta_{FI}(\cdot)$ which is defined as the pointwise solution to the following equation:*

$$\eta_{FI}(t) = \left\lceil A^{-1}\left(\frac{(v-p)+c^W\mathbb{E}[(\tau-t)|\tau>t]}{(v-\mathbb{E}[S])}\right)\right\rceil. \tag{1.8}$$

One might suspect that the utility obtained by an average customer endowed with full information is higher than the utility obtained by an average customer under the no-information equilibrium. However this is not always the case, as shown in the numerical study in Allon and Bassamboo (2008). Note that when we move from no-information to full-information, *all* the customers have more information. The utility obtained by a given customer in our model is driven not only by his own information but also by the actions of the other customers, which drive the availability of the product during the sales period. Further, these actions are driven by their own information set. When we move to full information, other customers are also making more informed decisions, thus the average customer may obtain lower utility.

1.4.3 Cheap Talk Equilibrium

In the last section, we fixed the strategy of the firm with regard to information sharing and studied the equilibrium emerging among the customers. In this section, we explore the game played between the firm and its customers, where the firm is allowed to use any information sharing strategy. In particular, the firm can choose full information as well as no information but is not restricted to do so. To define the single-retailer game formally, we shall start by defining the strategy of the customer followed by the strategy of the firm.

Let $\mathcal{M}$ be the Borel set which comprises of feasible signals that the firm can use. Let $y : \mathcal{M} \times [0,\infty) \mapsto [0,1]$ represent the strategy of the customers. Here, $y(m,t)$ is the probability that a customer arriving at time t, receiving a signal $m \in \mathcal{M}$, buys the product immediately. Thus, this customer waits for the sale period which starts at time τ with probability $1-y(m,t)$. Let the space of feasible strategies for the customer be denoted by $\mathcal{Y}$. Let $g : \mathbb{Z} \times [0,\infty) \times \mathcal{M} \mapsto \mathbb{R}$ represent the strategy of the firm. Here $g(q,t,\cdot)$ induces a probability measure on $\mathcal{M}$ from which the firm announces a realization, if the quantity on hand at time t is q. Thus, we will impose the condition that $\int_{\mathcal{M}} g(q,t,m)dm = 1$ for all $q \in \mathbb{Z}$ and $t \in [0,\infty)$. Let the space of feasible strategies for the firm be denoted by $\mathcal{G}$. Note that the quantity on hand process Q is determined by the customer's strategy as well as the firm's strategy g. Let $\mu_{g,y}(t)$ represent the distribution of the signal transmitted at time t if the firm follows strategy g and the customers follow strategy y. A r.v. with measure μ shall be represented by X_μ. Further, let the firms profit under the strategy pair g, y be written as $\Pi(g, y)$, and $Q_{g,y}(t)$ be the inventory on hand process under the strategy pair g, y.

Definition 1.4.4 *We say that the pair $(g, y) \in \mathcal{G} \times \mathcal{Y}$ forms a Markov Perfect Bayesian Nash Equilibrium (MPBNE) in the single-retailer game if and only if it satisfies the following two conditions:*

1. *For all $m \in \mathcal{M}$ and $t \in [0, \infty)$,*

$$y(m, t) \in \arg\max_{y \in [0,1]} y \left[(v - p) - \mathbb{E}[(v - s)A(Q_{g,y}(\tau)) - c^W(\tau - t) | \tau > t, X_{\mu_{g,y}(t)} = m] \right].$$

2. *Fixing the strategy of the customers y, the strategy of the firm g solves:*

$$g \in \arg\max_{\widetilde{g} \in \mathcal{G}} \Pi(y, \widetilde{g}).$$

The above definition requires that both the firm and the customers do not have any unilateral profitable deviation. Specifically, the first condition in the definition requires that fixing the strategy of the rest of the customers and the firm, a customer arriving at time t, should not have any profitable deviation. Similarly, the second condition requires that given the customer's action rule y as fixed, the firm maximizes its profit by using strategy g.

Next, we characterize the emerging equilibria in the single retailer game. We prove that it is impossible for the firm to credibly communicate any information to its customers. This result is equivalent to saying that the only type of equilibria that may arise in such a game are non-informative. Thus, it is either the case that the firm provides no information or the firm provides information, but the customers disregard it in making their decisions due to the lack of credibility on the part of the firm. The equilibrium language that emerges in this game does not carry any information, and is equivalent to babbling. We shall first define the class of equilibria which are non-informative, and hence referred to as *babbling* equilibria.

Definition 1.4.5 *We say that the pair $(y, g) \in \mathcal{Y} \times \mathcal{G}$ forms a babbling equilibrium if and only if the pair (y, g) forms a MPBNE and $y(m_1, t) = y(m_2, t)$ for all $m_1, m_2 \in \mathcal{M}$ and for all $t \in [0, \tau]$.*

This definition states that a MPBNE is a babbling equilibrium if the customer's actions in equilibrium do not depend on the information provided by the firm.

Note that Proposition 1.4.1 already established that such an equilibrium always exists in pure strategies in the single retailer game. We next show that babbling is the *only* type of equilibria that can arise in the single retailer game.

Proposition 1.4.3 (The impossibility result) *Under any MPBNE of the single-retailer cheap talk game, the customer's realized buying behavior satisfies the following*

$$y(X_{\mu_{g,y(t)}}, t) = y^*(t) \quad \textit{a.s.},$$

for almost all $t \in [0, \tau]$, *where there exists a babbling equilibrium where the customer purchases with probability* $y^*(t)$ *at time* $t \in [0, \tau]$.

Proof: Consider any pair $(y, g) \in \mathcal{Y} \times \mathcal{G}$ MPBNE of the above cheap talk game. We shall first show that at any point in time the firm would provide a signal that would maximize the probability of an arriving customer buying the product immediately. That is,

$$y(X_{\mu_{g,y(t)}}, t) = \max_{m' \in \mathcal{M}} y(m', t).$$

For this, consider condition 2 in the definition of the MPBNE in the single retailer game. It can be expressed using Markov Decision Process approach as follows: let $V(q, t)$ be the total expected profit starting from period t until the sales period and have q units on hand. Since the firm would maximize its revenue, $V(\cdot, \cdot)$ should solve:

$$\frac{\partial V(q,t)}{\partial t} = \max_{m \in \mathcal{M}} [\lambda y(m,t)(p + V(q-1,t)) + \lambda(1 - y(m,t))V(q,t) + h(t)\mathbb{E}[S]q - (\lambda + h(t))V(q,t)], \tag{1.9}$$

where $h(t)$ is the hazard rate of τ which defines the beginning of the sales period. The above can be reexpressed as

$$\frac{\partial V(q,t)}{\partial t} = \max_{m \in \mathcal{M}} y(m,t)\lambda[p + V(q-1,t) - V(q,t)] + h(t)\mathbb{E}[S]q - h(t)V(q,t). \tag{1.10}$$

Further, we have $V(q, t) \leq p + V(q - 1, t)$. Thus, we get the desired result that the support of $g(q, t, s)$ is a subset of $\arg\max_{m \in \mathcal{M}} y(m, t)$. So, we have

$$y(X_{\mu_{g,y(t)}}, t) = \max_{m' \in \mathcal{M}} y(m', t), \text{ a.s.}$$

Define $\overline{y}(m, t) = \arg\max_{m' \in \mathcal{M}} y(m', t)$ for all $m \in \mathcal{M}$ and $t \in [0, \infty)$. We can easily verify that the pair $(\overline{y}, g)$ is again a MPBNE. Further, by construction it is a babbling equilibrium. This completes the proof. ■

The above proposition shows that no matter what signalling rule the firm uses, the customers would simply ignore all the signals and make their buying decisions irrespective of any information provided. Thus, in this cheap talk game no credibility whatsoever can be created.

While a babbling equilibrium exists in all variants of the Crawford and Sobel cheap talk game, Allon et al. (2007) demonstrates that it may fail to exist in games with endogonized cheap talk. The result that there exists a pure strategy babbling equilibrium in a retail setting is driven by the fact that customers want to mimic other customers. This is in contrast to Allon et al. (2007) where, if no customer joins/purchases the service, an individual customer would like to join the service. See Section 1.5 for a detailed discussion.

Generalization of the Impossibility Result In this section, we consider the setting where the pricing and the timing is done endogenously by the firm. We assume that the valuation of the product at time t is given by $v(t)$. The firm chooses the regular season price p, the sales period price s and the beginning of the sales period τ. An equilibrium for this generalized cheap talk game can be defined in an analogous manner to Definition 1.4.5 where the strategy of the firm now includes the pricing and timing as well. Next we state the generalization of the impossibility result.

Consider any equilibrium of the generalized cheap talk game. Fix the pricing and the timing strategy of the firm. The signalling strategy of the firm and the buying/waiting behavior of the customer must also form an MPBNE equilibrium of a modified game where the pricing and the length of the regular season is exogenously fixed. Further note that Proposition 1.4.3 also holds for the setting where the valuation are decreasing. Thus this equilibrium must be non-informative. Further, note that if there is no equilibrium for the generalized cheap talk game, the result holds trivially. Thus we have the following general result.

Proposition 1.4.1 *There does not exist any informative equilibrium for the generalized cheap talk game.*

The fact that only babbling equilibria exist in the single retailer game suggests the inability to credibly disclose information is hampering any possibility of information sharing. We explore this issue next, examining whether it is possible to improve coordination between the retailer and its customers by offering several remedies and studying the resulting games.

1.4.4 Remedies and Discussion

Multiple channels of information While the previous section showed that the only equilibrium that emerges in the single retailer game is a babbling one, we next study a decentralized setting where the existence of a second information provider enables the retailers to gain “some” credibility.

There are numerous cases in practice where multiple channels sell inventory from the same pool of inventory and independently provide availability information; (this inventory may either be physically co-located, or virtually pooled). For example, *Dicks.com* and *Modells.com* – whose operations are both run by GSI commerce – compete on the same pool of potential customers yet provide information on the same pool of inventory for the same items. Many brick-and-mortar retailers such as *Barnes & Noble* and *Circuit City* allow the customer to check the availability at the different stores on their web-sites. Furthermore, Walmart.com, BN.com and Circuitcity.com have autonomy in managing their marketing and availability decisions. Demery (2004) explains, "Channels run under different responsibility centers and profit centers, so a dot-com, a catalog and brick-and-mortar store were run as separate businesses." We shall show that this multiplicity of information sources can actually help the firms to achieve some credibility. In cases in which such a system is not yet implemented, allowing customers to obtain information through multiple channels can be viewed as a remedy to the inability to communicate un-verifiable information with only a single retailer.

To study this multiple-retailer setting and to explore how much credibility "decentralization" can create in this setting, we shall next define the model and proceed to analyze it. We consider multiple autonomous sales channels of the same retailer or multiple sellers sharing a common inventory whose status the customer cannot see or verify. In this setting the sellers' signals are based on the common inventory and the customers make their buying decisions based on both signals. We assume that the utility function and profit of the firms are similar to the previous section with the following modification: the firms receive the profits from the products that are sold through them. Note that similar analysis can be carried out for more general systems, where the retailers carry some inventory "on-site" and share the rest.

To illustrate that an informative equilibria exists in this setting, we shall restrict ourselves to pure strategies. To describe the game formally, we denote the strategies of the firms by functions $g_1 : \mathbb{Z} \times [0, \infty) \mapsto \mathcal{M}$ and $g_2 : \mathbb{Z} \times [0, \infty) \mapsto \mathcal{M}$ to represent the signalling rule for the two sellers and $y : \mathcal{M} \times \mathcal{M} \times [0, \infty) \mapsto \{\text{“buy”}, \text{“buy-1”}, \text{“buy-2”}, \text{“wait”}, \text{“wait-1”}, \text{“wait-2”}\}$ to represent the purchasing behavior of the customer. Here $g_i(q, t)$ represents the signal given by the firm $i = 1, 2$ to a customer arriving at time t when the common inventory on hand is q at time t. Here $y(m_1, m_2, t)$ is "buy" if the customer arriving at time t decides to buy with equal probability from firm 1 and firm 2 when he receives the signals $m_1 \in \mathcal{M}$ and $m_2 \in \mathcal{M}$ from firm 1 and firm 2, respectively. The function y is "wait" if the customer decides to wait for the sales period and then buy from either one with equal probability. The action "wait-1" corresponds to the customer deciding to wait for the sales period and buy from retailer 1. The action "wait-2" is defined similarly where the customer buys from retailer 2 in the sales period. Similarly, the actions "buy-1" and "buy-2" correspond to the case when the customer decides to purchase from retailer 1 and 2 with probability one, respectively. Let $\mathcal{G}_1$ and $\mathcal{G}_2$ be the set of feasible strategies for the retailer 1 and 2. For $i = 1, 2$, let $\Pi^i(g_1, g_2, y)$ be the profit of the i^{th} retailer if retailer 1 follows strategy g_1, retailer 2 follows strategy g_2 and the customers follow strategy y.

For the purpose of this study, we shall restrict our attention to threshold induced strategies for the firms. We next define these strategies as follows:

Definition 1.4.6 *Let $\eta_i(t)$ $i = 1, 2$ be a decreasing function over the time interval $[0, \infty)$. The triplet of strategies (g_1, g_2, y) induced by η is defined as follows:*

$$g_i^{\eta_i}(q,t) = \begin{cases} M_1 & q \le \eta_i(t) \\ M_2 & \textit{otherwise} \end{cases} \tag{1.11}$$

$$y^{(\eta_1,\eta_2)}(m_1,m_2) = \begin{cases} \text{“}buy\text{”} & m_1 = m_2 = M_1 \\ \text{“}wait\text{”} & m_1 = m_2 = M_2 \\ \text{“}wait-1\text{”} & m_1 = M_2 \textit{ and } m_2 \ne M_2 \\ \text{“}wait-2\text{”} & m_2 = M_2 \textit{ and } m_1 \ne M_2 \end{cases} \tag{1.12}$$

Further, let Π_η be the total combined profit of the two firms under strategies $\eta_1 = \eta_2 = \eta$.

This definition is based on the following logic for the customer's action. Here M_1 corresponds to a "buy" state and M_2 corresponds to a "wait" state. Note that the the announcement M_2 that induces "wait", can actually be a lack of a signal (i.e. the firm is "silent" about the inventory status, and signals only if M_1 is used). Thus, if the firms agree about the information, the customer makes the decision as if there is just one signal. However if they disagree, then the customer decides not to buy and wait for the sales period. Further, during the sales period the customer (who came during the regular season) visits the firm that provided him the information that it has ample inventory (did not signal M_1) when the other firm did not provide a similar signal.

Next we define the MPBNE for strategies induced by threshold functions (η_1, η_2). For this let $Q_{\eta_1,\eta_2} = \{Q_{\eta_1,\eta_2}(t) : t \in [0, \infty)\}$ be the quantity on hand process, where $Q_{\eta_1,\eta_2}(t)$ is quantity on hand at time t under the strategies induced by (η_1, η_2).

Definition 1.4.7 *We say that the triplet $(g_1, g_2, y) \in \mathcal{G} \times \mathcal{G} \times \mathcal{Y}$ induced by (η_1, η_2) forms a MPBNE in the multi-retailer game if and only if it satisfies the following three conditions:*

1. *For all $m_1, m_2 \in \{M_1, M_2\}$ and $t \in [0, \infty)$, y satisfy the following*

(a) $y(m_1, m_2, t)$ *is "buy" if*

$$(v-p) \geq \mathbb{E}[(v-s)A(Q_{\eta_1,\eta_2}(\tau))-c^w(\tau-t)|g_i^{\eta_i}(Q_{\eta_1,\eta_2}(t),t) = m_i \text{ for } i = 1,2 \text{ and } \tau > t].$$

(b) $y(m_1, m_2, t)$ *is "wait", "wait-1" or "wait-2" if*

$$(v-p) < \mathbb{E}[(v-s)A(Q_{\eta_1,\eta_2}(\tau))-c^w(\tau-t)|g_i^{\eta_i}(Q_{\eta_1,\eta_2}(t),t) = m_i \text{ for } i = 1,2 \text{ and } \tau > t].$$

2. *Fixing* η_2 *(hence* $g_2^{\eta_2}$*) and* y, η_1 *solves:*

$$g_1^{\eta_1} \in \arg\max_{g \in \mathcal{G}} \Pi^1(g, g_2^{\eta_2}, y).$$

3. *Fixing* η_1 *(hence* $g_1^{\eta_1}$*) and* y, η_2 *solves:*

$$g_2^{\eta_2} \in \arg\max_{g \in \mathcal{G}} \Pi^2(g_1^{\eta_1}, g, y).$$

In Section 1.4.3 we showed that a babbling equilibrium always exists. This equilibrium trivially exists also in the multi-retail game. The next proposition shows that there also exists a MPBNE where the firms reveal complete information regarding their inventory to their customers. This MPBNE is induced by the threshold functions $\eta_i = \eta_{FI}$ for $i = 1, 2$, where η_{FI} is the function that induces the NE under-full-information defined in Section 1.4.2.

Proposition 1.4.4 *Let* $p > 2s$*. Then the strategy induced by* $\eta_i(\cdot) = \eta_{FI}(\cdot)$ *for* $i = 1, 2$ *forms a MPBNE.*

The importance of this result stems from the somewhat negative result obtained in Proposition 1.4.3, where it was shown that only a non-informative equilibria can exist in the single retailer game. Here, we show that the presence of another retailer sharing a common inventory can induce full revelation of the quantity in the common pool at any given time. Thus, we show that competition moved the information sharing from being completely non-informative to being fully-informative.

This result also stands in stark contrast to the existing literature on cheap talk games with multiple senders providing information regarding variability in a single dimension. The key driver for the existence of a fully revealing equilibrium even though the inventory status is one-dimensional, is the fact that the customer can "punish" the two senders differently given the signals. Even though both senders are identical, when faced with a signal which is off-the-equilibrium path, the customer punishes the senders in a differential manner. For example, if the quantity on hand is greater than $\eta(t)$, one firm announces "buy" and the other firm announces "wait": the customer punishes the firm announcing "buy" and rewards the one saying "wait" by purchasing in the sale period from the firm that announced "wait." In this manner, the customer punishes the firm deviating and rewards the other. Note that in some cases, such as when the equilibrium prescribes "buy," the customer punishes both firms if one firm deviates and tries to induce "wait." The intuition is that the customer may "need" to punish both firms to ensure that no firm tries to induce "buy" while the equilibrium prescribes "wait."

While the above proposition shows that the presence of competition or decentralization allows firms to credibly disclose information to their customers, one should note that decentralization may "destroy" the equilibrium as well if the gains of selling during the regular season are not high enough when compared to those gained during the sale season. Since both firms are competing on the same customer pool, it may create an incentive for a firm to deviate, and defer their customers to the sale season in the hope of exclusivity.

Next we pose the question whether there are any other informative equilibria (which are induced by some function η) that are not equivalent to the above described fully revealing MPBNE, yet provide the customer with some information regarding the availability level.

Proposition 1.4.5 *There exist two functions $\underline{\eta}$ and $\overline{\eta}$, such that for any $\eta_1 = \eta_2 = \eta$ which induces strategies that form a MPBNE in the multi-retailer game, we have $\underline{\eta} \leq \eta \leq \overline{\eta}$.*

The above proposition shows that there exist two functions $\underline{\eta}$ and $\overline{\eta}$ such that any threshold which induces a MPBNE must lie between $\underline{\eta}$ and $\overline{\eta}$. Figure 1.2 illustrates this result. Under any η that induces an equilibrium, at any point in time t, the signals provided by the retailers depend on whether the inventory on hand lie in the "buy" region or the "wait", corresponding to the region below and above the threshold, respectively. Note that the threshold function η must lie between $\underline{\eta}$ and $\overline{\eta}$ at each point in time. Furthermore, these envelopes themselves induce MPBNE. The exact characterization of these envelope thresholds is given in Allon and Bassamboo (2008). Here, we shall outline the intuition behind the characterization of these thresholds.

The informative equilibria corresponding to $\underline{\eta}$ and $\overline{\eta}$ exhibhit two extreme consumer behaviors: one in which maximum volume of purchases is induced during the regular season and one in which minimum volume is induced. Note that in both of these equilibria the firms do not reveal the actual inventory level and use *intentional vagueness*. For example, in the MPBNE induced by $\overline{\eta}$, the number of purchases is maximized by luring the customers to purchase in states of the inventory they would not buy had they known the exact information. This is accomplished

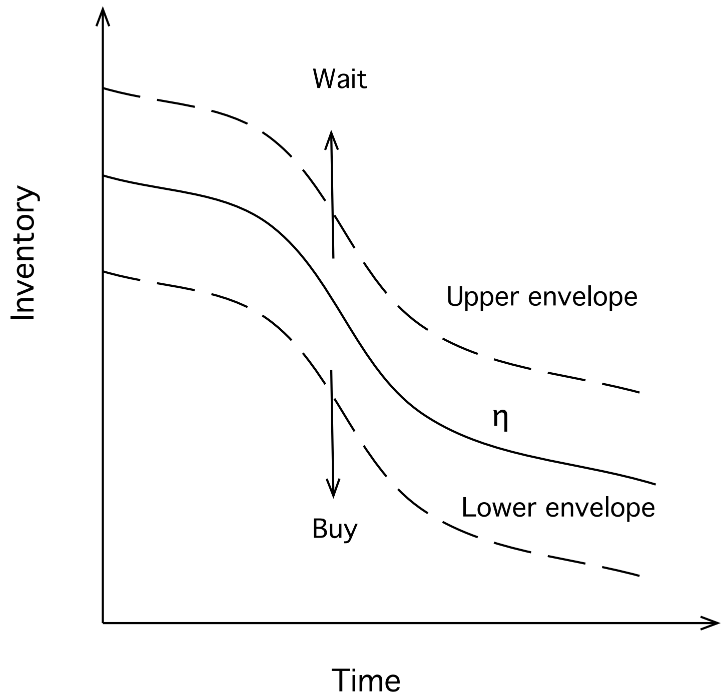

Figure 1.2: Various thresholds that induce MPBNE. For each threshold function, the area below the the threshold represents a "buy" region and the area above the threshold represents a "wait" region.

by giving one signal on the set over which the average utility obtained from waiting equals the utility obtained by buying the product immediately. To illustrate the idea behind intentional vagueness, consider the following scenario: Suppose at a certain point in time $t \in [0, \infty)$, $\overline{\eta}(t) = 5$, and say that $\tau > t$ and at this t had the customer known the exact status of the inventory they would have bought if the inventory was below 3. The firm uses the same signal as long as the inventory level is below 5, i.e., up to 3 (where the customer would have bought anyway) and also when the inventory is 4 (where the customer would *not* have bought). The firm, however, refrains from telling the truth for lower states, as it will not be able to induce customers to buy when the inventory level is 4. This is due to the fact that the customer would then be able to distinguish between the inventory being below 3 and the state being 4.

Noting the fact that there exist multiple equilibria, an important question to study is which equilibrium does the firm prefer among these. To answer this question we first note that any possible equilibrium must lie between the above mentioned envelopes. Thus, while there is clearly multiplicity of equilibria in this model, one can bound both the threshold functions that induce equilibria and the possible outcomes for the firms.

Next, we identify the equilibrium which maximizes the profit for the firms. To that end, we shall denote this threshold by η^*. Thus we have

$$\eta^* \in \arg\max_{\eta \in \mathcal{B}} \Pi_\eta,$$

where $\mathcal{B}$ is the set of all decreasing functions that can induce a MPBNE and Π_η is the combined profit made by the firms under the MPBNE induced by η. Noting the fact that if η_1 and η_2 both induce MPBNE, and $\eta_1 \leq \eta_2$, then the firms' combined profit is higher under the MPBNE induced by η_2, i.e., $\Pi_{\eta_2} \geq \Pi_{\eta_1}$. Then, we have the following corollary.

Corollary 1.4.2 *If $p > 2s$ then the profit maximizing MPBNE in the multi-retailer game is induced by $\overline{\eta}$, i.e., $\eta^* = \overline{\eta}$.*

The above result shows that the presence of decenteralized information by multiple parties *may* improve the firms' profits, if managing to induce, non-cooperatively, equilibrium using the threshold function $\overline{\eta}$.

Third party endorsement. In many settings, organizations can create credibility by being endorsed by others, often called "third-party endorsement". These third-parties, typically do not have a vested interest in the specific firm: they can either be firms that provide certification or generte rating, or non-profit consumer organization. The same role can be played by weblogs covering the specific industry, or bulletin boards where consumers can share information regarding their purchasing experience. One can show that these institutions may allow the retailers to credibly disclose availaibility information. These third party endorsements reduce the strategy space for the firm and improves both firm's credability and the firm's profit.

1.5 Summary

In this chapter, we survey the emerging literature of information sharing between the firms and its customers. A novel framework of *endogenized* cheap talk is developed. In these models, the customers are not only strategic in their actions but also in the way they interpret information, while the firm is strategic in the way it provides information. The developed framework helps answer questions concerning the ability (or the lack thereof) to communicate credibly unverifiable *real-time* information. This framework uses a game-theoretic construct to study this type of communication and discusses the equilibrium language emerging between the firm and its customers. We survey applications of this framework as applied to two models, which are central to the Operations Management literature: the first is a service provider model, and the second is a retail or finite inventory model. We show that one obtains diametrically opposite results with regard to information sharing in service systems and retail systems.

In the setting of a single retailer, the only equilibrium language that may emerge is the one in which no information is revealed to the customer. This result is in contrast to the service setting where a single service provider can "create" some credibility with respect to sharing real-time system information. Further, in the service setting, non-informative, pure strategies equilibrium may not exist. These differences in the nature of equilibrium emerge due to the following distinguishing features of the service and retail operations: a) In retail operations, the incentives of the customers and the firm are aligned for low inventory levels (i.e., both "agree" that the customer should purchase in these states), and misaligned for high inventory level (i.e., the firm would like the customers to purchase, however given that the inventory is high the customers can improve their utility by postponing the purchase to the sales season). However, in service operations, the service provider's and its customers' incentives are aligned both when the number of customers waiting in the system is "high" or "low." The only misalignment is when the number of customers is moderate. Since misalignment is limited in the service setting it helps the provider create some credibility. Thus, the one-sided-only agreement in retail operations games prevents the firm from creating any credibility when it is providing the information on its own. b) The non-existence of an equilibrium when no information is provided in the service setting is due to the "contrarian" behavior characteristic of queueing systems, i.e. customers prefer joining an empty system and resent joining a congested one. On the other hand, customer behavior in retail is one of mimicking, i.e. customers are more interested in buying if many customers buy during the regular season, due to the fear of low availability during the sale season.

One of the strongest phenomena common to both settings is the use of intentional vagueness. In the service setting, the firm might be vague either to lure customers to join the systems in states they would otherwise balk, or to create credibility. In the retail setting, when an informative equilibrium exists (e.g., when the information is provided by multiple autonomous retailers), the firm would always favor using a language that is intentionally vague.

1.6 The Past and the future

Recent literature in Operations Management analyzes and models the impact of strategic customers on managing operational systems. We begin by surveying this literature, both for queueing systems and inventory models.

Queueing models with strategic customers. The literature on queueing models with strategic customers began with Naor (1969), who studied a system in which strategic customers observe the length of the queue prior to making the decision whether to join or balk. There is a (partial) conflict of interest between the self-interested customer and the interests of the social-welfare-maximizing service provider. Naor (1969) shows that pricing can be used to achieve the first-best solution. The follow-up literature that extends Naor (1969) can be broadly divided into two: one that studies models where the firm offers different grades of services (see Mendelson and Whang (1990) and the recent paper by Afeche (2004)), and the other that focuses on competition in the presence of congestion-sensitive customers (see Cachon and Harker (2002) and the recent paper by Allon and Federgruen (2007)). All of these papers assume that the announcements made by the firm are long-term averages, (unlike real-time information), are credible, and are treated as such by customers.

Inventory models with strategic customer. The literature on inventory models with strategic customers can be broadly divided into two categories: a) models where no availability information is provided to the customer and, b) models where customer are provided complete information regarding availability.

Aviv and Pazgal (2007), which falls in the first category, studies pricing strategies for a retailer facing a stochastic arrival stream of customers. When customers arrive, they have no information about the current state of the inventory, and thus their model with fixed-discount strategy corresponds to our no-information model. Cachon and Swinney (2007) considers a model of a retailer that sells a product with uncertain demand over a finite selling season. The authors characterize the rational expectation equilibrium between the firm, who sets its initial quantity level, and the strategic customers, who choose whether to buy during the selling season or during the clearance season. Cachon and Swinney (2007) studies the impact of quick response and the interplay between the existence of strategic customers and this option. Su and Zhang (2007b) shows that the presence of strategic customers can impact the performance of a centralized supply chain when the customers form rational expectation regarding quantities and prices. They show that, while firms cannot commit to specific levels of inventory, decentralized supply chains can use contractual arrangements as indirect commitment devices to attain the desired outcomes with commitment.

Yin et al. (2007), and Su and Zhang (2007a) belong to the second category.

Yin et al. (2007) considers a retailer that announces the regular price and the sales-season clearance price at the beginning of the selling season, as in our model. In the presence of either myopic customers or strategic customers, the authors compare two display modes: one where the retailers displays all the available units (and corresponds to providing full information to the customers) and one where it shows only one unit. Customer treat this one unit as a verifiable proof that the firm has at least one unit in stock. The authors show that the retailers will earn higher expected profits under the “display one unit” format, when the customers are strategic. Su and Zhang (2007a) studies the role of availability and its impact on consumer demand by analyzing a newsvendor model with strategic customers that incur some search cost in order to visit the retailer. They contrast the rational expectations equilibrium in a game where the availability information is not provided to the customer with the scenario, where such information is provided. It is shown that the retailer can improve its profits in the latter. In order to deal with the lack of credibility of the above information, the authors study availability guarantees, in which the seller compensates the consumers in the event of stock-outs.

Delay announcements in other settings. There are several papers that study models in which either a service provider shares waiting time information or a make-to-stock manufacturer shares lead time information.

Hassin (1986) studies the problem of a price-setting, revenue-maximizing service provider that has the option to reveal the queue length to arriving customers, but may choose not to disclose this information, thus leaving the customers to decide whether to join the queue on the basis of the known distribution of the waiting times. The author shows that it may be - but not always - socially optimal to prevent suppression of information, and that it is never optimal to encourage suppression when the revenue maximizer prefers to reveal the queue length. Armony and Maglaras (2004b) analyzes a service system where arriving customers can decide whether to join, balk, or wait for the provider to call within a guaranteed time. The customers’ decisions are based on the equilibrium waiting time (which is equivalent to not providing any information). Armony and Maglaras (2004a) extends the above model to allow the service manager to provide the customers an estimate of the delay, based on the state of the system upon their arrival. The authors show that providing information on the estimated delay improves the system performance. Armony et al. (2007) studies the performance impact of making delay announcements to arriving customers who must wait before starting service in a many-server queue setting with customer abandonment. Customers who must wait are told upon arrival either the delay of the last customer to enter service or an appropriate average delay. Two approximations are proposed: (i) the equilibrium delay in a deterministic fluid model and (ii) the equilibrium steady-state delay in a stochastic model with fixed delay announcements. The authors show that within the fluid-model framework, under certain conditions, the actual delay coincides with the announced delay.

Duenyas and Hopp (1995) studies the problem of quoting customer lead times in a manufacturing environment, both under infinite and finite capacity. For the latter, the authors prove the optimality of different forms of control limit policies for the situations where the lead time is dictated by the market, and the the firms are able to compete on the basis of the lead time. Ata and Olsen (2007) studies a related problem for large systems under convex-concave cost structure.

Dobson and Pinker (2006) develops a stochastic model of a custom production environment with pricing, where customers have different tolerances for waiting. The authors model intermediate levels of information sharing (with a specific structure) ranging from none to complete state-dependent lead-time information, and compare the performance from the firm's and customer's perspectives. They show that for this specific structure it is not always the case that sharing information improves the profits of the firm. Guo and Zipkin (2007) studies a model in which customers are provided with information and make decisions based on their expected waiting times, conditional on the provided information. Three types of information are studied: (i) no information, (ii) queue length, and (iii) the exact waiting time (in systems in which such information is available). The authors provide examples in which accurate delay information improves or hurts the system performance.

1.6.1 Future Research

The framework surveyed in this chapter can be also applied to other operations management settings where the customers cannot credibly verify the information provided to them. One scenario worth exploring is the setting where the firm and the customer engage in "long cheap talk," i.e. the customer is periodically receiving information regarding the inventory. This is common in many retail settings where customers can request to be notified about the future availability of products, and service systems where the customer is informed repeatedly while waiting to be served. It is also worth exploring how this framework applies to fashion retail operations where the customer's utility depends either on the "exclusivity" of the item or its "trendiness," usually conveyed by the retailer.

Bibliography

Afeche, P. 2004. Incentive-compatible revenue management in queueing systems: Optimal strategic idleness and other delaying tactics Working paper, Kellogg School of Management, Northwestern University.

Allon, G., A. Bassamboo. 2008. Buying from the babbling newsvendor: Availability information and cheap talk. *Working paper, Kellogg School of Management, Northwestern University* .

Allon, G., A. Bassamboo, I. Gurvich. 2007. "We will be right with you": Managing customers with vague promises and cheap talk. *Working paper, Kellogg School of Management, Northwestern University* .

Allon, G., A. Federgruen. 2007. Competition in service industries. *Operations Research* **55**(1) 37–55.

Armony, M., C. Maglaras. 2004a. Contact center with a call-back option and real-time delay information. *Operations Research* **52** 527–545.

Armony, M., C. Maglaras. 2004b. On customer contact centers with a call-back options: Customer decisions, routing rules and system design. *Operations Research* **52** 271–292.

Armony, M., N. Shimkin, W. Whitt. 2007. The impact of delay announcements in many-server queues with abandonment. *Working Paper, Stern School of Business, NYU, NY* .

Ata, B., T. Olsen. 2007. Dynamic leadtime quotation under general customer utilities. *Working paper, Kellogg School of Management, Northwestern University* .

Austen-Smith, D. 1990. Information transmission in debate. *American Journal of Political Science* **34**(1) 124–152.

Aviv, Y., A. Pazgal. 2007. Optimal pricing of seasonal products in the presence of forward-looking consumers. *To Appear in MSOM* .

Cachon, G., R. Swinney. 2007. Purchasing, pricing and quick response in the presence of strategic consumers. *Working paper, Wharton School of Business, University of Pennsylvania* .

Cachon, P. C., P. T. Harker. 2002. Competition and outsourcing with scale economies. *Management Science* **48**(10) 1314–1333.

Crawford, V. P., J. Sobel. 1982. Strategic information transmission. *Econometrica* **50** 1431–1451.

Demery, P. 2004. The cross-channel ideal dismantling the application silos to support multi-channel selling. *Internetretailer.com* .

Dobson, G., J. Pinker. 2006. The value of sharing lead time information. *IIE Transactions* **38** 171–183.

Duenyas, I., W. J. Hopp. 1995. Quoting customer lead times. *Management Science* **41**(1) 43–57.

Farrell, J., R. Gibbons. 1989. Cheap talk can matter in bargaining. *Journal of Economic Theory* **48**(1) 221–37.

Farrell, J., M. Rabin. 1996. Cheap talk. *The Journal of Economic Perspectives* **10**(3) 103–118.

Guo, P., P. Zipkin. 2007. Analysis and comparison of queues with deifferent levels of delay information. *Management Science* **53**(6) 962–970.

Hassin, R. 1986. Consumer information in markets with random product quality: The case of queues and balking. *Econometrica* **54** 1185–1195.

Matthews, S. 1989. Veto threats: Rhetoric in a bargaining game. *Quarterly Journal of Economics* **104**(2) 347–370.

Mendelson, H., S. Whang. 1990. Optimal incentive-compatible priority pricing for the M/M/1 queue. *Oper. Res.* **38**(5) 870–883.

Naor, P. 1969. The regulation of queue size by levying tolls. *Econometrica* **37** 15–24.

Ren, Z. J., M. A. Cohen, T. H. Ho, C. Terwiesch. 2007. Information sharing in a long-term supply chain relationship: The role of customer review strategy. *Working paper, Kellogg School of Management, Northwestern University* .

Spence, A. M. 1973. Job market signaling. *Quarterly Journal of Economics* **87** 355–374.

Su, X., F. Zhang. 2007a. On the value of inventory information and availability guarantees when selling to strategic consumers. *Working paper, Haas School of Business, University of California, Berkeley* .

Su, X., F. Zhang. 2007b. Strategic customer behavior, commitment, and supply chain performance. *Working paper, Haas School of Business, University of California, Berkeley* .

Yin, R., Y. Aviv, A. Pazgal, C. S. Tang. 2007. Optimal markdown pricing: Implications of inventory display formats in the presence of strategic customers. *Working paper, Olin School of Business, Washington University* .