INSTITUTO DE ECONOMÍA

MAGÍSTER en ECONOMÍA

2020

Price-setting Decisions Under Experimentation and Rational Inattention

Gustavo Sepúlveda B.

www.economia.uc.cl

PONTIFICIA UNIVERSIDAD CATOLICA DE CHILE
INSTITUTO DE ECONOMIA
MAGISTER EN ECONOMIA

TESIS DE GRADO
MAGISTER EN ECONOMIA

Sepúlveda Benavente, Gustavo Arturo

Diciembre, 2020

PONTIFICIA UNIVERSIDAD CATOLICA DE CHILE
INSTITUTO DE ECONOMIA
MAGISTER EN ECONOMIA

PRICE-SETTING DECISIONS UNDER EXPERIMENTATION AND RATIONAL INATTENTION

Gustavo Arturo Sepúlveda Benavente

Comisión

Nicolás Figueroa

Javier Turén

Santiago, Diciembre de 2020

Price-setting decisions under experimentation and rational inattention*

by Gustavo Sepúlveda B.†

December, 2020

> *I think it is possible that some math models are a representation of philosophy in a single language. Beliefs updating through Bayes' theorem is a learning process contained between the subject and reality. Notwithstanding, the fundamentals of knowledge might be better understood in Kant's books than in this article, which I have painstakingly written.*

Abstract

How do firms acquire knowledge to maximize expected profits? In realistic market conditions, firms try to find out their demand facing imperfect information. There are two learning strategies through which firms can improve their information about markets to set optimal prices. On the one hand, a firm can experiment with price movements away from the optimal to update its demand elasticity knowledge, seeing how much its profit changes. On the other, a firm can pay limited attention to informative signals in the environment. In a two-stage model that implements experimentation and rational inattention, this paper studies the price patterns of a monopolistic seller who faces an uncertain demand curve and whose price-setting decisions are made through active learning in an imperfect-information scenario. Besides, it explores conditions under which the seller considers to complement or substitute these learning strategies.

*This document was written for the Macroeconomics Master's Thesis Seminar at the Pontificia Universidad Católica de Chile's Department of Economics. I am immensely grateful to my advisors Nicolás Figueroa and Javier Turén, for their guidance, valuable comments, time spent, and patience throughout the process of writing this paper. I want to thank my parents, Arturo Sepúlveda Navarrete and Mary I. Benavente Leigh, for their unconditional support all over my college life, through thick and thin. I am also grateful to my friends, especially Claudia De Goyeneche, Ela Díaz, Katia Everke, Sebastián Figari, and Jonathan Rojas, for their comments and ideas, emotional support, and helpful advice. Special thanks to Sebastián Figari for his hardware support.

†gasepulveda@uc.cl. Any errors or omissions in this thesis are my responsibility.

1. Introduction

In the most straightforward microeconomic approaches, monopolistic firms set prices to maximize their expected profit, given a perfectly-known demand curve. In this basic framework, firms take profit-maximizing decisions under perfect information, which leads to an optimal quantity sold and involves no cost to find out which is the optimal price. But it is an unrealistic situation; in the real world, the demand curve is not known with certainty by price-setting agents in markets. Hence, most of the time, firms need to learn about the market environment conditions to get useful information −which is informative of the true demand− so they can take an optimal price-setting decision that maximizes their profit. But how to learn is also a decision firms have to make. In these circumstances, they have two choices to discover what that demand is: either they can choose to experiment with price movements, so they try out the elasticity of demand and receive changes in profits, or they can pay attention to informative signals about the market environment. Nevertheless, acquiring information regarding the demand they face is not easy. In the first case, firms incur a cost via losing profits; in the second one, they are constrained by a bounded attention capacity, so they process a limited amount of informative signals.

Monopolistic firms can learn about two kinds of data: they can either learn of information actively generated on their own since that information is not freely available in the market, or they can learn from available information. The first type of data can be, for example, consumer preferences for specific products or qualities. In contrast, the second one can be the production level in a particular industry, interest rates, or GDP growth. We formulate a static model to understand the seller's price-setting behavior patterns. In the model, monopolists choose to learn of information actively generated by experimenting with price movements and decide to learn about available information by paying attention to available signals in the environment. In addition to studying pricing patterns, the model also shows how intensely the seller learns in one way or the other.

We call the first learning strategy as *experimentation*, and Mirman et al. (1993) have developed a simple framework. In this approach, a monopolist faces a random demand curve with an unknown mean demand, and its manager in charge of setting prices has a prior distribution over the possible mean demand curves. That is to say, he has a prior belief of which is going to be its realization.

The experimentation problem can be thought of as a two-period problem. In the first one, a firm is able to choose optimal prices given its manager's prior belief;[1] then, experimentation occurs: the firm decides to move prices away from the optimal to capture the actual demand curve through changes in profits. The outcome of this process is the posterior distribution over the possible mean demand curves. In other words, the firm updated its knowledge. In the second period, the manager can set a new optimal price due to this further information that experimentation allowed.

[1] Kind of a myopic optimal decision, considering that a prior belief would allow little knowledge about the market environment.

We call the second way of learning as *rational inattention*.[2] It builds on the assumption that agents are rational, but they cannot pay attention to all the available information they would like to. If you look from the outside as an omniscient observer, it would seem that agents are not making entirely rational decisions. In this context, agents pay limited attention and adopt as much available information as affordable for them, and only then they assess what actions to make. Under rational inattention, a monopolistic seller needs to process available information about the environment around him to choose, in an informed way, an optimal price plan given the unknown demand he faces. That information could be anything he considers valuable, such as interest rates, competitors' prices, stock levels, or input prices. The manager of the firm has a prior belief of the true demand in this approach as well, and he updates it with that information.

The rational inattention problem can also be thought of as a two-stage problem. In the first one, a manager has to determine the relevant information to get and process, given an "information processing capacity" constraint (Matějka, 2016). At this point, think of, for example, how much money a firm can allocate to maintain a marketing department doing surveys to know consumer preferences or a commercial department concerned about updating information on competitors' prices. That is, there is a decision about how to process information –say, an information processing mechanism– like surveys, paying a salary to a worker to visit competitor stores and snoop their prices, attending the Central Bank's monetary policy reports, or any other. In the second stage, the manager can set a maximizing-profit price given all this newly available information. In other words, he updated his prior belief of the actual state of demand into a posterior belief that allowed him to set a maximizing-profit price, given the information processing constraint.

Both strategies are treated in the literature of information theory as different frameworks. On the one side, the experimentation problem is framed in what literature calls *information production*; on the other, the rational inattention problem is formulated as *information processing*. Pomatto et al. (2020) use a simple intuition to differentiate both of them: while the rational inattention theory models an agent who pays attention to information that is freely available, experimentation models an agent that generates information –since that information is not freely available–. To understand why this distinction is important, think of a student trying to learn about the decimals of a particular number. If the decimals are freely available –but she does not know them–, an easy solution might be searching into the internet to find them out. However, if the information is not freely available, the student must generate those digits making, for example, a computing effort. Thus, the cost of acquiring knowledge is modeled in two different ways. In the rational inattention approach, paying for a flat internet plan makes it equally costly to learn about the first digit than about ten (here, the total cost is always the same independent of the amount of information). Still, suppose the student chooses to experiment. In that case, a computing effort makes it harder to learn about the tenth digit than the first one, as every decimal is equally costly (the total cost is increasing as you add more decimals). She would need a more expensive version of the software

[2] Christopher Sims proposed this framework in Sims (1998, 2003). Thenceforth, several authors' studies developed information theory concerning rational agents with limited attention.

to get more decimals. While under rational inattention the marginal costs of information pieces are decreasing, under experimentation are constant, which implies that getting closer to revealing the state of nature is more challenging under experimentation since the total cost of producing information increases faster than processing it (you can also think about research-time costs).[3]

Here is the dilemma: why would a firm experiment all the time if it has money enough to have a commercial –or sales– department able to learn about the market environment? Or asked backward, why would a firm always pay attention to information on the market environment if it can experiment? There is no reason to believe that a firm always chooses exclusively one or the other learning strategy to set optimal prices at all. It makes sense that they often do both when considering they have to treat two types of information; unavailable and available. The way to acquire any is different since the latter is "capturable" information from the environment, and the former is not (it has to be generated). In light of this, a decision-maker must adequately choose how and when he learns from one way or the other to maximize expected profits.

We present a two-stage model to understand optimal price-setting decisions of a monopolistic firm that has to actively learn about its environment when perfect information is not accessible. The monopolist faces two possible random demand curves with different levels and slopes. In the first stage, framed in experimentation, he has a prior belief about the high state of demand realization; he wishes to get more information about the actual demand to improve his belief and set more accurate prices in the next period. When the monopolist sets prices at this stage, he observes quantities that are realized from one demand curve or the other, with some probability drawn from a distribution. That said, he would like to set suboptimal prices to "separate" the believed demand curves so as to pull apart the distribution of quantities that come from each curve. This suboptimal pricing leads to a posterior belief –a better knowledge– of the realized state of demand. In the second stage, framed in rational inattention, the monopolist considers this posterior belief about the high state of demand as a knowledge input, and it becomes the prior belief of this second period. He then pays attention to all available and informative signals subject to an information capacity constraint (time or money restrictions representing limited attention) to get a new posterior belief that maximizes the second stage expected profits. These learning processes make the monopolist take the most accurate pricing decision in period two, allowing him to maximize the expected profits of the two periods.

Then, we show how the firm decides to learn according to his attention tightness. Three main results arise. First, when the seller has enough attention to process all the informative signals he wants to learn about the demand, he prefers learning by processing these signals (the rational inattention approach) using all the information channel's capacity and discards experimenting with suboptimal prices in period one. This situation allows the seller to reach a perfect-information scenario with complete-learning in which he sets monopolistic prices in period two since he knows with certainty the realization of the demand curve. Second, when the seller does not have enough

[3] Pomatto et al. (2020) compare the marginal costs of *information production* and *information processing* by showing a flipping-a-coin example.

attention capacity to reach this scenario, he uses both processing signals and experimenting to learn about the state of demand. In this situation, the monopolist sets introductory prices in period one and processes signals to make a more accurate pricing decision in period two. And third, when the seller does not have any attention capacity to process signals, only learns by experimenting in period one. A unique monopolistic but simple modeling, joining experimentation and rational inattention, can provide strong intuition on how firms adequately learn about the environment.

The rest of this paper is structured as follows: Section 2 contrasts the model to related literature. Section 3 motivates, introduces the information theory framework, characterizes every stage of the model, and solves numerical examples. Section 4 discusses conditions to assess substitutability between learning strategies and discusses what types of firms could use one learning approach or the other with greater intensity. Section 5 gives a brief warning about the modeling, and finally, Section 6 concludes and proposes extensions.

2. Related Literature

This paper presents a theoretical model of information acquisition framed in two kinds of learning literature; experimentation and rational inattention. First, the experimentation stage framework follows traditional literature on the two-period horizon model developed since Mirman et al. (1993). Second, we present a particular formulation of Matějka (2016)'s work for the rational inattention stage, in which the monopolist faces only two states of demand in a continuous space of signals. And third, we discuss the closest paper joining rational inattention and experimentation: Chaves (2019).

Mirman et al. (1993) present a model in which a monopolist faces two random mean demand curves with different slopes. He can experiment with price or quantity movements in period one to update his knowledge about that demand to set more accurate prices in period two, so he can maximize both periods' expected profits. In our experimentation problem, we add two small differences: in addition to taking into account two demand curves with different slopes, we also let these curves take on two different levels, and we force the monopolist always to choose prices, not quantities. The main inspiration for doing so is that, in real life, firms do not set production quotas unless they collude in a cartel; instead, they actively choose optimal prices. Besides, we impose a structure for the posterior beliefs distribution to make the numerical problem tractable.

Matějka (2016) solves the optimal information processing mechanism of a rational inattentive monopolist in a continuous and equally-spaced support of n signals and n states (to fix ideas, n different demand curves). He shows that, under an information processing constraint, a firm facing a continuous space of signals chooses only a few prices. The monopolist is just looking at the "big picture" of the state of demand as an optimal response to his inability to process all the available information he would like to. In our rational inattention problem, which is framed in a context where the seller faces only two states of demand in a continuous space of signals, we show that he does not choose the optimal monopolistic prices of each state. Instead, he selects intermediate prices. How further or closer they are from the monopolistic prices depend on how limited his attention is, represented by the *mutual information* channel (the information processing constraint). When the seller's attention capacity is larger, prices are closer to those monopolistic of each state of demand. Caplin et al. (2019) make a theoretical point on this result. They present a theory of optimal consideration set formation in which considered alternatives arise endogenously, based on prior beliefs and information processing costs (in our case, these information processing costs are represented as a constraint on the amount of information the seller can process). In addition to allowing optimal solutions for the states, that setting enables suboptimal actions, such as our model where prices are endogenously chosen –and they do not result in the states' monopolistic optima when facing limited attention–.

In information acquisition literature, it is no easy to find works that merge two active learning strategies to explain optimal decisions. In particular, there is no work merging experimentation and rational inattention in a static formulation to explain optimal pricing patterns. However, Chaves (2019) proposes a comparable approach to the formulated in this paper but another context. He models a two-stage problem in which an agent must choose whether to invest or not in each period. In his problem, initial experimentation (investing) enables the agent to access noise information about the state of nature, which can be improved by paying attention to signals in the environment. Beyond the context, there is an elementary difference to the model presented in our paper. His modeling requires experimenting in the first period to have access to noisy information in the second. Otherwise, the agent cannot refine the precision of the data. In that framework, information is not available at free disposal for agents until experimentation uncovers the state of nature's data. While experimentation is a mandatory stage in his work to acquire knowledge, our setting allows the agent to decide among experimenting if the cost of learning via *mutual information* is infinite (*i.e.*, $\kappa = 0$, as discussed in the next section), complementing strategies and learning exclusively under rational inattention.

3. Model

3.1. *Intuition and motivation*

Information is valuable in decision-making problems when agents face imperfect information. Firms must strive to obtain information about markets to set prices that maximize the expectation of their profits, but how to do it is challenging. In essence, they learn from informational signals, and there are two ways to achieve this aim. When there are no available signals about the actual state of demand in the environment, or when firms cannot capture these signals, they have to generate information by creating them endogenously. In this learning process, the decision involves choosing the precision of the created signals, which is the case of experimentation. Setting a suboptimal price leads to an observed quantity, that is the informative signal that tells the monopolist what demand curve he is likely facing. However, as you can foresee, the process to produce it is costly: whenever the firm sets suboptimal prices to obtain better profits in the future, it loses earnings today. Still, when there are available signals about the state of demand in the environment and firms can capture them, they can get information by choosing which signals to process endogenously. They decide what signals about the demand curve to pay attention to but facing an informational constraint that represents time or spending restrictions, which bounds the firms' attention capacity. That constraint makes the agents "rational inattentive".[4]

Let's focus on a simple case to understand why a firm actively chooses between the two learning strategies: a food company. When a food firm launches a new product, it usually sets an introductory price that is lower than which maximizes profit so that a broader audience can try it. Schlee (2001) shows that a seller lowers initial prices from the believed optimal to tempt potential buyers, increasing consumption and consumers' information about the quality of new products, which leads to an increase in future profits. Similar approaches have been analyzed by Milgrom & Roberts (1986), Bagwell (1990), and Guadalupi (2018). As discussed in Mirman et al. (1993), those firms that introduce new products have incentives to experiment in prices because they are most likely to face uncertainty in demand. This particular pricing action is an effective way to know about the unknown demand curve's hidden information. The information these firms capture is created instead of processed from the environment.

Notwithstanding, after making necessary efforts to create informative signals by introductory prices, firms can keep learning from available information in the market. This data can be anything relevant to the optimal pricing decision: the state of the competition environment (competitors' prices or new competitors' products that could affect demand), or the market conditions (interest rates, price level, or GDP growth). When firms face imperfect information and two different kinds of data to deal with, the learning strategies of experimentation and rational inattention are fundamental to following an optimal price plan.

The RI-experimentation model is a two-horizon problem: the first price-setting stage corresponds to experimentation and the second to rational inattention. The rest of the section proceeds as follows. First, we present the information theory concepts that frame the model. Second, we characterize the problem of a rational inattentive monopolist who wants to maximize profits when facing demand uncertainty, whose only way of acquiring knowledge about that demand is fronting a limited attention capacity and choosing the optimal way of receiving informative signals. Third, we characterize the problem of a monopolistic firm trying to maximize profits when facing demand uncertainty, whose only way of learning about that demand is experimenting through price movements. And fourth, we set up both maximization problems in a unique two-stage model.

[4] Insufficient attention and allocation of limited attention have been long studied in the economic literature of choice-making. Authors like Simon (1959, 1979) and Kahneman (1973) have studied how agents do not pay attention to all available information in the environment but use a selective mental cognition to simplify and deal with all available data. Moreover, Simon (1955) finds that agents consider alternatives until they are suitable, which implies that people or firms assimilate pieces of information that they consider the most important, just as Levitt & March (1988) point out.

3.2. *Framework*

This part explains how information is measured, how agents learn, and how that knowledge is measured. We first define what *entropy* is. Second, we present the example of a monopolist trying to maximize profits, detailing a learning process in which he has limited attention to learn from the environment. Then we move towards conceptualizing how literature represents this limited attention through a constrained information capacity channel with *mutual information*.

To introduce the basic concepts of information theory, it is necessary first to understand the duality between uncertainty and information. Consider that $\boldsymbol{x}$ is a random variable; anyone you want: the heads or tails when you flip a coin, the possible results of rolling a die, or the numbers in a lottery game. All of these outcomes have associated a probability of occurrence. For example, the probability of getting tails is $\frac{1}{2}$, and the probability of getting a six before rolling the die is $\frac{1}{6}$. The probabilities of a discrete random variable $\boldsymbol{x}$ are described by a probability mass function (*pmf*, henceforth), $g(\boldsymbol{x})$.

Any random variable is uncertain until the outcome is realized, and this uncertainty carries a specific amount of information. The following example clarifies it. Imagine having a die whose sides are all one, and think of the experiment "rolling the die": with perfect certainty, we know the result will be one. The random variable support (the six possible realizations) holds no uncertainty since the result is known before we roll the die. It does not give us information to know what the result of the rolling is. On the contrary, if its faces go from one to six, there is uncertainty about the possible outcome, which implies that $\boldsymbol{x}$ contains information: we only know what will happen until we roll the die. Measuring the random variable uncertainty is equivalent to measuring the amount of data it contains; the more uncertainty it has, the more information it holds.

To quantify the amount of information contained in a random variable, we have to define the Shannon's *entropy* concept, which measures the uncertainty of a source of information (*i.e.*, the amount of information that it holds). The *entropy* $\mathcal{H}(g(\boldsymbol{x}))$ of a discrete random variable $\boldsymbol{x}$, with *pmf* $g(\boldsymbol{x})$, is defined by

$$\mathcal{H}(g(\boldsymbol{x})) \equiv -\sum_{x \in X} g(\boldsymbol{x}) \log_2 g(\boldsymbol{x}),$$

where the log is to the base 2 and *entropy* is expressed in *bits*, which is the unit by default unless specified otherwise. If the base of the logarithm is e, it is measured in *nats*.

Let's see how the *entropy* behaves with the dice examples above. When each face is equal to one, with probability one the face will be one before rolling the die: *i.e.*, $g(\boldsymbol{x} = 1) = 1$. Then, $\mathcal{H}(g(\boldsymbol{x})) = -1 \log_2 1 = 0$ *bits*. As there exists no uncertainty about the result (it is already known), the random variable carries no information. The *entropy* is always minimal (*i.e.*, $\mathcal{H}(g(\boldsymbol{x})) = 0$) as long as there is no uncertainty about the realization of $\boldsymbol{x}$. However, when the faces go from one to six, $\mathcal{H}(g(\boldsymbol{x})) = \sum_1^6 \frac{1}{6} \log_2 \frac{1}{6} = 2.59$ *bits*, which implies that the random variable uncertainty does carry information. Note that in this particular situation with six possible outcomes, the *entropy* is maximal since every die face has an equal probability of occurrence (you cannot have more uncertainty about the result). If you roll a loaded die, whose face with the number six comes up with a probability of $\frac{1}{2}$ and its faces from one to five come up with the same probability ($\frac{1}{10}$ each), then $\mathcal{H}(g(\boldsymbol{x})) = \frac{1}{2}\log_2\frac{1}{2} + \sum_1^5 \frac{1}{10}\log_2\frac{1}{10} = 2.16$ *bits*. The random variable uncertainty is lower than in the non-loaded die since the number six has a higher probability of realization than the rest of the numbers. In other words, *a priori*, we know more about the possible outcome, and the random variable brings forth less information than if we roll the non-loaded die. *Entropy* is the concept that will allow us to quantify the knowledge provided by new valuable information through lowering the uncertainty of a random variable. In other words, it will help us to quantify how much we learn from this further information.

A monopolist facing rational inattention gathers information in sequences of small pieces (represented by "signals", which we denote by s) about many things and from, possibly, many sources. These signals can be any data telling him what will likely be the state of nature. For example, suppose it is a spring Sunday, and he wants to know tomorrow's weather to sell either sunshades or umbrellas at the street corner. He sells these products one Monday a month because he lives in a town that celebrates the Town Day of No Pollution, so people go to work by walking (an excellent opportunity to make a lot of money with umbrellas or sunshades). He wishes to maximize profits, but he has just a little van to transport its products and needs to know the climate to sell the appropriate goods. In that case, he can check the weather forecast on the internet, on the newspaper, on the radio, or on cable TV. If we consider that tomorrow's climate is our random variable, $\boldsymbol{x}$, these signals give information about its realization –the sunny or rainy day–. However, none of them provide perfect information; with some probability, they tell the truth, but they can also be wrong. Furthermore, the monopolistic seller has another issue to deal with: he has a limited attention capacity to process these signals. He would like to regularly pay attention to the four sources of information to know which is the most precise forecasting. Still, he cannot, because he does not have enough free time: he must use his limited time to learn about their accuracy. In other words, the seller must choose how to allocate his limited attention capacity to the most accurate signal source –the "best" forecast report– to load the van with a proper amount of sunshades and umbrellas.

The monopolist has a prior belief about every possible event of $\boldsymbol{x}$ –say, he allocates prior probabilities to the realization of a sunny and rainy day–, given by the *pmf* $g(\boldsymbol{x})$. This belief corresponds to the prior knowledge about tomorrow's weather; we can see it as the likelihood the monopolist allocates to the occurrence of any of them before he gets any evidence (*i.e.*, signals). The prior probability he assigns to the weather may be driven by, for example, the proportion of sunny and rainy days in the town in a year or season. After observing some signals, he will update his hypothesis with this new evidence to a posterior belief, given by the probability of seeing one of the climates conditional on the signal received (described by the *pmf* $f(\boldsymbol{x}|s)$). This posterior belief is the probability of observing a specific weather, given that the informative signal is correct (for instance, the probability of having a sunny day given that the forecast said so). He will also update the probability of error of the signal: the likelihood of observing the same weather given that the signal is incorrect (*e.g.*, the probability of having a sunny day given that the forecast said it would be rainy).

The updating procedure mentioned above depends on the seller's ability to incorporate the signals into the learning process; recall that he does not have all the free time he wants to learn about them. To this end, he needs to know the likelihood of observing signals, given that the hypothesis about the state of nature is true. For example, that is the probability that he would have got a sunny forecast on Sunday when he observes a sunny day on Monday. Posed in another way, if he had a crystal ball that predicts with no error on Sunday a sunny day for Monday, what is the probability of watching on Sunday the sunny forecast. In other words, he needs to know how "good" forecasts are (their quality or precision) to update his beliefs about tomorrow's weather after watching the forecast report. We can think of this as the accuracy of the mechanism that provides signals, and it is given by the probability of observing a signal conditional on a specific weather (described by the *pmf* $f(s|\boldsymbol{x})$). The monopolist uses part of his time to inquire information on the quality of the forecasts looking on the internet at, for instance, what percentage of the time they have adequately predicted the weather in the last weeks. With his limited time, he has to find the smartest way to research in order to know the reports' accuracy. Some are more precise forecasting than others, but he does not know which of them before doing some smart research. Later we will call this $f(s|\boldsymbol{x})$ as the information processing mechanism, in the case of the rational inattention approach and the quality of experimentation, in the experimentation approach. This learning process of transforming prior beliefs into posterior beliefs is given by Bayes' theorem,

$$f(\boldsymbol{x}|s) = \frac{f(s|\boldsymbol{x})g(\boldsymbol{x})}{\sum_{\boldsymbol{x}} f(s|\boldsymbol{x})g(\boldsymbol{x})} \equiv \frac{f(\boldsymbol{x}, s)}{\sum_{\boldsymbol{x}} f(\boldsymbol{x}, s)}, \tag{1}$$

where $f(\boldsymbol{x}, s)$ is the joint probability of signals and states. Bayes says that the posterior belief of a state of nature is a proportion between the joint probability of observing a specific state with the correct signal and the sum of joint probabilities of all possible scenarios of $\boldsymbol{x}$ (when the signal is right about the state plus the joint probabilities when the signal fails).

Let's go back to our example. The seller can face two states, $\boldsymbol{x} = x_s$ (a sunny day) and $\boldsymbol{x} = x_r$ (a rainy day); and he can receive two signals, s_s (the forecast says tomorrow it will be sunny) and s_r (the forecast says tomorrow it will be rainy). The posterior belief of a sunny day is given by

$$f(x_s|s_s) = \frac{f(s_s|x_s)g(x_s)}{f(s_s|x_s)g(x_s) + f(s_s|x_r)g(x_r)} \equiv \frac{f(x_s, s_s)}{f(x_s, s_s) + f(x_r, s_s)},$$

while the posterior probability of being cheated by the wrong signal (the posterior belief of the rainy day) is

$$f(x_s|s_r) = \frac{f(s_r|x_s)g(x_s)}{f(s_r|x_s)g(x_s) + f(s_r|x_r)g(x_r)} \equiv \frac{f(x_s, s_r)}{f(x_s, s_r) + f(x_r, s_r)}.$$

These two probabilities are enough to get their complements, $f(x_r|s_r)$ and $f(x_r|s_s)$, which are the probabilities of getting a rainy day when the forecast said so and getting a rainy day when the forecast said it would be sunny, respectively. You can also think of $f(x_r|s_r)$ as the posterior belief of the rainy day, instead of $f(x_s|s_r)$, but we will keep the most used notation in literature.

Imagine that the seller thinks it is more probable that tomorrow's weather is sunny rather than rainy, so his prior beliefs are $g(x_s) = 0.6$ and $g(x_r) = 0.4$,[5] before he sees any evidence. Suppose that, after his research, he discovers that cable TV forecasts are the most accurate for any possible weather realization, so he decides to watch only these reports. Thence, with a likelihood of 80%, the forecasts about sunny days are right, *i.e.*, $f(s_s|x_s) = 0.8$ (80% was the accuracy of cable TV reports in the last two weeks; he could not do a more prolonged investigation of weeks). Therefore, they are wrong with a likelihood of 20%; $f(s_r|x_s) = 0.2$ (20% of the time they foresee rainy days when it is not true). Also, consider that with a likelihood of 60%, the forecasts about rainy days are right, *i.e.*, $f(s_r|x_r) = 0.6$, which implies that they are wrong with a likelihood of 40%; $f(s_s|x_r) = 0.4$ (40% of the time they foresee sunny days when it is not true). Thus, the posterior belief for a sunny day is $f(x_s|s_s) = \frac{3}{4}$, and then $f(x_r|s_s) = \frac{1}{4}$. That is to say, when the TV forecast predicts a sunny day, the seller believes that with a probability of $\frac{3}{4}$ it will be sunny, which is better than his prior belief $g(x_s) = 0.6$ (however, with a probability of $\frac{1}{4}$ he thinks the signal is wrong). Whereas the posterior belief of a rainy day is $f(x_s|s_r) = \frac{1}{3}$, which implies that $f(x_r|s_r) = \frac{2}{3}$; with a probability of $\frac{1}{3}$ he believes that the report is wrong about its rainy forecasting, and with a probability of $\frac{2}{3}$ that it predicts correctly a rainy day. For any given prior belief, the forecasts' precision (how good signals are), $f(s|\boldsymbol{x})$, determine the posterior beliefs about tomorrow's weather (the state of nature); they improve the seller's knowledge and allow him to choose what products to load in the little van for tomorrow's sale after receiving the signal (a particular weather report).

We have already discussed how to measure information and how the seller learns and updates his knowledge. Now, it is necessary to quantify this amount of learning. Recall that, the less we know about our random variable $\boldsymbol{x}$, the more uncertainty it holds, which implies a higher *entropy*. The most common way to quantify the amount of learning is by measuring the expected reduction of *entropy* (uncertainty) of $\boldsymbol{x}$ due to the information content carried by the realization of another random variable; in particular, a realized signal s. This expected reduction in *entropy* between $\boldsymbol{x}$ and s is called *mutual information*, and it is defined by

$$\mathcal{I}[f(\boldsymbol{x}, s)] \equiv \mathcal{H}[g(\boldsymbol{x})] - \mathbb{E}_s \mathcal{H}[f(\boldsymbol{x}|s)],$$

where $\mathcal{H}[f(\boldsymbol{x}|s)] \equiv -\sum_{\boldsymbol{x}} \sum_s f(\boldsymbol{x}, s) \log_2 f(\boldsymbol{x}|s)$ is the *entropy* of the posterior distribution of $\boldsymbol{x}$ conditional on the observation of s, and the joint *pdf* $f(\boldsymbol{x}, s) = f(s|\boldsymbol{x})g(\boldsymbol{x})$. If signals are informative about the realizations of $\boldsymbol{x}$, as we should expect of a rational monopolist that pays –limited– attention to informative signals to maximize profits, then we should see that $\mathcal{H}[f(\boldsymbol{x}|s)] < \mathcal{H}[g(\boldsymbol{x})]$. That is to say, posterior uncertainty about $\boldsymbol{x}$ is lower than its prior uncertainty due to the knowledge bestowed by the signal s: it allows the seller to know with less uncertainty what will be

[5] In this particular situation, $g(x_r) = 1 - g(x_s)$ since there are only two states of nature.

the realization of $\boldsymbol{x}$. The difference between those entropies, *i.e.*, $\mathcal{I}[f(\boldsymbol{x}, s)]$, corresponds to the information content that the signal s provides.

For the last time, let's go back to our monopolist example. The prior beliefs about the weather were $g(x_s) = 0.6$ and $g(x_r) = 1 - g(x_s) = 0.4$, which imply that the *entropy* of the prior distribution is given by

$$\begin{aligned}\mathcal{H}(g(\boldsymbol{x})) =& g(x_s)\log_2 g(x_s) + (1 - g(x_s))\log_2(1 - g(x_s)) \\ =& \frac{6}{10}\log_2\frac{6}{10} + \frac{4}{10}\log_2\frac{4}{10} = 0.971 \ \textit{bits}.\end{aligned}$$

Whereas the *entropy* of the posterior distribution is given by

$$\begin{aligned}\mathcal{H}(f(\boldsymbol{x}|s)) =& f(x_s, s_s)\log_2 f(x_s|s_s) + f(x_s, s_r)\log_2 f(x_s|s_r) \\ &+ f(x_r, s_r)\log_2 f(x_r|s_r) + f(x_r, s_s)\log_2 f(x_r|s_s) \\ =& 0.48\log_2\frac{3}{4} + 0.12\log_2\frac{1}{3} + 0.24\log_2\frac{2}{3} + 0.16\log_2\frac{1}{4} = 0.8498 \ \textit{bits},\end{aligned}$$

which implies that

$$\mathcal{I}[f(\boldsymbol{x}, s)] = \mathcal{H}[g(\boldsymbol{x})] - \mathbb{E}_s\mathcal{H}[f(\boldsymbol{x}|s)] = 0.1211 \textit{bits}.$$

We just showed that $\mathcal{H}[f(\boldsymbol{x}|s)] < \mathcal{H}[g(\boldsymbol{x})]$ when the monopolist processes informative signals. In other words, the seller has now less uncertainty of which will be the weather tomorrow ($\boldsymbol{x}$) after the learning process; he acquired valuable information. In particular, paying attention to the TV forecast implies a reduction of *entropy* of 0.1211 *bits*.

This simple example was intended to clarify the learning process of a monopolist who has limited attention. We studied two main facts. First, the rational inattentive seller chooses the most informative signals in the learning process to update his beliefs, given limited attention. Second, the amount of learning is quantified by the reduction of *entropy* of $\boldsymbol{x}$ due to the realization of a signal through *mutual information*. However, for didactic purposes, we showed how much the seller arbitrarily learned about the weather reports' accuracy, $f(s|\boldsymbol{x})$, having a limited time that allowed him to research only about the past two weeks of forecasts (we did not solved a maximization problem to get the optimal $f(s|\boldsymbol{x})$). What if we wanted to model more or less time of research? The seller could have more or fewer days in his sample to assess the past forecasts' accuracy, respectively. As the monopolist is currently selling one Monday a month (and it is spring), he could get more information about the spring reports' precision if he had more time to study them, which would allow him to analyze their accuracy during the past three or four weeks rather than only two.[6] Suppose the forecast precision was a variable that we have to optimize analytically

[6]The idea of limited attention can be extended beyond limited time. We can also think about, for example, budgetary restrictions that constraint our capacity to acquire knowledge.

or numerically. In that case, we need to constraint the capacity of *mutual information* to reduce the *entropy* of $\boldsymbol{x}$ as a representation of limited attention.

The formal way of representing this bounded attention is by constraining the information capacity that the seller can process. In other words, he will face a restriction on the maximum amount of uncertainty reduction of $\boldsymbol{x}$ that he can achieve due to the knowledge of an informative signal, s. Formally,

$$\mathcal{I}[f(\boldsymbol{x}, s)] \equiv \mathcal{H}[g(\boldsymbol{x})] - \mathbb{E}_s \mathcal{H}[f(\boldsymbol{x}|s)] \leq \kappa$$

where κ is the maximum *entropy* reduction and represents how limited is the seller's attention (a higher κ represents more time available, which would allow him to have more days or weeks in his sample to assess the past forecasts' accuracy and vice versa). This is called the channel's information capacity. For any κ, he will intend to do the best research to learn about the forecasts' accuracy: he will try to use all his available attention to choose the most precise signals (the best forecast report), $f(s|\boldsymbol{x})$, that maximize his posterior beliefs accuracy, $f(\boldsymbol{x}|s)$.

In the illustrative example examined in this part, the seller attempts to learn about tomorrow's climate to make the best decision on what products to carry on his little van. He hypothesizes about which will be the state of nature tomorrow while receiving an informative signal today. The following models that this paper analyzes face the same timing. In the experimentation problem, the seller gets an informative signal (generated by his own while moving prices) about the unobserved state of demand in period one to set more accurate prices in period two. In the rational inattention model, the seller attempts to learn about an unobserved state of demand while processing informative signals at the beginning of the period (you can think of it as period one) to make the optimal price-setting decision at the end of it (you can think of it as period two).

In the context of continuous random variables we refer to the probability function as *pdf* (probability density function) instead of *pmf* (probability mass function). We also integrate the variables in the formulas above instead of using discrete sums. After studying how to quantify information, learning, how agents learn, and how to represent limited attention, we can build the model.

3.3. *The rational inattention decision-making*

This part presents the standard problem of a rational inattentive seller in the context of a continuous space of demand shocks and signals, developed by Matějka (2016). Then, it moves towards a particular formulation with two possible states of demand in a continuous space of signals. Subsequently, it solves a numerical example and presents the algorithm used.

3.3.1. **Formalization *(general problem).*** A monopolistic firm sets a price p to maximize his expected profits, Π.

$$\Pi = \Pi(\mathbf{x}, p), \tag{2}$$

where $\mathbf{x}$ is a vector of continuously distributed random variables that modify the profit function.

The vector $\mathbf{x}$ can model any stochastic components of the environment: micro variables such as competitors' prices, changes in consumer preferences or input costs, and macro variables such as interest rates, GDP growth, or price level. We can think that when the random vector $\mathbf{x}$ is realized, the seller will face a set of realized variables such as the mentioned above, which modify his demand and change $\Pi(\mathbf{x}, p)$. The profit function is concave, which means that there is only one profit-maximizing price p for any realization of $\mathbf{x}$.

The monopolist cannot observe the realized $\mathbf{x}$, so he has a prior belief described by a *pdf* $g(\mathbf{x})$ of its realization; therefore, he needs to process information to learn about it. The information about $\mathbf{x}$ is freely available: pieces of information are at free disposal in the market in the form of signals, and he needs to process them. However, he has limited attention. In this situation, the manager of the firm makes two decisions. First, he chooses how to allocate his attention, so he decides what signals to pay attention to. This attention allocation decision will improve his knowledge about the realization of $\mathbf{x}$ since it allows him to update his prior beliefs. Second, the manager chooses the profit-maximizing price given the information received and the realized posterior belief with a *pdf* $f(\mathbf{x}|s) \equiv \mu(\mathbf{x})$. In other words, the manager maximizes the following expected profits:

$$p[\mu(\mathbf{x})] = \arg\max_{\hat{p}} \int_{\mathbf{x}} \Pi(\mathbf{x}, \hat{p})\mu(\mathbf{x})d\mathbf{x}. \tag{3}$$

Notice that (because of a Bayesian argument) different attention allocation strategies, described by the *pdf* $f(s|\mathbf{x})$, generate different collections of $\mu(\mathbf{x})$; hence, different optimal prices.

The timing of events in a given period is: 1) The random vector $\mathbf{x}$ is realized. 2) The monopolist processes signals –information– about the realized $\mathbf{x}$ and acquires a posterior belief $\mu(\mathbf{x})$. 3) Due to this new knowledge, he sets a price p that maximizes the expected profit.

Nevertheless, acquiring information by paying limited attention to signals in the environment does not provide perfect knowledge about $\mathbf{x}$. If the firm were in an unlimited attention scenario, it could have perfect information about the realization of $\mathbf{x}$; and it would assign probability one to the realized $\mathbf{x}$ (equivalent to observing the actual state of nature). However, in a world of imperfect information, heaven on earth does not exist. The monopolist can choose the optimal way to process the information needed to set prices, but constrained by a maximum information capacity given by the channel of *mutual information.* This restriction is denoted as

$$\mathcal{I}[f(\mathbf{x}, s)] \equiv \mathcal{H}[g(\mathbf{x})] - \mathbb{E}_s\mathcal{H}[f(\mathbf{x}|s)] \leq \kappa, \tag{4}$$

in which κ represents the maximum knowledge achievable, $\mathcal{H}(g(\mathbf{x})) \equiv -\int g(\mathbf{x}) \log_2 g(\mathbf{x})d\mathbf{x}$ is the *entropy* of the prior knowledge and $\mathcal{H}[f(\mathbf{x}|p)] \equiv -\int\int f(\mathbf{x}, p) \log_2 f(\mathbf{x}|p)dpd\mathbf{x}$ is the *entropy* of the posterior knowledge. For a given prior distribution, the size of κ characterizes how limited the attention is: a smaller κ leads to less capacity to process information, which implies less precision of the posterior beliefs. Note that κ can take any value larger than zero, depending on the specific random variable distribution we decide to work out. For example, if we wanted to reduce the roll-the-die *entropy* in the preamble the most, we would need a value of $\kappa = 2.59$ since $\mathcal{H}(g(\boldsymbol{x})) = \sum_1^6 \frac{1}{6}\log_2 \frac{1}{6} = 2.59$. As discussed in the previous part, the seller chooses the posterior beliefs by selecting the most precise $f(s|\boldsymbol{x})$ –given his limited attention–, in a way that is the most favorable for $\Pi(\mathbf{x}, p)$.

We can argue that the joint distribution $f(\mathbf{x}, s)$ describes both the choice of attention allocation $f(s|\mathbf{x})$ and the posterior beliefs $f(\mathbf{x}|s)$, for any given prior belief $g(\mathbf{x})$, because of Bayes' theorem in equation (1) in his continuous setting. As we take the prior beliefs as given, deciding on $f(s|\mathbf{x})$ is equivalent to deciding on $f(\mathbf{x}, s)$. Nonetheless, even though attention allocation determines what signals to obtain, if the information capacity constraints $f(s|\mathbf{x})$, then the joint distribution is constrained as well. After deciding on an optimal mechanism of information processing –the one that determines the most accurate $f(s|\mathbf{x})$ with limited attention–, and after receiving a signal s, he can update his prior belief about the realization of $\mathbf{x}$ to a posterior belief $\mu(\mathbf{x}) \equiv f(\mathbf{x}|s)$, which allows the firm to set a maximizing-profit price according to optimization (3). Since one signal s leads to a specific posterior belief $f(\mathbf{x}|s)$, which determines the optimal pricing response $p(\mu)$ according to (3), then signals shape optimal pricing actions.[7] Therefore, as one signal leads to one action, the joint distribution $f(\mathbf{x}, s)$ describes both the choice of attention allocation and the optimal price-setting decision. Thus, $f(\mathbf{x}, s)$ together with $p(\mu)$ determine the joint distribution of the random vector $\mathbf{x}$ and chosen prices $f(\mathbf{x}, p)$. The implication of this mapping between signals and chosen prices is that the seller decides how to process information (signals) taking into account his optimal pricing responses to that information, $p(f(\mathbf{x}|s))$, which implies that he considers the profit outcomes resulting from $\Pi(\mathbf{x}, p)$ when optimizing $f(s|\mathbf{x})$.

At the beginning of the seller's problem, we mentioned the two decisions he makes. First, he chooses how to process information through a channel of a limited information capacity, *i.e.*, he decides how to allocate his limited attention. There exists a sequence of signals for any $f(s|\mathbf{x})$ that the seller chooses to achieve, as long as it satisfies the constraint of the *mutual information* channel in equation (4). Second, he chooses how to respond to the realized posterior knowledge, *i.e.*, he decides what price to set given the information received. Once a particular signal on $\mathbf{x}$ is realized, the seller chooses an optimal response, $p = \tilde{P}(s)$, maximizing the expected profits

$$\tilde{P}(s) = \arg\max_{p} \int_{\mathbf{x}} \Pi(\mathbf{x}, p) f(\mathbf{x}|s) d\mathbf{x}, \tag{5}$$

where the posterior knowledge $f(\mathbf{x}|s)$ is given by Bayes' theorem, for any $f(s|\mathbf{x})$. However, choosing the information processing mechanism, $f(s|\mathbf{x})$, and optimal prices is a simultaneous problem; both decisions are not independent. While deciding on $f(s|\mathbf{x})$, the seller takes into account his policy

[7] In particular, there is a one-to-one mapping between prices and signals. That is to say, one signal leads to one price. An optimal information processing strategy can never pay attention to two signals that lead to the same price. Given that the *entropy* function is concave, it is never optimal to choose two posterior beliefs of a different form that lead to the same optimal price. We could make an averaged combination of both distributions and get a third one that would lead to the same expected profit, but using less capacity of the *mutual information* channel. In consequence, the distribution of signals and prices are equal. For more details, check Matějka & McKay (2015), Appendix A.2 in Matějka (2016), and Jung et al. (2019).

function, $\tilde{P}(s)$. Hence, choosing how to process information can be written as

$$f(s|\mathbf{x}) = \underset{\hat{f}(s|\mathbf{x})}{\arg\max}\mathbb{E}[\Pi] = \underset{\hat{f}(s|\mathbf{x})}{\arg\max} \int_{\mathbf{x}} \int_{s} \Pi(\mathbf{x}, \tilde{P}(s))\hat{f}(s|\mathbf{x})g(\mathbf{x})d\mathbf{x}ds, \tag{6}$$

subject to equations (4) and (5). That is, he considers his policy function of prices given signals, $\tilde{P}(s)$, and the constrained *mutual information* channel (representing his limited attention). Note that $\hat{f}(s|\mathbf{x})g(\mathbf{x}) \equiv \hat{f}(\mathbf{x}, s)$; as discussed previously, we can maximize any of them because we take the prior belief $g(\mathbf{x})$ as given. Since we already showed that one signal leads to one price, we can substitute $f(s|\mathbf{x})$ by $f(p|\mathbf{x})$, where $p = \tilde{P}(s)$. Thus, we now move towards solving the whole optimization problem in terms of the joint distribution $f(\mathbf{x}, p)$.

Definition 1. ***The decision strategy of a rationally inattentive monopolist.*** *Let $\Pi(\mathbf{x}, p)$ be the seller's profit funcion, $g(\mathbf{x})$ his prior belief, and κ the information capacity. His decision strategy $f(\mathbf{x}, p)$ is a solution to the following maximization problem:*

$$f(\mathbf{x}, p) = \underset{\hat{f}(\mathbf{x},p)}{\arg\max}\mathbb{E}[\Pi(\mathbf{x}, p)] = \underset{\hat{f}(\mathbf{x},p)}{\arg\max} \int_{\mathbf{x}} \int_{p} \Pi(\mathbf{x}, p)\hat{f}(\mathbf{x}, p)dp, \tag{7}$$

subject to

$$\int_{p} \hat{f}(\mathbf{x}, p)dp = g(\mathbf{x}), \quad \forall \mathbf{x}, \tag{8}$$

$$\hat{f}(\mathbf{x}, p) \geq 0, \quad \forall \mathbf{x}, \forall p, \tag{9}$$

$$\mathcal{H}[g(\mathbf{x})] - \mathbb{E}_p\mathcal{H}\left[\hat{f}(\mathbf{x}|p)\right] \leq \kappa. \tag{10}$$

The firm chooses the $f(\mathbf{x}, p)$ that maximizes the expectation of his profit. Remember that, owing to Bayes' theorem, deciding on $f(\mathbf{x}, p)$ –given a prior belief $g(\mathbf{x})$– let us describe the posterior beliefs, $f(\mathbf{x}|p)$, and the distribution of prices charged for each $\mathbf{x}$, $f(p|\mathbf{x})$. Constraint (8) denotes the consistency of joint distributions with prior beliefs, restriction (9) requires non-negativity of the distribution, and inequality (10) is the information capacity constraint. Calling to mind, the latter constraint means that the reduction of *entropy*, due to the knowledge of an informative signal, cannot exceed the firm's ability to process data. How much the firm can learn about $\mathbf{x}$ is bounded by the seller's limited attention. However, notice that the posterior belief $f(\mathbf{x}|p)$ is a result of the optimization; we do not know it until the problem is solved. We need to rewrite the *mutual information* inequality of constraint (10) as

$$\mathcal{H}[g(\mathbf{x})] - \mathbb{E}_p\mathcal{H}\left[\hat{f}(\mathbf{x}|p)\right] = \int_{\mathbf{x}} \int_{p} \hat{f}(\mathbf{x}, p) \log_2 \left(\frac{\hat{f}(\mathbf{x}, p)}{g(\mathbf{x})f(p)}\right) dpd\mathbf{x}, \tag{11}$$

to solve the problem (proof in Appendix A.1), where $f(p)$ is the unconditional probability of setting prices (or receiving signals, equivalently). This maximization problem leads to an optimal information processing mechanism, $f(p|\mathbf{x})$, and the optimal chosen prices, which maximize expected profits.

The analytical solution to find $f(p|\mathbf{x})$ in this problem follows a logit model augmented for the unconditional *pmf* $f(p)$, that results from obtaining the first-order condition of Definition 1 with respect to $f(\mathbf{x}|p)$. You can check this result in Matějka (2016); it implies that the optimal pricing choice is state-dependent. Matějka & McKay (2015) study in detail how the rational inattention formulation leads to this result. However, since $f(p)$ is unknown, this optimization problem has not a closed-form solution but numerical. We can neither find optimal prices nor optimal joint distributions of prices and states analytically; we can only write the expression. In particular, Caplin et al. (2019) prove that, given a set of pricing actions which are chosen with positive probability, the rational inattention problem requires solving for the unconditional probabilities of the options ($f(p)$ in our case), to be a solution for Matějka & McKay (2015)'s posterior beliefs formulation. Nonetheless, it is a probability structure that we do not impose, which implies that we do not know what prices are chosen with positive probability until the maximization is numerically solved (the numerical resolution will endogenously find $f(p)$).

Following this general setting, we now present a formulation with two possible states of demand in a continuous space of signals.

3.3.2. **Formalization *(particular problem)*.** A monopolistic firm sets a price p to maximize his expected profits, Π.

$$\Pi = \Pi(\boldsymbol{x}, p), \tag{12}$$

where $\boldsymbol{x}$ is a random variable that can take on two values, $\bar{x}$ and $\underline{x}$, with $\bar{x} > \underline{x} > 0$. Particularly, $\boldsymbol{x}$ represents the possible state of the demand curve. These two realizable states are driven by market micro and macro variables, such as interest rates, GDP growth, input costs, competitors' prices or consumer preferences. For a given price, the profit function is increasing in $\boldsymbol{x}$, which means that $\Pi(\bar{x}, p) > \Pi(\underline{x}, p)$. When $\boldsymbol{x} = \bar{x}$, the monopolist faces a high state of demand, and when $\boldsymbol{x} = \underline{x}$, the monopolist faces a low state of demand. Such as in the general setting, the profit function is concave, which means that there is only one profit-maximizing price p for any realization of $\boldsymbol{x}$.

However, the monopolist does not know the state of demand −he cannot observe the realized $\boldsymbol{x}$−, so he has a prior belief described by a *pmf* $g(\boldsymbol{x}) \in \{g(\bar{x}), g(\underline{x})\}$ of its realization. Informative signals about $g(\boldsymbol{x})$ are available in the market, and he needs to process them. However, with limited attention, the seller makes two decisions. First, he chooses how to allocate his attention, so he decides what signals to pay attention to. This attention allocation decision will improve his knowledge about the realization of $\boldsymbol{x}$ since it allows him to update his prior beliefs. Second, he chooses the profit-maximizing price given the information received and the realized posterior belief with a *pmf* $f(\boldsymbol{x}|s) \equiv \mu(\boldsymbol{x})$. In other words, the firm maximizes the following expected profits:

$$p[\mu(\boldsymbol{x})] = \arg\max_{\hat{p}} \left\{\Pi(\bar{x}, \hat{p})\mu(\boldsymbol{x}) + \Pi(\underline{x}, \hat{p})(1 - \mu(\boldsymbol{x}))\right\}, \ \textit{for}\ \boldsymbol{x} \in \{\bar{x}, \underline{x}\}, \tag{13}$$

where $\mu(\bar{x}) = f(\bar{x}|s_{\bar{x}})$ and $\mu(\underline{x}) = f(\bar{x}|s_{\underline{x}})$, for any $s_{\bar{x}}$ and $s_{\underline{x}}$, so we have two different price responses. As in the general formulation, notice that (because of a Bayesian argument) different attention allocation strategies, $f(s|\boldsymbol{x})$, generate different collections of $\mu(\boldsymbol{x})$; hence, different optimal prices.

The timing of events in a given period is: 1) The state of demand $\boldsymbol{x}$ is realized. 2) The monopolist processes signals –information– about the realized state of demand and acquires a posterior belief $\mu(\boldsymbol{x})$. 3) Due to this new knowledge, he sets a price p that maximizes the expected profit.

As discussed, acquiring information by paying limited attention to signals in the environment does not provide perfect knowledge about $\boldsymbol{x}$. The monopolist can choose the optimal way to process the information needed to set prices, but constrained by a maximum information capacity given by the channel of *mutual information*. This restriction is denoted as

$$\mathcal{I}[f(\boldsymbol{x}, s)] \equiv \mathcal{H}[g(\boldsymbol{x})] - \mathbb{E}_s \mathcal{H}[f(\boldsymbol{x}|s)] \leq \kappa, \tag{14}$$

in which κ represents the maximum knowledge achievable. The *entropy* of the prior and posterior distributions can be written as

$$\mathcal{H}(g(\boldsymbol{x})) = -\left[g(\bar{x}) \log_2 g(\bar{x}) + (1 - g(\bar{x})) \log_2 (1 - g(\bar{x}))\right],$$

and,

$$\mathcal{H}[f(\boldsymbol{x}|p)] = -\int_p \left[f(\bar{x}, p) \log_2 f(\bar{x}|p) + f(\underline{x}, p) \log_2 f(\underline{x}|p)\right] dp,$$

respectively. Recall that $g(\underline{x}) = 1 - g(\bar{x})$ since the prior distribution is binary. In this particular situation, as the prior distribution is binary, the uncertainty about $\boldsymbol{x}$ is maximal when the prior knowledge allocates 50% of probabilities to the realization of each state; that is, $g(\bar{x}) = 0.5$ and $g(\underline{x}) \equiv 1 - g(\bar{x}) = 0.5$. If that is the case (the most extreme uncertainty scenario with binary beliefs), $\kappa = 1$ represents the perfect-information scenario (*i.e.*, an informative specific signal, either of the high or the low state of demand, leads to set monopolistic prices with probability one), while $\kappa = 0$ represents the prior belief knowledge (*i.e.*, original *entropy* of the state remains unchanged due to non-informative signals), since $\mathcal{H}[g(\boldsymbol{x})] = 1$. Whereby, for a given prior distribution, the size of κ characterizes how limited the attention is: a smaller κ leads to less capacity to process information, which implies less precision of the posterior beliefs. For any other binary prior distribution, the perfect-information scenario is achievable with $\kappa < 1$.

In the general formulation, we had two extensive, meaningful discussions that hold in this particular modeling. First, one signal leads to one pricing action. Second, the joint distribution $f(\boldsymbol{x}, p)$ describes both the attention allocation choice and the optimal price-setting decision.

At the beginning of the seller's problem, we mentioned the two decisions he makes. First, he chooses how to process information through a channel of a limited information capacity, *i.e.*, he decides how to allocate his limited attention. Second, he chooses how to respond to the realized posterior knowledge, *i.e.*, he decides what price to set given the information received. Once a particular signal on $\boldsymbol{x}$ is realized, the seller chooses an optimal response, $p = \tilde{P}(s)$, maximizing the expected profits

$$\tilde{P}(s) = \arg\max_{p} \left\{ \Pi(\bar{x}, p)\mu(\boldsymbol{x}) + \Pi(\underline{x}, p)(1 - \mu(\boldsymbol{x})) \right\}, \text{ for } \boldsymbol{x} \in \{\bar{x}, \underline{x}\}, \tag{15}$$

where the posterior knowledge $\mu(\boldsymbol{x}) \equiv f(\boldsymbol{x}|s)$ is given by Bayes' theorem, for any $f(s|\boldsymbol{x})$. In particular, $\mu(\bar{x}) = f(\bar{x}|s_{\bar{x}})$ and $\mu(\underline{x}) = f(\bar{x}|s_{\underline{x}})$, for any $s_{\bar{x}}$ and $s_{\underline{x}}$, so we have two different policy functions. However, choosing the information processing mechanism, $f(s|\boldsymbol{x})$, and optimal prices is a simultaneous problem; both decisions are not independent. While deciding on $f(s|\boldsymbol{x})$, the seller takes into account his policy functions, $\tilde{P}(s)$. Hence, choosing how to process information can be written as

$$f(s|\mathbf{x}) = \arg\max_{\hat{f}(s|\mathbf{x})} \mathbb{E}[\Pi] = \arg\max_{\hat{f}(s|\mathbf{x})} \int_s \left[\Pi(\bar{x}, \tilde{P}(s))\hat{f}(s|\bar{x})g(\bar{x}) + \Pi(\underline{x}, \tilde{P}(s))\hat{f}(s|\underline{x})g(\underline{x}) \right] ds, \tag{16}$$

subject to equations (14) and (15). That is, he considers each policy function of prices given signals, $\tilde{P}(s)$, and the constrained *mutual information* channel (representing his limited attention). Note that $\hat{f}(s|\boldsymbol{x})g(\boldsymbol{x}) \equiv \hat{f}(\boldsymbol{x}, s)$, for $\boldsymbol{x} \in \{\underline{x}, \bar{x}\}$; we can maximize any of them because we take the prior belief $g(\boldsymbol{x})$ as given. Since we already showed that one signal leads to one price, we can substitute $f(s|\boldsymbol{x})$ by $f(p|\boldsymbol{x})$, where $p = \tilde{P}(s)$. Thus, we now move towards solving the whole particular optimization problem in terms of the joint distribution $f(\boldsymbol{x}, p)$.

Definition 2. ***The decision strategy of a rationally inattentive monopolist.*** *Let $\Pi(\boldsymbol{x}, p)$ be the seller's profit function, $g(\boldsymbol{x})$ his prior belief, and κ the information capacity. His decision strategy $f(\boldsymbol{x}, p)$ is a solution to the following maximization problem:*

$$f(\boldsymbol{x}, p) = \arg\max_{\hat{f}(\boldsymbol{x},p)} \mathbb{E}[\Pi(\boldsymbol{x}, p)] = \arg\max_{\hat{f}(\boldsymbol{x},p)} \int_p \left(\Pi(\bar{x}, p)\hat{f}(\bar{x}, p) + \Pi(\underline{x}, p)\hat{f}(\underline{x}, p) \right) dp, \tag{17}$$

subject to

$$\int_p \hat{f}(\boldsymbol{x}, p)dp = g(\boldsymbol{x}), \quad \text{for } \boldsymbol{x} \in \{\underline{x}, \bar{x}\}, \tag{18}$$

$$\hat{f}(\boldsymbol{x}, p) \geq 0, \quad \text{for } \boldsymbol{x} \in \{\underline{x}, \bar{x}\}, \forall p, \tag{19}$$

$$\mathcal{H}[g(\boldsymbol{x})] - \mathbb{E}_p \mathcal{H}\left[\hat{f}(\boldsymbol{x}|p)\right] \leq \kappa. \tag{20}$$

The firm chooses the $f(\boldsymbol{x}, p)$ that maximizes the expectation of his profit. Constraint (18) denotes the consistency of joint distributions with prior beliefs, restriction (19) requires non-negativity of the distribution, and inequality (20) is the information capacity constraint. As we do not know $f(\mathbf{x}|p)$ until the problem is solved, we rewrite the *mutual information* inequality of constraint (20) as

$$\mathcal{H}[g(\boldsymbol{x})] - \mathbb{E}_p\mathcal{H}\left[\hat{f}(\boldsymbol{x}|p)\right] = \int_p \left[\hat{f}(\bar{x},p)\log_2\left(\frac{\hat{f}(\bar{x},p)}{g(\bar{x})f(p)}\right) + \hat{f}(\underline{x},p)\log_2\left(\frac{\hat{f}(\underline{x},p)}{g(\underline{x})f(p)}\right)\right] dp, \tag{21}$$

to solve the problem (general proof in Appendix A.1), where $f(p)$ is the unconditional probability of setting prices (or receiving signals, equivalently). This maximization problem leads to an optimal information processing mechanism, $f(p|\boldsymbol{x})$, and the optimal chosen prices, which maximize expected profits. Since we do not define $f(p)$, the solution to the problem is numerical.

3.3.3. **Solving the model & algorithm.** Consider the following demand curve, $p = \boldsymbol{a} - \boldsymbol{b}Q$, where $Q = \left[\frac{\boldsymbol{a}}{\boldsymbol{b}} - \frac{p}{\boldsymbol{b}}\right]$ are quantities. Hence, the profit function can be written as:

$$\Pi(\boldsymbol{d}, p) = p\left[\frac{\boldsymbol{a}}{\boldsymbol{b}} - \frac{p}{\boldsymbol{b}}\right], \tag{22}$$

where $\boldsymbol{a} \in \{\underline{a}, \bar{a}\}$ is the level and $\boldsymbol{b} \in \{\underline{b}, \bar{b}\}$ the slope, with $\bar{a} > \underline{a} > 0$ and $\bar{b} > \underline{b} > 0$. With no loss of generality, the curve can only take on two states: while the duple $\{\underline{a}, \bar{b}\}$ represents the low state of demand, $\underline{x}$, the duple $\{\bar{a}, \underline{b}\}$ represents the high state of demand, $\bar{x}$. Simply put, there is a high, elastic demand curve and a low, inelastic demand curve. To simplify the analysis, the price is net of marginal costs. It is straightforward to see that the profit function as a function of posterior beliefs is strictly convex for any $\mu \in (0, 1)$,

$$\frac{\partial^2 \Pi(p[\mu])}{\partial \mu^2} > 0.$$

That is to say, the information is valuable, and then the monopolist can learn and seeks to maximize expected profits.

We define the support of the pricing action distribution between the optimal monopolistic prices that maximize the high and low state of demand. Hence, we set up $\mathbf{p}$ as the vector of eligible prices for any $\kappa \in [0, 1]$, where each element of $\mathbf{p}$ satisfies that $p \in [p^*_{\underline{x}}, p^*_{\bar{x}}]$. In other words, the firm faces a set of signals and has to decide precisely which of them to pay attention to maximize the expectation of its profit. Matějka (2016)'s formulation allows to our particular model that the firm will not necessarily go for optimal monopolistic prices when facing two states. *A priori*, there are not two specific signals leading to monopolistic prices.

We obtain a numerical solution to this problem using the `Matlab`'s `fmincon` local solver with two algorithms: the `sequential quadratic programming (sqp)`, which is an iterative method for constrained nonlinear optimizations, and `GlobalSearch`, which generates start points that will likely improve the best local minimum findings when the solver is running. `sqp` is more precise (but slower in terms of time) than other algorithms and supports `GlobalSearch`, which is required to get a globally optimal solution. We discretize the domain of p and $\boldsymbol{x}$ by introducing 30x2 grid points; thirty prices and two states of demand. Numerically, we consider the following. At the beginning of the period, the prior belief support is given by $\mathbf{a} \in \{10, 20\}$ and $\mathbf{b} \in \{\frac{1}{2}, 1\}$. The duple $\{10, 1\}$ corresponds to the low state of demand, $\underline{x}$, and $\{20, \frac{1}{2}\}$ to the high, $\bar{x}$. Besides, the monopolist knows nothing about the possible realization of the state, so he allocates prior probabilities to each state of the form $\mu_0^H = \frac{1}{2}$ and $(1 - \mu_0^H) = \frac{1}{2}$. The vector of signals (number of prices) to attend is 30. Note that, under a perfect-information scenario, the monopolistic prices are given by $\frac{\boldsymbol{a}}{2}$, for $\mathbf{a} \in \{\underline{a}, \bar{a}\}$. Therefore, $p \in \{\frac{\underline{a}}{2}, \dots, \frac{\bar{a}}{2}\}$. In this particular case, $p \in \{5, \dots, 10\}$. That is to say, there are 28 possible prices in addition to those monopolistic. The vector of prices consists of thirty equally-spaced values.

We set a 30x2 profit matrix with the functional form of equation (22). Each matrix position is denoted as $\Pi(i, j)$, where $j \in \{1, 2\}$ represents the column position and $i \in \{1, \dots, 30\}$ the row position. The first column contains the functional form evaluated in the $\bar{x}$ state, while the second the functional form evaluated in the $\underline{x}$ state. In each row, these functional forms are evaluated in one of the thirty prices distributed as $p \in \{5, \dots, 10\}$. Every element of the profit matrix is multiplied by an unknown denoted as $f(i, j)$ to build the objective function of equation (17), where $j \in \{1, 2\}$ denotes the high or low functional form of the profit matrix, and $i \in \{1, \dots, 30\}$ denotes the position in the vector of prices in which the profit matrix is evaluated. $f(i, j)$ represents the joint distribution of prices and states that we want to maximize since it describes the optimal information processing mechanism and the optimal price-setting decision. Restrictions (18) and (19) are imputed as linear constraints, whereas (21) as nonlinear. Note that, numerically, there is an extra constraint because of equation (21): joint distributions have to be consistent with the prior beliefs in equation (18), but also have to be consistent with $f(p)$. That is to say, we are adding

$$\sum_{x} \hat{f}(\boldsymbol{x}, p) = \hat{f}(\bar{x}, p) + \hat{f}(\underline{x}, p) = f(p), \quad \forall p,$$

which completes the constraints setting of the optimization problem. The numerical solution with this configuration should take less than five minutes, depending on hardware capacity. For $\kappa = 0.5$, Table 1 shows the numerical solution for the optimal prices and joint distributions.

Table 1. Optimal prices and joint distributions $f(\boldsymbol{x}, p)$, for $\kappa = 0.5$:

$f(\boldsymbol{x}, p)$	$\underline{x}$	$\bar{x}$
$p_l^* = 5.17$	0.377	0.008
$p_h^* = 9.48$	0.1229	0.4919

Thus, expected profits are $\mathbb{E}[\Pi(\boldsymbol{x}, p)] = 109.38$. This profit expectation is greater than $\mathbb{E}[\Pi(\boldsymbol{x}, p)] = 109.04$, which results from running the optimization considering two signals that lead to monopolistic prices. The monopolist is better off with optimal intermediate prices rather than with monopolistic prices. Additionally, notice that this model only reports profits after the firm learned about the environment, as we disregard period one's myopic earnings when the optimization problem begins.

These results are interesting for two reasons. First, optimal prices are only two. Second, optimal prices are not the ones that maximize the high or low demand level; they are different from those monopolistic. The optimal price for high (low) demand that the firm sets is lower (higher) than the optimal monopolistic price. What could be driving those results? When the firm's manager faces many different signals and has to decide which of them to pay attention to, he only chooses two. He is allocating attention to a couple of signals that lead to only two intermediate prices. The formulation allows the firm to ≪fear from uncertainty≫: as it is utopian to have perfect knowledge of the environment through limited attention, setting monopolistic prices could be extremely costly in terms of expected profits. Given that signals always make the firm take decisions subject to errors –say, the monopolist receives the signal that with high probability tells him the state of demand is high, but with low probability he is cheated–, a good way not to suffer a huge negative shock in profits, if he makes the wrong price-setting decision, is to get a bit away from the monopolistic optimal prices.

For $\mu_0^H = \frac{1}{2}$, Figure 1 shows the expectation of profits, $\mathbb{E}[\Pi(\boldsymbol{x}, p)]$, as a function of κ, which represents how limited is the seller's attention. Although it seems obvious, as a higher κ represents more capacity of acquiring information since we relax the *mutual information* constraint, the larger κ, the bigger the expected profits, $\mathbb{E}[\Pi(\boldsymbol{x}, p)]$. Alike, for larger values of κ, prices get closer to monopolistic prices. Considering that we are computing the constraint in *bits* ($\log_2$), this two-state distribution's *entropy* is maximal at $\kappa = \frac{1}{2}$, which implies that we are facing the "worst" scenario: perfect knowledge is only achievable with $\kappa = 1$ (on the contrary, $\kappa = 0.1$ is close to a formulation in which the firm can almost learn nothing with *mutual information*). The less we know about the state of demand at the beginning of the period, the harder is to reach the perfect-information scenario for any given κ.

Figure 1. $\mathbb{E}[\Pi(\boldsymbol{x}, p)]$ in axis y as a function of $\kappa \in (\frac{1}{10}, 1)$, for $\mu_0^H = \frac{1}{2}$:

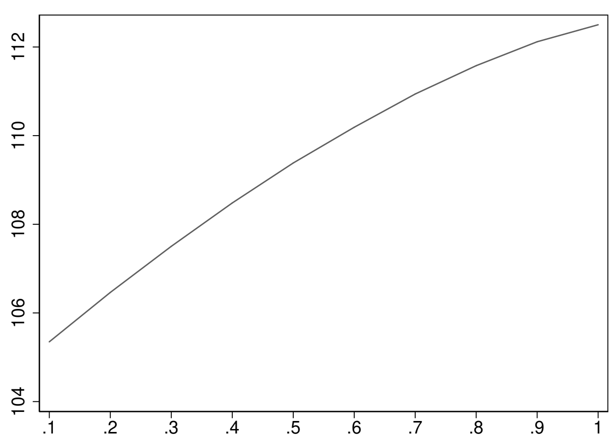

3.4. *The experimentation decision-making*

This part characterizes a two-stage standard model of a firm that learns exclusively by experimenting with price movements about two states of the demand. Subsequently, it solves a numerical example.

3.4.1. **Formalization.** A monopolistic firm maximizes its expected discounted sum of profits of periods one and two, $\mathbb{E}_0\left[\Pi_1\left(p_1, \boldsymbol{x}\right) + \delta\Pi_2\left(p_2, \boldsymbol{x}\right)\right]$. We establish $\delta = 1$ for simplicity. The demand curve is given by

$$q_t = \psi\left(p_t, \boldsymbol{x}\right) + \epsilon_t, \quad \textit{for}\ t \in \{1, 2\}, \tag{23}$$

where q_t are quantities, and p_t are prices. Prices are assumed to be net of marginal costs to simplify the analysis. $\psi\left(p_t, \boldsymbol{x}\right)$ is decreasing in p_t, and $\psi(0, \boldsymbol{x}) > 0$. The parameter $\boldsymbol{x}$ is a random variable; it represents the state of demand and can take on two values, $\boldsymbol{x} \in \{\underline{x}, \bar{x}\}$, with $\bar{x} > \underline{x} > 0$. These two realizable states are driven by market micro and macro variables, such as interest rates, GDP growth, input costs, competitors' prices or consumer preferences. The random variable ϵ_t is a white noise characterized by a *pdf* $f(\epsilon)$ with expected value zero, known variance, and satisfies the monotone likelihood ratio property (MLRP) –*i.e.*, $f'(\epsilon)/f(\epsilon)$ is a continuous and non-increasing function–. To fix ideas, we consider for now that $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$, which implies that the expected demand function is $\mathbb{E}_t[q_t] = \psi\left(p_t, \boldsymbol{x}\right)$. The monopolist knows the distribution of the noise and has a prior belief $g(\bar{x}) = \mu_0^H$ that $\boldsymbol{x} = \bar{x}$.

The firm wants to know if the state of the demand curve is good or bad, which is an incentive to experiment. In each period t the firm chooses a price p_t and, at the end of the period, it observes a quantity q_t. The observed quantity is a function of the price, the unknown parameter $\boldsymbol{x}$ and the realization of ϵ_t. After observing the realized q_1, the monopolist updates his prior beliefs through the Bayes' theorem:

$$\mu_1^H\left(p_1, q_1\right) = \frac{g\left(q_1 | p_1, \bar{x}\right)\mu_0^H}{g\left(q_1 | p_1, \bar{x}\right)\mu_0^H + g\left(q_1 | p_1, \underline{x}\right)\left(1 - \mu_0^H\right)}, \tag{24}$$

while

$$\mu_1^L\left(p_1, q_1\right) = \frac{g\left(q_1 | p_1, \underline{x}\right)\mu_0^H}{g\left(q_1 | p_1, \underline{x}\right)\mu_0^H + g\left(q_1 | p_1, \bar{x}\right)\left(1 - \mu_0^H\right)}. \tag{25}$$

However, notice that $g\left(q_1 | p_1, \boldsymbol{x}\right)$ is the probability of observing a certain quantity given the state of demand and the chosen price. Since the noise $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$ is the only source of uncertainty (where $\epsilon_t = q_t - \psi\left(p_t, \boldsymbol{x}\right)$), it draws the distribution of probabilities of q_t given p_t and $\boldsymbol{x}$. Then, $g\left(q_1 | p_1, \boldsymbol{x}\right)$ is equivalent to writing $f(q_1 - \psi\left(p_1, \boldsymbol{x}\right))$. In other words, the probability of observing a specific noise in quantity is equivalent to the likelihood of observing the realized quantity, given the price and the state of demand. Thus, we can write (24) as

$$\mu_1^H\left(p_1, q_1\right) = \frac{f(q_1 - \psi\left(p_1, \bar{x}\right))\mu_0^H}{f(q_1 - \psi\left(p_1, \bar{x}\right))\mu_0^H + f(q_1 - \psi\left(p_1, \underline{x}\right))(1 - \mu_0^H)}, \tag{26}$$

and likewise (25). The posterior beliefs are now random variables drawn by the distribution of ϵ_1.

This formulation gives a notion of the direction in which experimentation occurs. The goal of the price-setting decision in period one is to make an informative observation of the quantity so that the firm can determine, with some probability, which $\boldsymbol{x}$ is facing. In period one, the firm wishes to set a suboptimal price to capture the most separation between mean demand curves as possible, given its beliefs. The further these curves are at a given price, the less the noise distributions intersect, which allows the firm to know with more certainty if the observed quantity is resulting from one mean curve or the other.

Now we turn to the optimization problem. In the second period, the firm solves:

$$L(\mu_1) \equiv \max_{p_2} \left[\mu_1 p_2 \psi\left(p_2, \bar{x}\right) + (1-\mu_1) p_2 \psi\left(p_2, \underline{x}\right)\right], \quad \textit{for } \mu_1 \in \{\mu_1^H, \mu_1^L\}, \tag{27}$$

where $L(\mu_1)$ is the maximized profit of period two as a function of $p_2(\mu_1)$, given the posterior belief resulting from period one μ_1 (in other words, the second period prior belief). In the first period, the monopolist solves:

$$p_1^E = \arg\max_{p_1} \left\{\mu_0^H p_1 \psi\left(p_1, \bar{x}\right) + \left(1-\mu_0^H\right) \psi\left(p_1, \underline{x}\right) + \mathbb{E}_q L(\mu_1(p,q))\right\}. \tag{28}$$

As the distribution of probabilities of q_t is drawn by $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$, every realization of ϵ_1 leads to a particular posterior μ_1, which forces us to take the expectation of profits in period two. In particular, the posterior belief is a random variable for any given price, with a distribution $G(\mu_1, p_1)$. Thus, maximization (28) can be rewritten as:

$$p_1^E = \arg\max_{p_1} \left\{\mu_0^H p_1 \psi\left(p_1, \bar{x}\right) + \left(1-\mu_0^H\right) \psi\left(p_1, \underline{x}\right) + \int_0^1 L(\mu_1) dG(\mu_1, p_1)\right\}. \tag{29}$$

3.4.2. **Solving the model & algorithm.** The difficulty of finding out p_1^E is to set up a tractable expression for $G(\mu_1, p_1)$ to capture the main conclusions of the optimization problem. Since $G(\mu_1, p_1)$ follows a normal distribution of n realizable posterior beliefs, it requires a grid to build all possible realizations of μ_1, which becomes an intractable numerical maximization (particularly considering that the experimentation problem is the first stage of the joint learning strategy). Consider the following demand curve:

$$\psi\left(p_t, \boldsymbol{x}\right) \equiv q_t = \left[\frac{\boldsymbol{a}}{\boldsymbol{b}} - \frac{p_t}{\boldsymbol{b}}\right] + \epsilon_t, \quad \textit{for } t \in \{1, 2\}, \tag{30}$$

where $\boldsymbol{a} \in \{\underline{a}, \bar{a}\}$ is the level and $\boldsymbol{b} \in \{\underline{b}, \bar{b}\}$ the slope, with $\bar{a} > \underline{a} > 0$ and $\bar{b} > \underline{b} > 0$. With no loss of generality, the curve can only take on two states: while the duple $\{\underline{a}, \bar{b}\}$ represents the low state of demand, $\underline{x}$, the duple $\{\bar{a}, \underline{b}\}$ represents the high state of demand, $\bar{x}$. Simply put, there is a high, elastic demand curve and a low, inelastic demand curve. This setting implies that the lower the price is, the most separated the mean curves are, which allows the monopolist to learn more as he sets lower prices. Also, note that the profit function as a function of posterior beliefs is strictly convex for any $\mu_1 \in (0,1)$, such as in the rational inattention approach, which is the formal requirement to have valuable information and experiment.

At the beginning of the section (part 3.2), we had a broad discussion on how the accuracy of the mechanism that provides signals, $f(s|\boldsymbol{x})$, determines the posterior beliefs' precision. In the experimentation problem, signals are endogenous: the signal that tells the seller which mean demand is likely facing is the quantity realized after setting prices in period one. He generates the signal by setting suboptimal prices. That is to say, how precise are the posterior beliefs owing to price experimentation depends on $g\left(q_1|p_1,\boldsymbol{x}\right) = f(\epsilon_1)$. This mechanism leads to a collection of posterior beliefs (μ_1^H and μ_1^L) since $\epsilon_t \sim \mathcal{N}(0,\sigma_\epsilon^2)$. We need to simplify this setting to make the optimization problem tractable.

Assumption 1. ***(quality of experimentation).*** *For any given p_1, with probability σ the monopolist observes the actual q_1, and with probability $(1-\sigma)$ he observes the wrong q_1.*

The already stated assumption captures the idea of observing the mean quantity of the demand curve with some probability, which turns us aside from having two whole distributions of posterior beliefs to only two posterior beliefs. In other words, we set the accuracy of experimentation as $\sigma = f(q_1 - \psi\left(p_1, \bar{x}\right))$ and $1-\sigma = f(q_1 - \psi\left(p_1, \underline{x}\right))$. Let q^H be the signal of the high state of demand, and let q^L be the signal of the low state of demand. As a consequence, the posterior beliefs can be written as

$$\mathbb{P}\left(\bar{x}|q^H\right) \equiv \mu_1^H = \frac{\mu_0^H\sigma}{\mu_0^H\sigma + \left(1-\mu_0^H\right)(1-\sigma)}, \tag{31}$$

and,

$$\mathbb{P}\left(\bar{x}|q^L\right) \equiv \mu_1^L = \frac{\mu_0^H(1-\sigma)}{\mu_0^H(1-\sigma) + \left(1-\mu_0^H\right)\sigma}, \tag{32}$$

respectively. These posterior beliefs imply that

$$\begin{aligned}\int_0^1 L(\mu_1)dG(\mu_1,p_1) =& L(\mu_1^H)(\mu_0\sigma + (1-\mu_0^H)(1-\sigma)) \\ &+ L(\mu_1^L)(\mu_0^H(1-\sigma) + (1-\mu_0^H)\sigma),\end{aligned} \tag{33}$$

which is a solution for equation (29). Note that $\mu_0\sigma + (1-\mu_0^H)(1-\sigma)$ is the likelihood of attaining the high state of demand, and corresponds to the sum of the joint probabilities $g(\bar{x},q^H) = \mu_0\sigma$ and $g(\underline{x},q^H) = (1-\mu_0^H)(1-\sigma)$. That is to say, the likelihood of attaining the high state of demand is the sum of probabilities of all the scenarios in which the seller receives the signal of the high demand, q^H. Similarly, $\mu_0^H(1-\sigma) + (1-\mu_0^H)\sigma$ is the likelihood of attaining the low state of demand, and corresponds to the sum of the joint probabilities $g(\bar{x},q^L) = \mu_0^H(1-\sigma)$ and $g(\underline{x},q^L) = (1-\mu_0^H)\sigma$. The likelihood of attaining the low state of demand is the sum of probabilities of all the scenarios in which the seller receives the signal of the low demand, q^L. Also, note that we already have the functional form that solves the second period of the optimization problem in equation (29):

$$L(\mu_1^H) \equiv \max_{p_2^H} \left(\mu_1^H p_2^H \left[\frac{\bar{a}}{\underline{b}} - \frac{p_2^H}{\underline{b}} \right] + (1-\mu_1^H) p_2^H \left[\frac{\underline{a}}{\bar{b}} - \frac{p_2^H}{\bar{b}} \right] \right), \tag{34}$$

and,

$$L(\mu_1^L) \equiv \max_{p_2^L} \left(\mu_1^L p_2^L \left[\frac{\bar{a}}{\underline{b}} - \frac{p_2^L}{\underline{b}} \right] + (1-\mu_1^L) p_2^L \left[\frac{\underline{a}}{\bar{b}} - \frac{p_2^L}{\bar{b}} \right] \right). \tag{35}$$

While the first period in that equation is simply solved by (30). However, when the monopolist decides not to experiment, he sets myopic prices given his prior beliefs in period one, which are determined by

$$p_1(\mu_0^H) = \arg\max_{p_1} \left(\mu_0^H p_1 \left[\frac{\bar{a}}{\underline{b}} - \frac{p_1}{\underline{b}} \right] + (1-\mu_0^H) p_1 \left[\frac{\underline{a}}{\bar{b}} - \frac{p_1}{\bar{b}} \right] \right). \tag{36}$$

We also need a functional form for the quality of experimentation, σ, as a decreasing function in prices. Namely, we need to make the monopolist learn about the state of demand as he lowers prices. For this purpose, we define the function as

$$\sigma(p) = m - np, \tag{37}$$

where m is an intercept and n a slope; we will then set *ad hoc* values to those parameters to model the problem adequately. It is easy to see that $\sigma'(p) < 0$, which implies that experimentation goes into the right direction.

We obtain a numerical solution to this problem by using `Matlab`. However, it is straightforward to solve and is not hardware intensive; no solver is needed. We discretize the domain of $\boldsymbol{x}$ by introducing two grid points (two states of demand), and forty for p (prices). Numerically, we consider the following. At the beginning of the period, the prior belief support is given by $\mathbf{a} \in \{10, 20\}$ and $\mathbf{b} \in \{\frac{1}{2}, 1\}$, and it is assumed to be the same for both periods. The duple $\{10, 1\}$ corresponds to the low state of demand, $\underline{x}$, and $\{20, \frac{1}{2}\}$ to the high, $\bar{x}$. Besides, the monopolist knows nothing about the possible realization of the state, so he allocates prior probabilities to each state of the form $\mu_0^H = \frac{1}{2}$ and $(1-\mu_0^H) = \frac{1}{2}$. Thus, the myopic price of equation (36) is $p_1(\mu_0^H) = 8.33$. We let the vector of prices be between the monopolistic for every state, $p \in \{5, \ldots, 10\}$. This vector contains equally-spaced values. The set up for $\sigma(p)$ considers $m = 1.46$ and $n = 0.096$, which implies that $\sigma(p) \in \{\frac{1}{2}, \ldots, 0.98\}$ for $p \in \{5, \ldots, 10\}$. Three formal conclusions arise with this setting for $\sigma(p)$: the price-setting decision makes the monopolist get more information of the state of demand as he lowers prices, he cannot unlearn in the process, and he cannot get perfect knowledge.

We first set the vector for σ as $\sigma(i, 1)$ with the functional form of equation (37), where $i \in \{1, \ldots, 40\}$ represents the row position. Each row of σ is evaluated in one of the forty (equally-spaced) prices distributed as $p \in \{5, \ldots, 10\}$. Given the prior beliefs and the vector for σ, we build the posterior beliefs vector of equations (31) and (32) with the form $\mu_1^H(i,1)$ and $\mu_1^L(i,1)$, where $i \in \{1,\ldots,40\}$ represents the row position. Each row of $\mu_1^H(i,1)$ and $\mu_1^L(i,1)$ is evaluated in one of the forty values of $\sigma(i,1)$. For period one, we set a 40x2 profit matrix with the functional form of equation (30). Each matrix position is denoted as $\Pi(i,j)$, where $j \in \{1,2\}$ represents the column position and $i \in \{1,\ldots,40\}$ the row position. The first column contains the functional form of the $\bar{x}$ state, while the second the functional form of the $\underline{x}$ state. In each row, these functional forms are evaluated in one of the forty prices distributed as $p \in \{5,\ldots,10\}$. Finally, we build the 40x1 vector of period one expected profits as in the first middle of equation (29). Optimal prices in period two result from analytically deriving equations (34) and (35), and we set them up in a 40x2 matrix (keeping the already stated syntax order). As the values for the posterior beliefs are already known for each σ, doing so is simple. Then, we set the 40x2 profit matrix of period two using (34) and (35) with the optimal prices for each position i,j. We also build the probabilities of attaining the states of equation (33) as a 40x2 matrix. Finally, we multiply each element of the profit matrix columns, which contains $L(\mu_1^H)(i,1)$ and $L(\mu_1^L)(i,2)$, by each element of the probabilities-of-attaining-the-states matrix (to be precise, an element-wise multiplication), and then we sum by rows to get the 40x1 expected profits vector of period two. The last procedure remaining to know the experimentation price in period one is summing both periods' profits by rows, getting the maximum, and finding out which is the first-period price (looking at the position in the price vector) that leads the best expected profits achievable.

The optimal experimentation price in period one that maximizes expected profits is $p_1^E = 7.94$, which is lower than $p_1(\mu_0^H) = 8.33$, the myopic price. Optimal posterior beliefs after experimentation are $\mu_1^H = 0.69$ and $\mu_1^L = 0.31$ (*i.e.*, $\mathbb{P}(\underline{x}|s^L) = 0.69$). Note that this symmetry in probabilities is a gadget of $\mu_0^H = \frac{1}{2}$; if we set up any other prior belief, this symmetry gets lost. Finally, total profits with this optimal experimentation price corresponds to $\mathbb{E}_0\left[\Pi_{1+2}^E\right] = 209.28$, which is higher than if the monopolist decides to set myopic prices in both periods, $\mathbb{E}_0\left[\Pi_{1+2}\right] = 208.33$. When the monopolist knows nothing about the state of demand at the beginning of the problem, setting suboptimal prices in period one is useful to increase total expected profits; earnings more than offset the firm's cost of these introductory prices (the earnings loss in period one). Information is valuable, and capturing it by setting suboptimal prices allows the seller to make a more accurate pricing decision in period two to maximize total expected profits.

3.5. *The joint learning decision-making*

This part relies on the two previous models. We set up the experimentation problem as the first stage and the rational inattention problem as the second; Section 5 discusses why a reverse order is not a plausible formulation.

3.5.1. **Formalization.**

The monopolistic seller's problem has the same demand setting of the experimentation strategy. The firm maximizes its expected discounted sum of profits of periods one and two, $\mathbb{E}_0\left[\Pi_1(p_1,\boldsymbol{x}) + \delta\Pi_2(p_2,\boldsymbol{x})\right]$. We establish $\delta = 1$ for simplicity. The demand curve is given by

$$q_t = \psi\left(p_t, \boldsymbol{x}\right) + \epsilon_t, \quad \textit{for}\ t \in \{1, 2\}, \tag{38}$$

where q_t are quantities, and p_t are prices. Prices are assumed to be net of marginal costs to simplify the analysis. $\psi\left(p_t, \boldsymbol{x}\right)$ is decreasing in p_t, and $\psi(0, \boldsymbol{x}) > 0$. The parameter $\boldsymbol{x}$ is a random variable; it represents the state of demand and can take on two values, $\boldsymbol{x} \in \{\underline{x}, \bar{x}\}$, with $\bar{x} > \underline{x} > 0$. We consider $\sigma(p)$ as the quality of experimentation of Assumption 1, keeping the same properties discussed in the previous part.

The model considers two learning periods. In the first one, the firm has the chance to experiment with suboptimal prices. We showed that this procedure leads to two posterior beliefs. Strictly speaking, it leads to two binary posterior belief distributions, given by the pdf_s $g^H(\boldsymbol{x})$ and $g^L(\boldsymbol{x})$, where $g^H(\bar{x}) = \mu_1^H$, $g^H(\underline{x}) = 1 - \mu_1^H$, $g^L(\bar{x}) = \mu_1^L$, and $g^L(\underline{x}) = 1 - \mu_1^L$. In the second period, the firm can learn under rational inattention by paying attention to available signals in the environment. This stage considers the posterior distributions resulting from learning by experimenting in the first period, $g^H(\boldsymbol{x})$ and $g^L(\boldsymbol{x})$, as the prior distributions.

In the second period, the firm solves the rational inattention profit-maximizing problem for the prior distribution $g^H(\boldsymbol{x})$, denoted as $L\left(\mu_1^H\right)$:

$$\hat{L}\left(\mu_1^H\right) = \mathbb{E}\left[\hat{\Pi}\left(\boldsymbol{x}, p\right)\right] = \max_{f(\boldsymbol{x},p)} \int_p \left[\Pi\left(\bar{x}, p\right) f\left(\bar{x}, p\right) + \Pi\left(\underline{x}, p\right) f\left(\underline{x}, p\right)\right] dp, \tag{39}$$

subject to

$$\int_p f(\boldsymbol{x}, p) dp = g^H(\boldsymbol{x}), \quad \textit{for}\ \boldsymbol{x} \in \{\underline{x}, \bar{x}\}, \tag{40}$$

$$f(\boldsymbol{x}, p) \geq 0, \quad \textit{for}\ \boldsymbol{x} \in \{\underline{x}, \bar{x}\}, \forall p, \tag{41}$$

$$\mathcal{H}[g^H(\boldsymbol{x})] - \mathbb{E}_p \mathcal{H}\left[f(\boldsymbol{x}|p)\right] \leq \kappa, \tag{42}$$

and the rational inattention profit-maximizing problem for the prior distribution $g^L(\boldsymbol{x})$, denoted as $L\left(\mu_1^L\right)$:

$$\hat{L}\left(\mu_1^L\right) = \mathbb{E}\left[\hat{\Pi}\left(\boldsymbol{x}, p\right)\right] = \max_{f(\boldsymbol{x},p)} \int_p \left[\Pi\left(\bar{x}, p\right) f\left(\bar{x}, p\right) + \Pi\left(\underline{x}, p\right) f\left(\underline{x}, p\right)\right] dp, \tag{43}$$

subject to

$$\int_p f(\boldsymbol{x}, p) dp = g^L(\boldsymbol{x}), \quad \textit{for}\ \boldsymbol{x} \in \{\underline{x}, \bar{x}\}, \tag{44}$$

$$f(\boldsymbol{x}, p) \geq 0, \quad \textit{for}\ \boldsymbol{x} \in \{\underline{x}, \bar{x}\}, \forall p, \tag{45}$$

$$\mathcal{H}[g^L(\boldsymbol{x})] - \mathbb{E}_p \mathcal{H}\left[f(\boldsymbol{x}|p)\right] \leq \kappa, \tag{46}$$

following the same notation of Definition 2. To have tractable expressions for the *mutual informa-*

tion constraints, we rewrite equations (42) and (46) as

$$\mathcal{H}[g^H(\boldsymbol{x})] - \mathbb{E}_p\mathcal{H}\left[f(\boldsymbol{x}|p)\right] = \int_p \left[f(\bar{x},p)\log_2\left(\frac{f(\bar{x},p)}{\mu_1^H f(p)}\right) + f(\underline{x},p)\log_2\left(\frac{f(\underline{x},p)}{(1-\mu_1^H)f(p)}\right)\right] dp, \quad (47)$$

and

$$\mathcal{H}[g^L(\boldsymbol{x})] - \mathbb{E}_p\mathcal{H}\left[f(\boldsymbol{x}|p)\right] = \int_p \left[f(\bar{x},p)\log_2\left(\frac{f(\bar{x},p)}{\mu_1^L f(p)}\right) + f(\underline{x},p)\log_2\left(\frac{f(\underline{x},p)}{(1-\mu_1^L)f(p)}\right)\right] dp, \quad (48)$$

respectively.

In the first period, the firm solves

$$\begin{aligned} p_1^E = \arg\max_{p_1}\{&\mu_0^H p_1\psi\left(p_1,\bar{x}\right) + \left(1-\mu_0^H\right)\psi\left(p_1,\underline{x}\right) \\ &+\hat{L}(\mu_1^H)(\mu_0^H\sigma + (1-\mu_0^H)(1-\sigma)) \\ &+\hat{L}(\mu_1^L)(\mu_0^H(1-\sigma) + (1-\mu_0^H)\sigma)\}, \end{aligned} \quad (49)$$

where μ_0^H is the prior belief at the beginning of period one described by $g(\boldsymbol{x}) \in \{g(\bar{x}), g(\underline{x})\}$, with $g(\bar{x}) = \mu_0^H$ and $g(\underline{x}) = 1-\mu_0^H$.

The firm chooses the p_1^* that maximizes the profits expectation sum of the two periods. Choosing p_1^* implies determining $\sigma(p)$, and consequently, μ_1^H, and μ_1^L. These two posterior beliefs become the prior beliefs of the second stage. Recall that the rational inattention stage provides an optimized profit value –due to learning via *mutual information*– that does not consider its myopic first stage earnings, so its profit result enters as an input to the second period of the experimentation stage through $L(\mu_1^H)$ and $L(\mu_1^L)$. Since there is no closed-form for the rational inattention stage solution, the joint learning problem does not have one either. However, a numerical solution is easily attainable because each stage is already solved.

Three possible scenarios arise with this setting. First, if the seller has some attention capacity and chooses to experiment in the first period, he could refine his knowledge by paying attention to available signals in period two. Second, if the seller has no attention capacity, he could only choose to experiment. And third, the seller could decide not to experiment and only to pay attention to informative signals. Although this part solves a particular resolution of the two-stage model, Section 5 studies those scenarios in detail.

3.5.2. **Solving the model & algorithm.** The solution to the problem is easy to obtain by backward induction. First, let's consider the experimentation setting to get particular results to this problem. To that end,

$$\mathbb{E}_1[\psi\left(p_1,\boldsymbol{x}\right)] \equiv q_1 = \left[\frac{\boldsymbol{a}}{\boldsymbol{b}} - \frac{p_1}{\boldsymbol{b}}\right]. \quad (50)$$

We consider $\mathbb{E}_1[\psi\left(p_1,\boldsymbol{x}\right)]$ because the problem is under Assumption 1. At the beginning of period one, the prior belief support is given by $\mathbf{a} \in \{10, 20\}$ and $\mathbf{b} \in \{\frac{1}{2}, 1\}$, which is assumed to be the same for the two periods. The duple $\{10, 1\}$ corresponds to the low state of demand, $\underline{x}$, and $\{20, \frac{1}{2}\}$ to the high, $\bar{x}$. Besides, the monopolist knows nothing about the possible realization of the state, so he allocates prior probabilities to each state of the form $\mu_0^H = \frac{1}{2}$ and $(1 - \mu_0^H) = \frac{1}{2}$. Thus, the myopic price is still $p_1(\mu_0^H) = 8.33$. We let the prices be between the monopolistic for every state, $p \in \{5, \ldots, 10\}$, in a grid of 40 equally-spaced values. The set up for $\sigma(p) = m - np$ considers $m = 1.46$ and $n = 0.096$, which implies that $\sigma(p) \in \{\frac{1}{2}, \ldots, 0.98\}$ for $p \in \{5, \ldots, 10\}$. As in the pure experimentation problem, this setting for $\sigma(p)$ implies that the pricing decision makes the monopolist get more information of the state of demand as he lowers prices; he cannot unlearn in the process, and he cannot get perfect knowledge as well.

Secondly, let's consider the rational inattention setting. Call to mind that the profit function can be written as

$$\Pi(\boldsymbol{x}, p) = p \left[\frac{\boldsymbol{a}}{\boldsymbol{b}} - \frac{p}{\boldsymbol{b}} \right]. \tag{51}$$

The support of the parameters of the demand curve is the same as in the experimentation problem. The duple $\{10, 1\}$ corresponds to the low state of demand, $\underline{x}$, and $\{20, \frac{1}{2}\}$ to the high, $\bar{x}$. The vector of signals (number of prices) to attend is 40. As in the original setting, they are equally-spaced between the monopolistic optima, *i.e.*, $p \in \{5, \ldots, 10\}$. Finally, in the *mutual information* constraint, $\kappa = 0.5$. Note that the functional form of the profit function is the same for each stage.

We keep the experimentation code. The fundamental step to solve the sequential problem is to replace the expected profits resulting from the second period of experimentation with the ones resulting from the rational inattention problem. To that end, we run the rational inattention stage using as prior belief each element of $\mu_1^H(i, 1)$ and $\mu_1^L(i, 1)$, where $i \in \{1, \ldots, 40\}$. We have two vectors of forty posterior beliefs: the rational inattention stage has to be solved eighty times. With these results, we build the profits matrix of period two as $L(\mu_1^H)(i, 1)$ and $L(\mu_1^L)(i, 2)$. Finally, we find the maximum value of the total expected profits vector given by the experimentation code and search for the same position i, j in the vector of prices: that position contains the experimentation price in period one that maximizes the total expected profits. This numerical process could take at least one hour, depending on hardware capacity.

The optimal experimentation price in period one that maximizes expected profits is $p_1^* = 8.07$, which is lower than $p_1(\mu_0^H) = 8.33$, but higher than in the pure experimentation problem, $p_1^E = 7.94$. In other words, the monopolist experiments less –he is closer to the myopic price– because now he has another instance to learn and improve expected profits, in which he pays limited attention to informative signals about the state of demand. Optimal posterior beliefs after experimentation are $\mu_1^H = 0.68$ and $\mu_1^L = 0.32$ (*i.e.*, $\mathbb{P}(\underline{x}|q^L) = 0.68$). Finally, total profits with this two-stage procedure corresponds to $\mathbb{E}_0\left[\Pi_{1+2}^*\right] = 214.22$, which is higher than if the monopolist decides not to learn via *mutual information* but only experiment, $\mathbb{E}_0\left[\Pi_{1+2}^E\right] = 209.28$. That is to say, the rational inattention stage improves the profit expectations when taking as prior beliefs the posterior beliefs that result of experimenting in period one.

Since at the beginning of period one the uncertainty about the state of demand is maximal (*i.e.*,

$\mu_0^H = \frac{1}{2}$), the *entropy* of the prior distribution equals $\mathcal{H}[g(\boldsymbol{x})] = 1$. This implies that the perfect information knowledge is only achievable when the seller has "unlimited" attention, which in this particular case is represented by $\kappa = 1$. With $\kappa = 1$, optimal prices and profits are the following: $p_1^* = 8.33$ and $\mathbb{E}_0\left[\Pi_{1+2}^*\right] = 216.66$. In this situation, the monopolist chooses not to experiment in period one. He sets myopic prices the first period to learn everything in the second period through *mutual information* as the perfect-information scenario is easy to attain given "unlimited" attention.

4. Strategies Analysis

As you might have noticed in the previous section, the seller complements the use of both learning strategies when facing a binding information acquisition constraint in the second stage. However, he substitutes their use when facing a "big enough" capacity to learn from available signals in the environment. Specifically, if the monopolist can learn everything via *mutual information*, he rules out learning by experimentation. To better understand how these learning strategies behave, it is necessary to examine all possible scenarios that the monopolist can face. We need to consider total expected profits for all combinations of κ (which represents how limited the seller's attention is) and prior beliefs. In other words, when looking at those optimal profits, we will be able to analyze the strategies behavior with different levels of the seller's attention for different prior beliefs of the states. To that end, we consider total expected profits of the two-stage model for all combinations of $\kappa \in [\frac{1}{10}, 1]$ and $\mu_0^H \in [\frac{1}{2}, 0.98]$. The case for $\mu_0^H \in (0, \frac{1}{2}]$ is homologous and symmetrical, so we do not present its results.

Firstly, it is necessary to differentiate between two circumstances regarding expected profits: two zones of profits show up for each prior belief that depend on the value of the information constraint, κ (to make this analysis, we consider that the seller has some attention capacity, *i.e.*, $\kappa > 0$). There will be a region of perfect-information profits (a complete-learning zone), in which the seller learns exclusively under rational inattention, and another region of partial-knowledge profits (a partial-learning zone), which is reachable by using both strategies. In the first one, the seller sets myopic prices in period one and optimal monopolistic prices in period two since he perfectly learns about the demand realization. In the second one, the seller cannot attain perfect information, so he struggles to obtain as much information as possible to set the most accurate prices according to the information received. In particular, there is a theoretical condition that lets us know precisely the zone under which the monopolist gets the perfect-information profits. Suppose the monopolist can capture all the relevant information necessary to set maximizing-profit prices with *mutual information* (*i.e.*, attending to available signals). In that case, he does not need to experiment and sets prices in the second period as if he were in a perfect-information scenario due to a slack information constraint: κ is big enough to set monopolistic prices in period two. The *entropy* of the prior belief distribution determines this slackness condition.

Bear in mind that $\mathcal{H}[g(\boldsymbol{x})] = \mu_0^H \log_2(\mu_0^H) + (1 - \mu_0^H)\log_2(1 - \mu_0^H)$, $\forall \mu_0^H \in [\frac{1}{2}, 0.98]$. Remember that the prior distribution is described by $g(\boldsymbol{x}) \in \{g(\bar{x}), g(\underline{x})\}$, with $g(\bar{x}) = \mu_0^H$ and $g(\underline{x}) = 1 - \mu_0^H$. Then, we can get all values of *entropy* for each prior belief (we measure the uncertainty of each prior distribution). Figure 2 shows this exercise. The black line represents the critical values that κ can take for every prior belief ($\kappa = \mathcal{H}[g(\boldsymbol{x})]$, for each $g(\bar{x}) = \mu_0^H$) so as to learn exclusively with *mutual information*: the region over it outlines a zone of perfect information where the seller learns only by processing available signals. The monopolist will not learn by experimenting with introductory prices in that zone since he has a "wide attention" to get all the information needed to maximize expected profits, given his prior probabilities; he can absorb all the uncertainty of the prior belief distribution with his capacity to process data. This zone of perfect information is described by $\kappa \geq \mathcal{H}[g(\boldsymbol{x})]$, for any $g(\bar{x}) = \mu_0^H$.

Figure 2. $\mathcal{H}[g(\boldsymbol{x})]$ in axis y as a function of $\mu_0^H \in [\frac{1}{2}, 0.98]$:

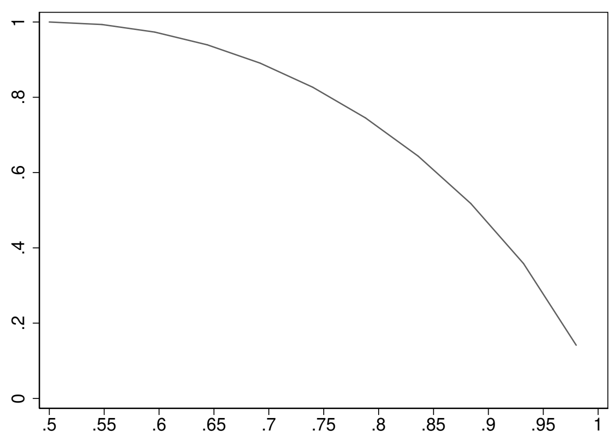

Proposition 1. ***(perfect-knowledge profits).*** *Let $\mathcal{H}[g(\boldsymbol{x})] - \mathbb{E}_s\mathcal{H}[f(\boldsymbol{x}|s)] \leq \kappa$ be the mutual information constraint. For any binary prior distribution $g(\boldsymbol{x})$, the complete-learning profits of the two-stage model as a function of κ and $g(\bar{x}) = \mu_0^H$ are attainable exclusively by mutual information, and they are described by $\kappa \geq \mathcal{H}[g(\boldsymbol{x})]$, with $\kappa \in (0, 1]$.*

Proof. (numerical). We solve numerically all the possible scenarios of expected profits that the seller can face as a function of $\kappa \in \{\frac{1}{10}, \ldots, 1\}$ and $\mu_0^H \in \{\frac{1}{2}, \ldots, 0.98\}$. Therefore, we run the RI-experimentation model with a grid of 10 values for κ and 10 for μ_0^H. Since we already have the code that solves the optimal profits of the model for any κ and any μ_0^H, the numerical procedure is straightforward; we make a nested loop to fill a profit matrix in which every row contains a profit value evaluated in a specific κ, and every column has a profit value evaluated in a particular μ_0^H. Simple but time-consuming: the numerical process takes around 110 hours, depending on hardware capacity. Absolute results for the profit expectations are shown in Table 2.

However, these results provide little clarity on what we seek to assess: the profit matrix does not supply intuitions whether the seller can reach or not perfect-information profits. As we already proposed the region of expected profits under perfect information for each prior belief (*i.e.*, that satisfies $\kappa \geq \mathcal{H}[g(\boldsymbol{x})]$, for any $g(\bar{x}) = \mu_0^H$), we normalize every matrix element of each profit column by its perfect-information profit to build normalized earnings. For example, profits of column 1 were solved with $\mu_0^H = \frac{1}{2}$, and we know that perfect-information profits in that column are only achievable with $\kappa = 1$. Profits solved with $\kappa = 1$ are contained in row one, so we divide each profit value of the first column by 216.66 to normalize perfect-information profits to 1. The procedure to normalize earnings is same for the rest of the columns: we can divide the profit values of each column by any profit satisfying $\kappa \geq \mathcal{H}[g(\boldsymbol{x})]$ in that column. To make it simpler, we just divide the values of each column by every profit in row one since it can represent the perfect-information profit for any $\mu_0^H \in \{\frac{1}{2}, \ldots, 0.98\}$, as you can see in Figure 2.

Table 2. Profits as a function of μ_0^H (columns) and κ (rows):

216.6639	235.7074	254.9086	274.2431	293.6998	313.2712	332.9468	352.7157	372.5708	392.5017
216.6092	235.5633	254.7268	274.0942	293.6408	313.2711	332.9459	352.7158	372.5708	392.5017
216.1883	235.1552	254.3647	273.8316	293.4383	313.1810	332.9461	352.7158	372.5708	392.5018
215.6344	234.5910	253.8289	273.3383	293.0788	312.9738	332.8863	352.7158	372.5708	392.5021
214.9740	233.9306	253.2004	272.7497	292.5439	312.5446	332.6713	352.7158	372.5708	392.5017
214.2167	233.1784	252.4840	272.0765	291.9378	311.9992	332.2406	352.5666	372.5709	392.5017
213.3775	232.3609	251.6942	271.3228	291.2441	311.3896	331.7106	352.1665	372.5709	392.5022
212.4578	231.4628	250.8268	270.5052	290.4832	310.6928	331.1080	351.6718	372.3244	392.5017
211.4692	230.4969	249.8909	269.6199	289.6562	309.9429	330.4328	351.1055	371.8901	392.5017
210.4103	229.4615	248.8974	268.6751	288.7712	309.1213	329.6989	350.4626	371.3846	392.3898

Note. The figure displays profits as a function of: increasing values of $\kappa \in \{\frac{1}{10}, \ldots, 1\}$ from the bottom row to the top row, and increasing values of $\mu_0^H \in \{\frac{1}{2}, \ldots, 0.98\}$ from the first column to the last.

Figure 3(a) reports the normalized expected profits as a function of μ_0^H and κ in axis z, where 1 represents the zone of perfect information (complete-learning) by intense yellow (the flat area). Figure 3(b) shows the support of these expected profits.

As you can visually notice, the critical κ shapes the zone of perfect-knowledge profits in intense yellow (it proves the discussion of Figure 2). This area determines explicitly complete-learning through *mutual information*, and it is only achievable when the seller has enough attention capacity to learn everything about the demand realization by paying attention to informative signals. In this context, he sets myopic prices in period one and monopolistic prices in period two; he does not need to set suboptimal prices in the first period, given that he can perfectly learn about the environment in the second period. The case for $\mu_0^H \in (0, \frac{1}{2}]$ is homologous and symmetrical, so the proof works likewise.

Figure 3. Normalized expected profits as a function of κ and μ_0^H:

(a) Normalized profits:

(b) Normalized profits support:

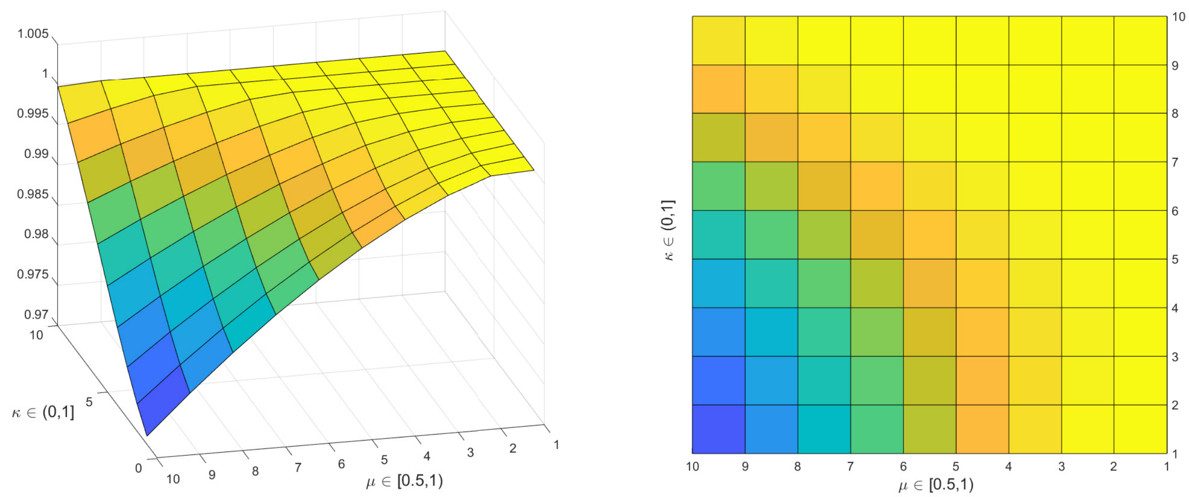

Note. Figure 3(a) displays normalized profits as a function of: increasing values of $\kappa \in \{\frac{1}{10}, \ldots, 1\}$ from the origin in axis $\mathtt{y}$, and increasing values of $\mu_0^H \in \{\frac{1}{2}, \ldots, 0.98\}$ from the origin in axis $\mathtt{x}$. Figure 3(b) displays the normalized profit support as a function of: increasing values of $\kappa \in \{\frac{1}{10}, \ldots, 1\}$ from the bottom row to the top row, and increasing values of $\mu_0^H \in \{\frac{1}{2}, \ldots, 0.98\}$ from the first column to the last.

■

However, what happens outside the intense-yellow zone in Figure 3 needs careful analysis (equivalently, beneath the black line in Figure 2). We would then like to see if there is any zone in which the monopolist complements the use of both strategies and attains the perfect-information profits. Easy, although not least, we also have to analyze the regions where the monopolist cannot reach the perfect-information scenario and the area where he only experiments.

Four conclusions arise from these results. First, the intense-yellow zone determines complete-learning exclusively under rational inattention. Let's analyze in detail what happens for every $\mu_0^H \in [\frac{1}{2}, 0.98]$ in that area. Recall that, as the prior belief distribution is binary, the *entropy* is maximal, $\mathcal{H}[g(\boldsymbol{x})] = 1$, at $g(\bar{x}) = \mu_0^H = \frac{1}{2}$. When the prior belief about the high state of demand gets closer to $\mu_0^H = \frac{1}{2}$ (it is either probable to face the high or low demand), the seller's knowledge about the demand realization is more uncertain, which means that he needs more attention capacity to process information to reach the perfect-information profits. The seller can absorb all the uncertainty of the prior beliefs (he learns everything about the state of demand) when κ is larger than the prior distribution uncertainty. Instead, when we move towards $\mu_0^H = 1$, the seller's knowledge about the demand realization is less uncertain since it is more probable to face the high state of demand (*a priori* he knows more). It implies that he needs less attention capacity to process information to reach the perfect-information profits. In other words, if the attention capacity to process informative signals is big enough to learn all that he does not know (represented by the uncertainty of the prior distribution), then the seller can attain a perfect-knowledge scenario. Thus, he sets monopolistic prices in period two and does not experiment in period one since it is not necessary. For any κ over the black line in Figure 2, the firm faces an exclusion zone: learning under rational inattention completely substitutes the use of the experimentation strategy.

Let's now analyze what happens for a given value of κ. When we take κ as given, it is indisputable that profits deviate more from perfect-information earnings when prior uncertainty about the state of demand is higher (*i.e.*, μ_0^H closer to $\frac{1}{2}$), because it is harder for the monopolist to acquire information when prior knowledge is scarcer. On the contrary, as we move towards $\mu_0^H = 1$, the perfect-information profits are easier to attain, because the seller needs less capacity of that particular κ to "absorb" the prior uncertainty (his prior knowledge is higher).

Nonetheless, how realistic this scenario is, is debatable. We can think of κ as the installed capacity of a firm to attend to available information in the environment (how much attention it can pay). For example, its sales department's ability to survey customers about their preferences or its capacity to learn about competitors' prices. Only huge companies might have the ability to be close enough to the complete-learning scenario; they could have the ability to learn from the competitors even when launching new products (avoiding introductory prices). However, how close they could be is a task even hard for empirical analysis. We can take one fact for granted: small and medium-sized enterprises do not have such an attention capacity because they do not have enough resources. They do not have the buildings, the networks, and they cannot keep many employees researching markets, which implies less market-research time. These firms cannot exploit learning under rational inattention as they would like to (to set monopolistic prices or close enough to them).

Second, the seller never complements the use of both strategies to get expected profits of the complete-learning scenario. There exists no zone in which the firm implements the two learning methods and reap those profits. When looking at the pricing decisions in period one, the firm starts experimenting only if it does not have enough ability to process informative signals. That is to say, when the seller has enough capacity to learn from the market (represented by a slack *mutual information* constraint), paying attention to the signals is the monopolist's rational answer to set optimal prices. He has no incentives to experiment because all his available attention makes the information acquisition costless; he just has to use his attention capacity to process signals. In this situation, experimenting is unfavorable because he has to pay for generating signals (the observed quantity) via losing profits, which is a cost that he does not pay while using his attention capacity. For example, a firm with enough capacity to know its consumers' preferences by surveying them to set optimal prices for the next period would not ever set introductory prices today. Its manager is not willing to do it since he can use the tools the firm already has to capture information from the market without losing profits due to suboptimal prices. To fix the idea, the company already pays its workers a salary, and they can survey the firm's customers without incurring extra costs.

Third, the seller complements the use of both strategies if he faces a combination of κ and prior beliefs outside the intense-yellow area, in which case he can only get partial-knowledge profits (partial-learning). Since he cannot pay attention to all the informative signals he would like to, he also decides to experiment to refine his knowledge about the state of demand. This combined strategy is his best answer to uncertainty when attention is severely bounded. We could think of medium-sized companies as an example of this situation; they have a limited capacity to learn from market signals. They do not have enough resources to keep big sales departments or a considerable amount of employees to capture available information from the environment. Still, they must learn about the demand when launching new products or capture switching macro or micro-variables to maximize expected profits. They can achieve this aim by setting introductory prices, generating information about the faced demand, instead of paying attention to signals.

However, there is a critical zone hard to assess because of severe hardware limitations. Although the partial-learning profits evaluated in combinations of prior beliefs and values of κ further apart from $\kappa = \mathcal{H}[g(\boldsymbol{x})]$ show both learning by experimentation and *mutual information*, we can not capture experimentation in the zone of profits right around $\kappa = \mathcal{H}[g(\boldsymbol{x})]$ (right below the black line in Figure 2 or right leaving the intense-yellow zone in Figure 3). We used forty grid points to discretize the price vectors to get expected profits in the two-stage model, which is not enough to capture tiny price movements. As we move towards that κ's critical values, we need more grid points to get the experimentation price. In particular, we would need at least 300 grid points "smartly distributed" around the myopic price of period one in the price vector to be sure of getting it, which would imply running the rational inattention stage 600 times (even more, considering the non-smartly-distributed grid points to complete the price vector between the monopolistic prices). Besides, we would need a profit matrix larger than 10x10. Since it took around 110 hours to get Figure 3, we could only evaluate isolated elements of profits close to the critical zone in a Server with more hardware capacity. On the contrary, when we move further apart from those critical values, introductory prices deviate more from those myopic, so they are easier to capture by a small grid of prices. In other words, the seller experiments more as he moves away from the complete-learning area. For a given $\kappa \in \{\frac{1}{10}, \ldots, 1\}$, the experimentation is more aggressive –introductory prices deviate more from myopic prices– when moving towards $\mu_0^H = \frac{1}{2}$, given that prior uncertainty is greater (the seller knows less about the demand realization). For a given $\mu_0^H \in \{\frac{1}{2}, \ldots, 0.98\}$, the experimentation is more aggressive when moving towards $\kappa = 0$, since the learning capacity to capture informative signals is more bounded.

Fourth and final, and as it may seem obvious, the closer κ is to 0, the closer the profits are to those of pure experimentation, for any prior belief. As the firm cannot process informative signals (*i.e.*, $\kappa = 0$), the only way to improve its demand knowledge is setting introductory prices. In our setting, the firm experiments for any $\mu_0^H \in [\frac{1}{2}, 0.98]$. However, the experimentation prices deviate more from those myopic as we move towards $\mu_0^H = \frac{1}{2}$; therefore, when we are getting closer to $\mu_0^H = 0.9$ it is harder to capture the experimentation pricing numerically.

Proposition 2. ***(intensity of experimentation).*** *Let $\kappa = 0$. For any binary prior distribution $g(\boldsymbol{x})$, the difference between the myopic and experimentation prices of the two-stage model is decreasing in $\mu_0^H \in [\frac{1}{2}, 1)$, and increasing in $\mu_0^H \in (0, \frac{1}{2}]$.*

Proof. (numerical). We solve the two-stage model with arbitrarily five increasing values of μ_0^H for $\kappa = 0$; $\mu_0^H \in \{0.5, 0.6, 0.7, 0.8, 0.85\}$. We denote myopic prices as p_1^*, experimentation prices as p_1^E, the difference between myopic and experimentation prices as *Diff*, and grid points of the price vectors as *GD*. Myopic prices result from solving equation (36). Recall that prices in the resolution of the model are defined as $p \in [5, 10]$. The grid of the price vectors is equally-spaced for $\mu_0^H \in \{0.5, \ldots, 0.8\}$; however, the grid of the price vectors for $\mu_0^H = 0.85$ contains 385 grid points for $p \in \{9, 10\}$ and 15 for $p \in \{5, 9\}$. Results are shown in Table 3; the case for $\mu_0^H \in (0, \frac{1}{2}]$ is homologous and symmetrical, so we do not present its results, and the proof works likewise.

Table 3. Decreasing difference of prices, $p_1^* - p_1^E$, in μ_0^H:

μ_0^H	0.5	0.6	0.7	0.8	0.85
p_1^*	8.33	8.75	9.1176	9.4444	9.5946
p_1^E	7.9293	8.5354	9.0452	9.4264	9.5885
Diff	0.4007	0.2146	0.0724	0.018	0.0061
GP	100	100	200	350	400

■

These results are intuitive. The more the seller knows about the demand realization, the less he experiments. When his prior beliefs hold less uncertainty about the state of demand, he needs to generate less information to set more accurate prices in period two. For example, when the seller guesses that each state has a 50% chance to occur, but he cannot pay attention to signals in the market since $\kappa = 0$, then the deviation from myopic prices is maximal to learn as much as possible about the demand curve. When the seller guesses that the high state of demand has an 80% chance to occur, then the deviation from myopic prices is little because he already knows which is the most likely realization of the demand curve, so he does not need to set very low prices to get information. When moving towards $\mu_0^H = 1$, the experimentation strategy gets worthless. Furthermore, it is possible that even continuous and dense grids do not capture experimentation when the prior belief gets very close to $\mu_0^H = 1$ because the seller knows too much about the realization of the state, which implies that experimentation might not offset the profit loss in period one by allowing a more accurate price-setting decision in period two. In particular, when $\mu_0^H = 1$, there is no experimentation since the monopolist has perfect information about the demand realization.

We can think of firms that frequently set introductory prices as small-sized sellers with no capacity to learn from market signals. Whenever they need to start selling products (or when they face demand shocks), the only strategy to capture their demand curves and learn about them is by lowering prices. They would like to get available information from the market, but they have no attention capacity to do it.

5. Discussion

This brief section discusses why a reverse order in the two-stage model is not a plausible formulation. The model does not support experimenting in the second stage since it changes the optimization problem's nature; to understand why we need to focus on the rational inattention strategy. At this stage, the requirement to have a numerical solution is that the payoffs are determined by the profit function, which implies an objective function linear in the joint distributions. This setting allows an easy way to get the optimal joint distributions because the payoffs depend only on the grid of prices. Therefore, chosen prices depend on the state. Since we solve sequential optimization problems by backward induction, when the rational inattention step is the last period to solve, then the numerical formulation satisfies that requirement (but it does not when it is not the last one). When the experimentation problem is the second stage to solve, we need to find an expression for optimal prices as a function of given probabilities (that we would maximize in the first period). These expressions are the input for the rational inattention stage in the backward induction problem. Consequently, the profit function payoffs in period one now depend on the prices as a function of given probabilities that we are bringing from period two, not on a simple grid of prices. This drawback gives rise to prices that do not depend on the state, a maximization problem nonlinear in the joint distributions, and unknown profit function payoffs, making the optimization problem analytically intractable.

6. Conclusion

This paper put forward a two-stage model implementing experimentation and rational inattention and addressed a monopolistic seller's optimal pricing decisions who faced an uncertain demand curve. We examined how the firm's limited attention affects its decision on what strategy to put into practice to learn from the uncertain demand. Scarce attention is the crucial factor the seller considers when deciding whether to learn by experimentation or by paying attention to informative signals in the environment. We have shown that he prefers using his attention capacity to learn about the market before setting introductory prices since the latter strategy is costly in terms of profits.

Beyond what we show in this paper, interesting further research is feasible about this model. Many information theory models in macroeconomics distinguish between *idiosyncratic* and *aggregate* types of data and consider that firms can learn both kinds of information. An interesting approach would be, on one side, taking the demand curve level as aggregate data, which is available information (such as production level, interest rates, GDP growth, or monetary policy shocks) that the firm can process via rational inattention (in particular, *mutual information*). On the other, the demand curve slope can represent idiosyncratic data (for instance, the particular demand for a certain quality of a product or consumer preferences), which is specific information to the firm that is not freely available. For this purpose, the firm can generate it via experimentation. That modeling could sophisticate and make more tangible, for instance, intuitions behind the agents' learning in the Lucas Islands Model.[8] At it, Robert Lucas Jr. tries to explain how agents discern between monetary policy shocks and changes in demand, and plenty of literature was developed in macroeconomics after publishing his framework. Nevertheless, the studied therein focuses on its methodological impact, using imperfect information and rational expectations, rather than how agents acquire knowledge to explain optimal pricing patterns –a void that this model might conceptually fill–.

Finally, another plausible, exciting extension to the model would be to set it up as a full dynamic formulation. The model presented in this paper can explain the price patterns of a monopolist that actively learns of two types of information at a static level, but it cannot explain the price series over time. It would be interesting to explore how prices behave when there is coordination between the two learning strategies over time, leading to a new possible way to assess price stickiness. Furthermore, a dynamic formulation of a model that considers the demand curve level as aggregate data and the slope as idiosyncratic data can also be interesting and attainable.

[8] Robert Lucas Jr. presented this simple but influential framework in a series of three papers in the 1970s, particularly in Lucas (1973).

References

Bagwell, K. (1990). Informational product differentiation as a barrier to entry. *International Journal of Industrial Organization*, *8*, 207–223.

Caplin, A., Dean, M., & Leahy, J. (2019). Rational inattention, optimal consideration sets, and stochastic choice. *Review of Economic Studies*, *86*, 1061–1094.

Chaves, J. C. (2019). The less I know the better? A model of rational attention and experimentation. *Social Science Research Network*, *Documento CEDE No. 2019-43*.

Guadalupi, C. (2018). Learning quality through prices and word-of-mouth communication. *Journal of Economics and Managment Strategy*, *27*, 53–70.

Jung, J., Kim, J., Matějka, F., & Sims, C. (2019). Discrete actions in information-constrained decision problems. *Review of Economic Studies*, *86*, 2643–2667.

Kahneman, D. (1973). *Attention and effort*. Prentice-Hall.

Levitt, B., & March, J. (1988). Organizational learning. *Annual Review of Sociology*, 319–340.

Lucas, R. E. (1973). Some international evidence on output-inflation tradeoffs. *American Economic Review*, *63*(3), 326–334.

Matějka, F. (2016). Rationally inattentive seller: sales and discrete pricing. *Review of Economic Studies*, *83*, 1125–1255.

Matějka, F., & McKay, A. (2015). Rational inattention to discrete choices: A new foundation of the multinomial logit model. *American Economic Review*, *105*(1), 272–298.

Milgrom, P., & Roberts, J. (1986). Price and advertising signals of product quality. *Journal of Political Economy*, *45*, 393–406.

Mirman, L. J., Samuelson, L., & Urbano, A. (1993). Monopoly experimentation. *International Economic Review*, *34*(3), 549–563.

Pomatto, L., Strack, P., & Tamuz, O. (2020). The cost of information. R&R at the *American Economic Review*.

Schlee, E. (2001). Buyer experimentation and introductory pricing. *Journal of Economic Behavior and Organization*, *44*, 347–362.

Simon, H. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, *69*, 99–118.

Simon, H. (1959). Theories of decision-making in economics and behavioral science. *American Economic Review*, *49*(3), 253–83.

Simon, H. (1979). Information processing models of cognition. *Annual Review of Psychology*, *30*, 363–396.

Sims, C. A. (1998). Stickiness. *Carnegie-Rochester Conference Series on Public Policy*, *49*, 317–356.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, *50*, 665–690.

A. Incidental proof

A.1. Mutual Information

Proof. Using the definitions of *entropy* and *conditional entropy* in a continuous setting, we have

$$\mathcal{H}[g(\mathbf{x})] = -\int_{\mathbf{x}} g(\mathbf{x}) \log_2 g(\mathbf{x}) d\mathbf{x}$$

and,

$$\mathcal{H}[f(\mathbf{x}|p)] = -\int_{\mathbf{x}} \int_{p} f(\mathbf{x}, p) \log_2 f(\mathbf{x}|p) dp d\mathbf{x}.$$

Therefore, we can write

$$\begin{aligned}
&\mathcal{I}[f(\mathbf{x}, p)] = -\int_{\mathbf{x}} g(\mathbf{x}) \log_2 g(\mathbf{x}) d\mathbf{x} + \int_{\mathbf{x}} \int_{p} f(\mathbf{x}, p) \log_2 f(\mathbf{x}|p) dp d\mathbf{x}, \\
\iff &\mathcal{I}[f(\mathbf{x}, p)] = -\underbrace{\int_{p} f(p|\mathbf{x})}_{1} \int_{\mathbf{x}} g(\mathbf{x}) \log_2 g(\mathbf{x}) dp d\mathbf{x} + \int_{\mathbf{x}} \int_{p} f(\mathbf{x}, p) \log_2 f(\mathbf{x}|p) dp d\mathbf{x}, \\
\iff &\mathcal{I}[f(\mathbf{x}, p)] = -\int_{\mathbf{x}} \int_{p} \underbrace{f(p|\mathbf{x}) g(\mathbf{x})}_{f(\mathbf{x}, p)} \log_2 g(\mathbf{x}) dp d\mathbf{x} + \int_{\mathbf{x}} \int_{p} f(\mathbf{x}, p) \log_2 f(\mathbf{x}|p) dp d\mathbf{x}, \\
\iff &\mathcal{I}[f(\mathbf{x}, p)] = -\int_{\mathbf{x}} \int_{p} f(\mathbf{x}, p) \log_2 g(\mathbf{x}) dp d\mathbf{x} + \int_{\mathbf{x}} \int_{p} f(\mathbf{x}, p) \log_2 f(\mathbf{x}|p) dp d\mathbf{x}, \\
\iff &\mathcal{I}[f(\mathbf{x}, p)] = \int_{\mathbf{x}} \int_{p} f(\mathbf{x}, p) [\log_2 f(\mathbf{x}|p) - \log_2 g(\mathbf{x})] dp d\mathbf{x}, \\
\iff &\mathcal{I}[f(\mathbf{x}, p)] = \int_{\mathbf{x}} \int_{p} f(\mathbf{x}, p) \log_2 \frac{f(\mathbf{x}|p)}{g(\mathbf{x})} dp d\mathbf{x},
\end{aligned}$$

which by Bayes' theorem is

$$\mathcal{I}[f(\mathbf{x}, p)] = \int_{\mathbf{x}} \int_{p} f(\mathbf{x}, p) \log_2 \frac{f(\mathbf{x}, p)}{g(\mathbf{x}) f(p)} dp d\mathbf{x}.$$

■

This result gives us a tractable *mutual information* expression to solve the problem.