Transaction Costs in the Trading of Variable Quality Commodities

Working Paper*

Yoram Barzel Aurora Stephany

December 6, 2017

Abstract

Measurement is a part of every transaction. When buyers can inspect non-uniform goods, they will try to get the best ones. Measurement is always costly, and some resources will be dissipated; mutually beneficial trade will be forgone in the absence of measurement information. In this paper, we develop a search model in which a seller has a batch of goods of varying quality. Pricing each good individually is prohibitively expensive, so they are all sold at the same price. In the absence of trust between buyer and seller, the former doesn't trust the latter with the selection of units, because the seller would gain by giving him the worst units in order to improve the remaining distribution. So the buyer will spend some resources "picking and choosing", that is, inspecting items until finding an acceptable one. This inspection is costly and will result in trade below the optimal levels. We prove several results: 1., in such circumstances the good will be necessarily sold at a price above the average value, with buyers still participating in the market if they can inspect until they find items from the top of the distribution. 2. the distribution will decay over time, and the seller will be forced to lower the price. 3.greater dispersion in the quality of the distribution will deepen the problems mentioned above, 4. when possible, the seller will spend resources making individual inspection costly while making general inspection (assessment of the whole sample) cheaper for buyers, 5. the seller will be willing to incur a cost in order to increase the uniformity of the goods; and 6. when there are different types of buyers, the ones with lower cost of inspecting could drive away the ones with higher inspection cost. When this problem is serious enough, the market for the commodity may not exist. The results in this paper help explain a wide variety of observed phenomena in the markets: why are oranges displayed in a pyramid, why pre-selected fruit is cheaper, why supermarkets with a heterogeneous customer base tend to sell either very uniform produce or pre-bagged produce, why is there a "second hand" vegetable market, and why so many resources are spent making retail produce homogeneous.

*The authors are very grateful for comments received from Doug Allen, Chris Anderson, Patrick Bajari, Lee Benham, Mike Freimuth, Cheryl Long, Scott Masten, Deirdre McCloskey, Qing Zhang, and participants of the 2017 SIOE conference. All errors are our own.

1 Introduction

Every year, almost one third of agricultural produce is discarded while still being edible, at various points of the production and distribution chain[1] [?]. . A mutually beneficial exchange requires, before anything else, a measurement: the buyer and the seller must "measure" the merchandise; that is, assess its characteristics, in order to even consider whether the exchange is beneficial or not. Despite being part of every transaction, this initial step, the acquisition of information, is often overlooked in economic models [2]. This omission, we think, leaves very common economic phenomena unexplained: for example, if measurement was indeed costless buyers would establish at the time of the purchase the true value of the item, so there would be no need for quality seals, brand names, or reputation; because buyers would decide whether to purchase or not by comparing this true value to the price.

Measurement becomes all the more relevant to trade for cases in which commodities are not homogeneous [3]. Indeed, if goods are not all identical, they will probably not be equally valuable, so the ability to tell them apart - measure them - becomes central to the economic problem.

In this paper, following the framework proposed by Barzel(1982) [?] , we will focus on the measurement issue. We propose that it should not be brushed aside when modeling trade: for measurement, its costs, and the subsequent optimization associated with it have a profound effect not only on the materialization or not of each exchange, but also on the way that markets are organized.

According to Buzby et al. (2015) [?] the supermarket "shrink"-the inventory that is not sold-in the United States for 2005 and 2006 averaged 11.4 percent for fresh fruit and 9.7 percent for

[1] See Gustavsson *et. al.* (2011

[2] In most models, economic agents are assumed to have either perfect information about an issue or no information at all. The possibility of a cost function for the acquisition of information is not a common feature of information-centered models, and agents who optimize in all other instances seem to accept the given structure of information without any attempt to improve their options by paying to have better information.

[3] Strictly speaking, one could argue that commodities are always heterogeneous, since absolute uniformity is a physical abstraction. However, we refer here to cases in which the variation between units can be detected with technology readily available to buyer or seller.

fresh vegetables by weight, and about 4.2 billion pounds of edible produce are lost every year[4]. So most of the "shrink" does not happen due to mishandling, theft or spoilage, some produce is simply not sold at the posted price and ends up being discarded. See FAO(2011) [?]

On the other hand it is a common observation that fruit sold in bags is less uniform than fruit sold by units, and bagged fruits are sold at a lower price per unit. A rational consumer would not be systematically fooled, so it is not likely that bagged fruits are an attempt to deceive the buyer. This is confirmed by the multiplicity of recent attempts by supermarkets to sell pre-selected or bagged produce while acknowledging their heterogeneity but insisting on their being of equal quality *on average*[5].

We construct a search model with non-uniform commodities, and explicitly incorporate the cost of measuring the quality of each unit. We show that if the cost of inspecting units is low enough, buyers will pick and choose . The seller is better off making the selection for the buyers; but sellers have an incentive to give each buyer a low quality unit and improve the average quality of his remaining batch, and buyers are aware of this incentive. So,if the seller is unable to commit (be it through reputation, brand name, guarantee, legal liability or some other mechanism) then the buyer has no reason to trust him and will need to inspect the unit.

But when units are inspected, only those at the top of the distribution are picked, which over time leads to a deterioration of average quality. Inspection is also inefficient because it is a cost to the buyer that does not become a revenue to the seller. An arrangement that forces buyers to pick at random- and convinces them that other buyers are also picking at random - could increase the total gains from the transaction, and in some cases even represent a Pareto improvement. We prove several results that occur because of heterogeneity of the goods, and then some results that occur when buyer heterogeneity is also considered.

[4]According to the Food and Agriculture Association, in developed countries most food wastage occurs at the retail level, unlike in developing countries where it occurs at the farm selection level

[5]See for example, Wal-Mart's brand "I'm Perfect" of apples, which have some cosmetic problems but are nutritionally equivalent and sold at a discount. There is no reason why this apples would be sold in a bag rather than simply at a cheaper unit price than other apples. Our model explains this fact.

Our first set of results follows from the model even if only one type of buyer is considered. By only one "type" of buyer we mean that all buyers are considered to have the same inspection costs. While not at all realistic, this initial assumption allows us to discern early on which of our results are due simply to commodity heterogenity. First, if the cost of inspection is low enough, buyers will inspect instead of picking at random, which means that only the best units will be picked, so optimal pricing strategy by the seller is to set a price *higher* than the average value of the commodity. Second, because of buyer picking and choosing, the average quality of the distribution will decay over time, so price will necessarily have to decrease at least once. Third, greater dispersion intensifies the problems above because it leads to more inspection. Fourth, it is desirable for the seller to make individual inspection costly to the buyer, but general inspection cheap. Fifth, the seller (or the producer) will be willing to incur a cost in order to make the goods more homogeneous.

We then extend the model to consider buyers with different inspection costs. This could be due simply to different opportunity costs of time, although other factors -like eyesight and training- are likely to affect inspection costs. When we include the possibility that buyers are different, the ones with lower cost of inspecting have higher expected gains from the transaction and could drive away the ones with higher inspection cost. This, however, is overall inefficient, because the buyers with higher inspection costs are the ones more likely to pick at random, which, as mentioned before, yields a higher total benefit from trade.

In section 2 we look into some characteristics of markets that would be puzzling if measurement costs were not considered. In section 3 we briefly survey the existing literature on information and transaction costs, particularly the framework proposed by Barzel(1982). In section 4 we develop our model and derive our hypothesis as testable implications. In part 5 we explain how our results can be tested and validated, and how they help explain the puzzles suggested in section 2. Part 6 offers conclusions and some unanswered questions that may motivate further research.

2 Some Motivating Observations

2.1 How is Fruit Sold

Sometimes in supermarkets certain fruits are sold in sealed bags while simultaneously being displayed in units to allow buyers to pick and choose. Other supermarkets offer only prepackaged produce, while some offer only loose fruit to be selected by the customer. How do profit-maximizing retailers choose which of these alternatives will be made available for each type of fruit or vegetable in each store? Casual observation suggests that high-end grocers offer near identical units and allow buyers to select their purchase, retailers in an airport seldom give the buyer much choice, large supermarkets often offer both presentations, street vendors will pick in front of the buyer.

The bagged fruit tends to be less uniform than the displayed one. In fact, produce is more likely to be sold in presealed bags when it is less uniform: there have been several recent attempts to market "ugly fruit": siamese carrots, deformed cucumbers, or smaller apples are now part of Wal-Mart's offer, with products such as "I'm Perfect" apples and "Spuglies" potatoes being sold in prepackaged bags and initiatives such as Imperfect Produce offering to deliver less pretty produce at a discount. As of this writing, they claim to have "rescued" 7.6 pounds of food[6]. Although these mechanisms offer savings by accepting less homogeneous food, they eliminate the possibility of a buyer cherry-picking to find the best units.

An orange pyramid is a display commonly found in many farmer's markets, supermarkets, and street vendors. Although it is a volume-efficient arrangement, a pyramid is not a particularly easy display to achieve, so perhaps there is a reason for it being the preferred shape across so many types of retailers in so many different markets.

[6]See https://www.imperfectproduce.com

2.2 The Fruit Not Sold

The large quantities of edible food that are thrown away before they reach the consumer have been a source of alarm by both profit-maximizing supermarkets and environmental and sustainability crusaders. According to the Food and Agriculture Organization, the annual total waste of edible food in 2012 was 1.6 billion tonnes, out of a total production of 6 billion tonnes. High income countries present much larger percentages of food waste[7].

This food waste can occur at several points of the supply chain: producers (farmers) discard as much as one third of their crop (both pre and post harvest) right away because it fails to meet uniformity and other quality standards[8] [**?**]. Supermarkets throw away large quantities of unsold produce. Buzby *et. al.* (2014) estimates for several varieties of fruit the proportion that is thrown away by supermarkets while still being edible. They estimate that in 2010-11 on average 11% of oranges 9% of grapefruits and 3% of lemons were discarded while still edible.

Over the past few decades, the uniformity of the produce that reaches consumers has increased considerably, due to pre sorting and genetic modifications. This evolution has come in some cases at the expense of flavor, nutritional value or even portability or longevity[9] .

2.3 other examples

Wood panels in hardware stores Fishing contracts (sell all catch to one processor) Marble for sculptors Cattle??? Little dogs??

land where some oil or natural gas is possible (buyer has to buy a bunch of acres before determining which ones are the ones with oill in them) specialty coffee

online songs: iTunes charges the same price for any song. This is not a rival good so I'm not sure this example is worth exploring much.

[7]FAO 2013

[8]See FAO 2011

[9]Some varieties of apples, although allegedly tastier and more durable, are no longer grown because the fruits were not similar enough (Johnson, apples). Cavendish bananas are reportedly very bland in flavor and less durable than Gros Michael bananas.

potted plants, the last ones are ugliest

3 The Literature on Information Costs

The transaction costs approach, also referred to as the property right approach, focuses on the persistent incompleteness of ownership over assets. An "owner" never truly has full rights to use, prevent others from using, and benefit from "his" assets. Whatever rights are not enforced, are left in the public domain and economic agents will spend resources to capture them[10] . These costs are transaction costs.

A common source of transaction costs -one could argue the main one- is the inability to measure the attributes of assets, or, put another way, the costliness of gathering information. Coase-like solutions to externalities would be more easily implemented if information asymmetries were easier to mitigate. And when assets in a certain category are not uniform, it does not suffice to measure one to know all the rest.

Note that quality variability does not pose a problem for classical economic models under costless commodity information. All that is needed to extend the analysis to accommodate quality variability is adding a quality dimension to the D-S model. Under costly information, however, this is not so straightforward.

Specific characteristics, i.e., the proportions of different attributes in different commodity specimens, vary, and getting informed about them is expensive. As a result, sellers do not have complete economic rights over their commodities even if they are their legal owners, as they are unable to extract their full value in trade. The competition to capture these rights is dissipating. How is the non-uniformity to be incorporated in the D-S model and what are its consequences?

How are non-uniform commodities to be exchanged? Most straightforwardly, each specimen may be sold individually, perhaps in auction. But auction is expensive, especially for commodities

[10] For example,WHAT IS THE BEST EXAMPLE HERE

with attributes that are not easy to observe at the time of the transaction, such as durability. Moreover, under auction, multiple individuals measure each specimen, which duplicates information. In any case, auction is impractical for most commodities (retail tomatoes, movie theater tickets).

The problem of selling non-uniform commodities becomes less acute if they are sorted or graded. But besides being costly, the optimal levels of sorting and grading are not such as to make commodities uniform, so some variability remains within each category, and welfare-dissipating activities may still take place.

Microeconomic models of buyer behavior under incomplete information have for the most part focused on price, rather than commodity, information[11] : buyers know the characteristics of the good they intend to buy but not the price in each store, so they spend some time searching[12]. In Salop and Stiglitz (1977) [?] buyers differ in their ability to observe and interpret market prices, Although costly information is considered, this information is about prices, not commodities, since all goods are identical. It is unclear how price dispersion can prevail over time with only one uniform good being traded. It seems that over time a new class of agents would emerge and specialize in bargain hunting and information gathering, and price dispersion would be arbitraged away.

The second class of information models regards information about commodities, not (only) prices. For this to be meaningful, goods cannot be assumed to be homogeneous. Differentiation is, however, treated as a strategic variable in most models where goods are not uniform. In a space of characteristics, producers decide where to be so as to maximize profit.

[11] This branch of microeconomics was started by Stigler (1961) [?] The Economics of Information, in which he notes: *"Price dispersion is a manifestation -and indeed, it is the measure- of ignorance in the market. Dispersion is a biased measure of ignorance because there is never absolute homogeneity in the commodity if we include the terms of sale within the concept of the commodity. Thus some automobile dealers might perform more service, or carry a larger range of varieties in stock, and a portion of the observed dispersion is presumably attributable to such differences. But it would be metaphysical, and fruitless, to assert that all dispersion is due to heterogeneity"*. Although he claimed price dispersion was a measure of ignorance, his model does not in fact explain price dispersion. But his work displays early discomfort with the overly simplifying Marshallian and Walrasian assumptions of perfect information, which is so blatantly contradicted by even the most casual observation.

[12] For a survey of these early models, see Rothschild(1973) [?].

When it comes to variation in quality, as opposed to variation in price, some early models deal with strategic behavior by the seller or double moral hazard surrounding the quality level, but to our knowledge the measurement issue itself is not included in any of these models. Laffont and Maskin (1987)[?] propose a model with a similar setup to ours: Quality can be observed by the buyer after purchase (say, during consumption) but guarantees are not possible because quality is not verifiable by a third party. Buyers do not come to the market often enough for reputation effects to take place. This set of assumptions is quite similar to ours, except for one fundamental issue: that of measurement. In their model, the seller observes quality perfectly at zero cost and buyers are unable to observe it at all. This "all or nothing" approach to information gathering makes it impossible to include the measurement itself as an optimization margin subject to the usual treatment –spend resources gathering information as long as the marginal cost of acquiring it is less than the expected welfare gain from having it.

As will be shown in our model, including measurement as part of the decision process of both buyer and seller makes it possible to describe a wider range of observed phenomena than that found in the literature describing interactions between buyer and sellers.

4 Model

4.1 The Marketplace under Study

In this section we develop our model. First, it is convenient to pinpoint exactly what is the situation being modeled and what are our assumptions. The transaction we study is of a seller (he) who has a large batch of a certain good. For example, oranges. They are not all identical, of course, but they are all sold at the same price. The seller can choose the oranges for the buyer (she), induce her to pick at random, or allow her to pick and choose[13] . In our setting, the buyer does

[13] All three of these alternatives are observed in real markets: in an online purchase lets the seller choose the units that the buyer ultimately receives. Packaged foods, becuse they look similar to the buyer, are not individually inspected, though of course nobody expects the contents of all cereal boxes to be *exactly* identical, but the differences are not enough to induce inspection. Finally, in farmer markets, sometimes buyers are allowed to pick and choose

not trust this seller, so she will want to select the units herself (either at random or by inspection), to avoid being given the worse oranges.

The buyer walks past the stand and gets a very vague idea of the quality of the goods being sold, as well as their price. She can then spend some effort in learning more about the selection, look at it from close up, and get a better idea about the batch. Once she has done that, she might decide to buy an orange. She can pick an orange at random, or again spend some effort inspecting an individual unit and deciding whether to buy it or not. Should this unit prove unacceptable, she will return it and pick another, inspect it, and so on. Once she finds an acceptable unit, she buys it and leaves.

4.2 The Type of Diversity under Study

The objects of economic transactions -natural resources, manufactured goods, agricultural produce, labor, etc.- are never strictly homogeneous[14]. When considered from individuals' viewpoints, variability can come in two forms[15] , which have been called "horizontal" and "vertical". By "horizontal variation" we refer to variation in those attributes about which the order of preference across individuals differ, such as color (some people like green shirts and some prefer pink shirts); or taste (some consumers prefer their oranges very sweet, others prefer them tart). In contrast, "vertical variation" is about those for which there is consensus in preferences: for example, purer gold is better gold to everyone, fresher fish is better, a more durable toaster is preferred, a more

the units they will purchase (though not always).

[14]Commodity variability is pervasive. One obvious reason for variability is Darwinian; diversity in living matter enhances survival, so all biological entities are diverse. Minerals in nature are seldom pure, and purifying them is expensive, so impurities are common; moreover, the origin of crude oil, natural gas and coal is organic and thus they are also diverse. The features of the terrain surrounding us are diverse, as is the weather. The value of units in the same apartment house differ for the ones facing north from those facing south; and changes as well depending on whether they occupy a corner and which floor they are on. Restaurant seats are not all equal and not all seats at a movie theater or at the stadium are equally valuable. Dairy products and many others deteriorate as time passes, as do fashion goods. Top law firms seldom invite all the lawyers starting with them to become partners even if their starting salaries were the same.

[15]See for example Waterson (1989) [**?**]. The term "horizontal variation" has been attributed to Hotelling (1929) [**?**], who pointed out that his location model for competing sellers needed not be applied to spatial competition, but could also be used to understand competition by attributes.

fuel efficient car is better, and an orange with less seeds is preferable always. In this paper we explore heterogeneity in the vertical sense.

Once we recognize the lack of uniformity, we see that in many cases it will be too costly for the seller to measure the characteristics of each specimen for sale. Even in very commoditized markets, some variability across specimens remains[16] . If variation is large, buyers will attempt to capture the most desirable units.

A seller may attempt to separate the units into differently priced categories but this is costly. No matter how narrowly defined the category some variation will remain.[17] So a seller must decide, first, how many categories to use, and second, how to deal with the remaining heterogeneity within each category. In this paper we shall focus on the second decision: what to do about heterogeneity within a single category.

4.3 The Transaction under Study

Our seller, who has a batch of non-identical goods to be sold, has three options. Any selection mechanism observed in a real market falls into one of these:

- Choose for the buyer, that is, pick one unit and give it to the buyer who then pays and leaves.
- Make the buyer choose at random, by making it difficult to pick and choose or by offering units that are so uniform that individual inspection is pointless.

[16] Plants, animals and their products may be cultivated or raised to make them more uniform and in addition, they may be sorted into grades; still, it is seldom worthwhile making them strictly uniform. Apples are graded, as are seats in stadiums, but in both cases variability within grade is far from trivial. Manufactured commodities can be made more uniform by using more uniform raw materials and by more rigorous quality control. As these methods seem subject to increasing cost, variability usually remains. Indeed, manufactured products would not require warranties had they been uniform. Moreover, even though new automobiles are guaranteed, some turn to be more "lemony" than others. These observations bring out, among other things, the fact that even in our mechanized world and in spite of the effort to enhance uniformity, the quality of a large number of commodities is far from uniform.

[17] Except the trivial case when each unit is its own category

- Allow the buyer to pick and choose.

If the seller makes the selection, asymmetric information problems ensue: he has incentives to give the buyer a low-quality unit and improve the remaining distribution; but the buyer knows these incentives exist. If she received below-average quality, she would not know if she was cheated or unlucky [18]. Several mechanisms exist that deal with this moral hazard issue: repeated interaction between buyer and seller, a guarantee by the seller, a brand name or a third party guarantee, when possible, can mitigate the asymmetric information problem. But these mechanisms are costly and it might be preferred to let the buyer choose to convince her she is not being cheated.

If the seller manages to make the buyer choose at random, the good is effectively homogeneous *ex ante*. But it might be tricky to induce random choice without opening the possibility for picking and choosing. It is possible, either by increasing uniformity in the production process, which is costly; or by making inspection costly to the buyer[19], but in some cases the only way to make the buyer be sure that she is getting a random (or better) unit, is by allowing her to pick.

It is this third scenario that we explore in this paper. If the buyer is allowed to inspect the units and select one, she will spend effort choosing the best units. With a low inspection cost, the gain from inspecting comes from the asymmetry in outcome brought about by the threshold rule: the buyer gets rewarded for finding an exceptionally good unit, but not punished for finding an exceptionally bad one, because all units below the threshold are equally bad to the buyer. An immediate problem is that the inspection cost, unlike the price paid for the unit, is not transferred

[18]This is a well studied problem . In Akerlof's famous paper, if "lemons" and good cars are indistinguishable by the buyer and the difference in value is large enough, the market could unravel, leaving only lemons; with sellers unwilling to part with their good cars at the price the buyers are willing to pay for an uncertain car. See Akerlof (1970) [**?**] and the literature it engendered.

[19]For example, in many grocery stores, buyers will get nuts and other dried goods in bulk from a clear gravity dispenser which allows them to see the nuts, but they can only open their bag under the dispenser and not choose which units they will take. Should the output be unacceptable, of course, the buyer can discard the whole bag and start again, but we suspect this can only be repeated a few times before the store clerk shows up. The astute reader will point out that this is done also to prevent the buyers from damaging the merchandise during inspection, but this reason does not explain why the container is see-through. In fact, to better preserve dried food, an opaque container is usually better. But then buyers would not know which distribution they are drawing a random sample from, and will be less willing to buy

to the seller, it is dissipated[20]. And there is a second problem: if some buyers pick and choose but others do not, the random-pickers get the worse units, which lowers their expected gain and could potentially drive them, the most desired customers, away. This is the situation we will model below.

We start with an attribute which we will call quality, q. This attribute is such that buyers prefer more to less. The seller has a large batch of specimens that are not identical, in that they have different levels of q. Pricing them individually is out of the question, for it would entail large menu costs and still the buyer would need to inspect. So these oranges are all sold at the same unit price [21], p, which the seller chooses and posts.

We first look at buyers' individual behavior, and then explore the seller's optimal behavior if he faces many buyers who have different inspection costs.

4.4 Individual Buyer Behavior

The buyer passes by the stand and gets very rudimentary information about the merchandise, that is, learns the price p and a preliminary idea of the average quality of the specimens, $\overline{q}$. If this seems acceptable, she spends effort C in scanning the display more carefully and getting a more precise idea of the quality distribution, $f(q)$. Then she can either pick a unit at random, pay, and leave, or pick a unit, inspect it (at a personal cost s) to learn its particular quality level, after which she can either buy it or put it back. If she does put it back, she can pick up another, and either buy it without inspection; or inspect it, then either keep it or put it back, and so on. The flow chart in Figure **??** below ilustrates her possible actions.

To determine our buyer's behavior, we will look at her consumer surplus when she makes each decision. She maximizes her (expected) consumer surplus, which is the difference between the value of the unit she ultimately gets (if any) minus the costs incurred: the price paid to the seller and *the total transaction costs:* the cost of scanning if she decides to do it, and the inspection

[20]Even worse if multiple buyers inspect the same item

[21]Note that if the seller were selling the oranges by the pound rather than by unit, we could consider a pound of oranges to be the unit that is sold at price p, and again, each pound of oranges would be of variable quality. Thus at this point, it is of no consequence to assume that the good is sold by unit.

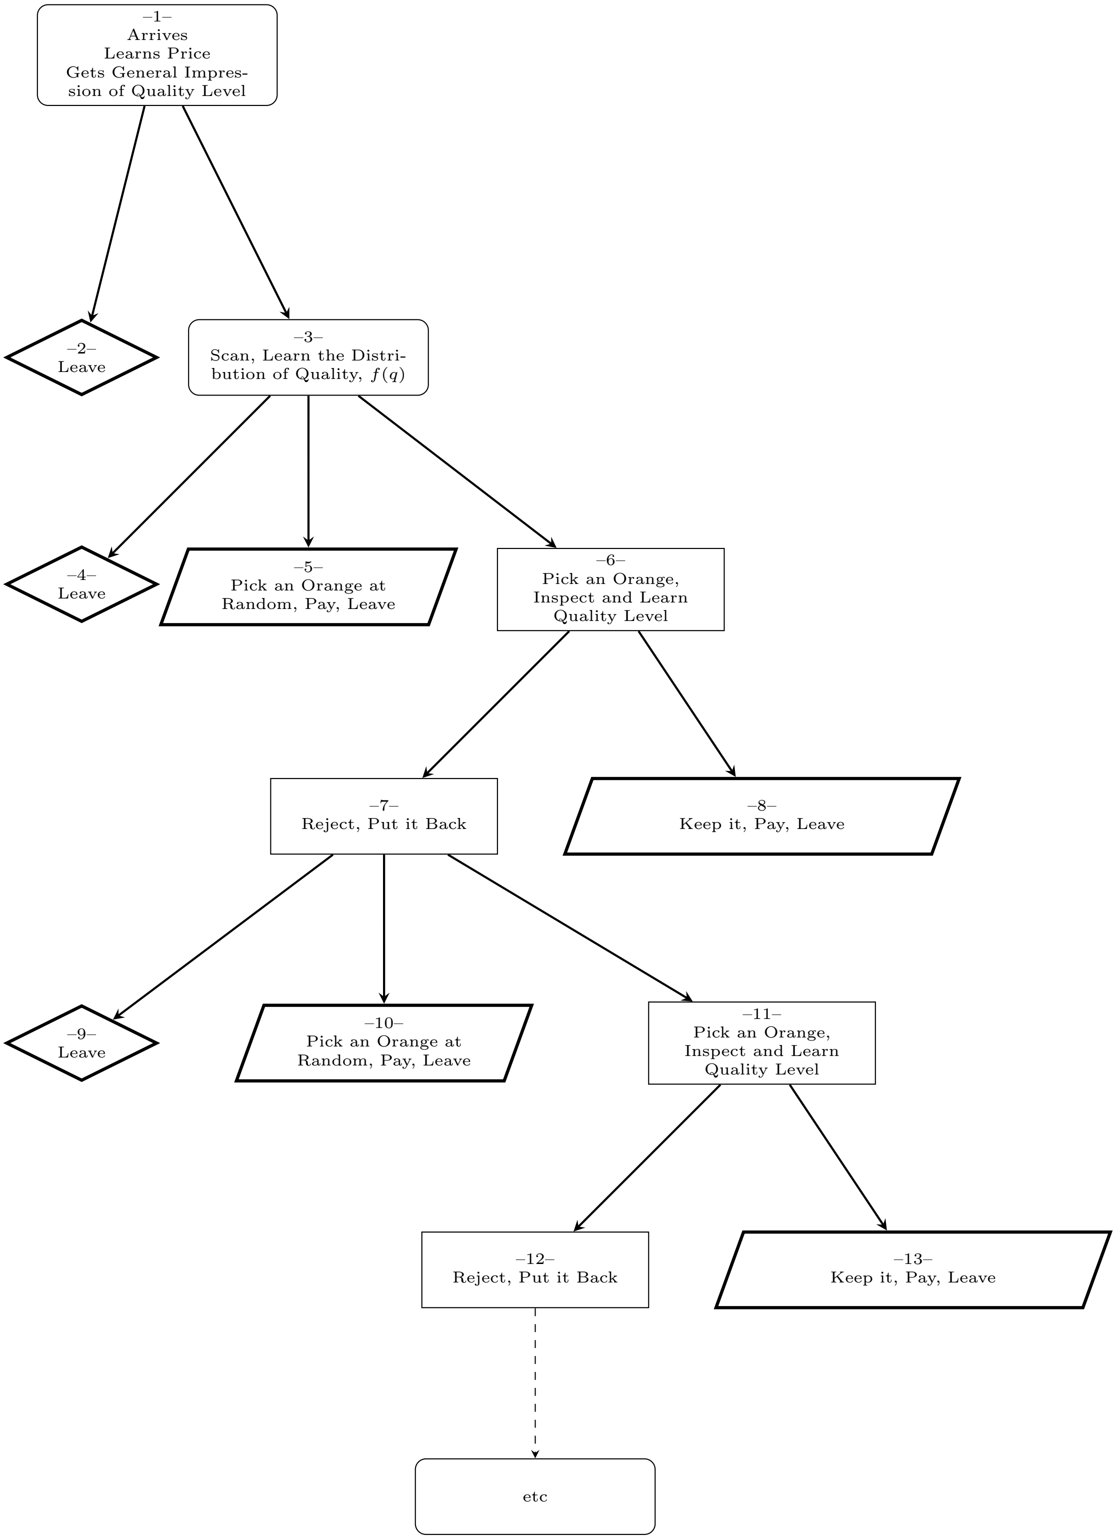

Figure 1: Buyer's Decision Tree

cost for each unit that she inspected.

With this decision tree, the transaction can only end in a "leave" node (nodes $2, 4, 9, 14, 19, ..., 4+5n, ...$) or in a "pick at random, pay, and leave" node (nodes $5, 10, 15, ..., 5n, ...$) or in a "keep an inspected unit, pay and leave" node (nodes $8, 13, 18, ..., 8 + 5n, ...$). In the diagram, the nodes that are final have been drawn with a thicker edge. So her consumer surplus will depend on where in the decision tree she leaves.

- If she leaves without scanning, she has not spent effort or money, and she gains nothing. So her consumer surplus in this case is zero.
- If she scans, but then decides to leave, her consumer surplus is $CS = -C$.
- If she scans, and then decides to pick at random, her consumer surplus is $CS = q - p - C$, where q is the realized quality of the unit, which she only finds out later at the time of consumption.
- If she scans and decides to inspect units, her consumer surplus depends on how many oranges she ends up inspecting: if she were to inspect n units and ultimately purchases a unit of quality q would be $CS = q - p - ns - C$

Suppose she has just scanned the distribution and is in Node 3. She will decide what to do next based on the expected consumer surplus she stands to gain from then on (when cost C is already sunk). If she leaves, she can expect $CS = 0$. If she picks at random, her expected consumer surplus is $E(CS) = E(q) - p$.

If she decides to start inspecting, this is a recursive problem : right after rejecting an inspected orange, her possible actions are the same as right after scanning the distribution. Whatever her expected net gain was when she first decided to start inspecting oranges, it is the same right after discarding an orange (considering that the scanning cost C is sunk by now). This means, among other things, that if she has inspected once, after rejecting an orange she will not leave empty-handed or choose at random, because she would not have started inspecting if those alternatives were more valuable. In fact, some of the nodes in Figure 1 are irrelevant, because no optimal decision-making would ever touch them. This is the case for all nodes of the form $4+5n$ for $n \geq 1$ or or $5n$ for $n \geq 2$; that is, the nodes in which she inspects, rejects, and then does something other than inspect again.

As has been well established[22] , the solution to this type of recursive optimization problem is a threshold A^* such that she will keep the unit if $q > A^*$, and reject it if $q < A^*$. So we will consider what is the threshold that maximizes the expected gain, and compare the maximized expected gain if choosing with the expected gain if leaving or picking at random.

Now, if q is distributed with density $f(q)$ and cumulative distribution $F(q)$, then the expected quality of an *accepted* unit when using a threshold A is given by:

$$E(q|q > A) = \frac{\int_A^\infty tf(t)dt}{1 - F(A)} \tag{1}$$

If she uses an acceptance threshold A, with probability $F(A)$ the next unit will be rejected. So if she decides to start inspecting, she will incur the inspection cost for the first unit with probability 1, but the inspection cost for the second unit only if the first is rejected; the inspection cost for the third only if both the first and the second are rejected, and so on. So, at the time when she decides whether or not to start inspecting, the expected total inspection costs she will incur is:

$$TC(A, s) = s\left(1 + F(A) + [F(A)]^2 + [F(A)]^3 + [F(A)]^4 + ...\right) = \frac{s}{1 - F(A)} \tag{2}$$

And so her expected gain from the whole transaction is

$$E\left(CS(A)\right) = E(q|q > A) - p - \frac{s}{1 - F(A)} \tag{3}$$

Equation **??** shows the expected Consumer Surplus that corresponds to a given threshold A. To optimize, the buyer will choose A^* , *ex ante*, so as to maximize this expression. In the

[22]See for example Bertsekas (1976)

Appendix, we illustrate this threshold and how it affects the buyer's expected Consumer Surplus. This maximization problem has First Order Conditions.

$$(1 - F(A^*))\,[E\,(q \mid q \geq A^*) - A^*] = s \tag{4}$$

This condition is quite intuitive[23] : at the quality threshold, by definition, the buyer is indifferent between keeping the orange in hand (which is of quality A^*) or picking a new one, which will be acceptable with probability $(1 - F(A^*))$ and will yield an expected improvement upon A^* of $E\,[q \mid q \geq A^*] - A^*$.

Plugging Equation **??** in Equation **??** we obtain that the expected consumer surplus of choosing node 6 (inspect), when using the optimal threshold $A*$. is

$$E\,(CS(A^*)) = E(q|q > A^*) - p - \frac{s}{1 - F(A^*)} = A^* - p \tag{5}$$

We can now return to the problem the buyer faces in node 3, right after learning the quality distribution $f(q)$. She must choose between 1)leaving, 2)picking at random, and 3)inspecting, that is, she must choose between 0, $E(q) - p$ and $A^* - p$, with A^* characterized as in **??**. Note that, by now, the cost of scanning to learn f is sunk, so it does not factor in her decision once she is in node 3. It will, as we shall see later, matter for her decision in Node 1, when she either leaves or scans the merchandise.

We now know that she will:

1) leave if both $E(q)$ and A are less than p.

2) pick at random if $E(q)$ is greater than or equal to p and greater than A.

3) and will start inspecting if A is greater than p and greater than $E(q)$.

This yields the first prediction of the model:

[23] See Appendix for derivation

Theorem 1. *The unit price p will determine whether she stays or leaves the stand, but not whether she inspects or picks at random.*

This is is important, because it means that sellers will not be able to induce random choosing by changing price.

Note that the acceptance threshold $A*$ will also be independent of the price, and will depend only on the distribution of q and the unit price s. This means that price influences the decision of whether to buy or not, but not whether to inspect and which threshold to use. This decision depends on two factors: the cost of inspection s and the distribution $f(q)$. That these factors are the determinants of the inspection decision should be clear from Equation **??**.

Theorem 2. *The optimal threshold $A^*(s, f)$, is determined in the following manner:*

1. *The optimal threshold $A^*(s, f)$ is decreasing in s.*
2. *The optimal threshold A^* is higher the more disperse the density function is in its right tail.*

The intuition for 1 that when inspection is cheaper, the prospective cost of continuing the search is lower, so the buyer will be satisfied with a higher level of quality[24] .

For 2, the intuition is that when there is greater dispersion, the buyer could get a large reward for finding an exceptionally good unit, but an exceptionally bad one will be simply returned upon inspection. In fact, this asymmetry is what drives buyers to pay the inspection cost, and thus causes the dissipation.

The next step is to go further up the decision tree, namely, at node 1, when the buyer must decide whether to incur cost C to scan the distribution. This of course depends on what she expects to get as Consumer Surplus if she ends up getting an orange. Given the price and the scanning cost, she needs to compare it with whatever expectations she has about this seller's oranges.

[24] See Appendix for proof and illustrations.

Regarding the information she has to begin with (at Node 1), we assume that she acquires some information at near-zero cost, because presumably she has encountered oranges before, or the seller has displayed some of them very prominently. So she has some beliefs (not necessarily correct) at the beginning of this transaction[25]. In any case, given a certain expectation, she is more likely to stay and start scanning when the scanning cost C is lower.

Now that we have looked at the buyer's behavior, we can briefly look at the seller's behavior if he only faces one type of buyer. Then we will extend the model to many buyer types (that is, different inspection costs) and again look at seller's optimal pricing strategy.

4.5 Seller's Behavior with One Type of Buyer

Let's first suppose that the only variable that the seller can control is the price he posts.

With only one type of buyer, the seller simply needs to choose the highest price possible that does not dissuade the buyer from trading. In the previous section we proved that the price influences the buyer's decision to stay or leave, but not her decision to pick at random or inspect.

Recall that the buyer's Expected Consumer Surplus, when considered from Node 3 (after scanning) is $A^* - p$. So, the maximum price that this seller could charge at that point[26] is A^*.

As more and more buyers pick and chose, however, the remaining distribution of oranges will decay. In fact, no oranges of quality below A^* are sold. Eventually, the seller will need to lower the price. With buyers that pick-and-choose, the price of the oranges will always be above the average quality (value). The resulting propositions are presented in Theorem 3 below.

Theorem 3. *If inspection costs are low enough to induce picking and choosing, we have the following testable implications*[27] *:*

*1 Unit price will be above **average** quality*

[25]In a separate paper (Barzel and Stephany(2017), we model the seller's optimal choice of initial information conveyed to the buyer. In this paper, however, the focus is on the consequences of quality non-uniformity and buyer heterogeneity.

[26]We are sidestepping for now the fact that the buyer needs to stay and scan the merchandise in order for the transaction to occur, so whatever the value of $\overline{q}$ (her expected value for the quality of the orange she would ultimately get), must be at least as large as $p + C$

[27]Formal proofs to all these statements are in the Appendix

2 The distribution will decay over time

3 The price will have to drop.

4 Under this scheme, some units will not be sold.

So, to the extent that changing prices is costly, the seller doesn't want s to be small; he wants it to be large enough to dissuade buyers from picking and choosing. If the seller has any control over s, he will want to make it large enough to bring the threshold A^* below the average quality $E(q)$, because then the buyer will pick at random.

If the seller had control over s, it would suffice to make it large enough to make the buyers pick at random. The problem is that, in many cases, the cost of inspecting individual units is related to the cost of scanning the distribution and making *inspection* costly could entail making *scanning the distribution* costly. Recall that a large scanning cost is likely to drive buyers away. In general, the seller wants s to be large and C to be small, which is why in many occasions it is possible to see the merchandise, but individual inspection - picking and choosing - is made difficult. Orange pyramids are an example of that.

Our prediction that some units will not be sold under this scheme refers to the fact that the seller will eventually have to change the selling scheme (for example, liquidate all the remainder in bulk) or dispose of part of the merchandise (either give it away or dispose of it). A clear example of this phenomenon is studied in Orbach and Einav (2007), who claim that the uniform pricing observed in movie theaters is a consequence of industry structure and legislation, and it is likely that movie theaters would profit by charging different prices for different movies and different prices for different show times or different prices for different seats. In this case, the heteroegeneity is artificially imposed (because the uniform pricing is the result of a combination of market power from the producers and legislation), but several of our results apply as well. For example, we predict that the bottom of the distribution is will not be sold at the uniform price and indeed, movie theater attendance is 3.5 times higher on weekend days than the rest of the week.

4.6 The Market with Many Buyers

So far all of our results are a consequence of the heterogeneity in the merchandise and the possibility of the buyer inspecting in order to capture the best units. In this section, we will consider a market where buyers differ in their inspection costs. This type of heterogeneity is not difficult to find in the real world: people have different opportunity costs for time and different skills, so the effort associated with choosing as well as the value of the time spent inspecting are likely non-uniform.

Let us consider a continuum of buyer types, indexed by $i \in [0, 1]$, and let's assume without loss of generality, that their inspection costs are increasing in i. For each buyer, there will be an optimal acceptance threshold $A^*(s_i)$, where s_i is the inspection cost of buyer i. For convenience, we will denote this threshold A_i^*. By Theorem 2, A_i^* is decreasing in i.

The seller chooses a price below A_0 or else he won't sell any units[28] . However, for any price that he chooses, the Consumer Surplus will be largest for those buyers with the lowest inspection cost, because they will have a larger expected gain from inspecting. In fact, if the price he sets is above A_1, some of the potential buyers will be discouraged from participating.

Note that the total welfare dissipated equals the total effort spent scanning and inspecting. Any transaction cost that buyers spend is lost as profit to the seller. The seller would like to induce ever smaller costs to the buyer so that their willingness to pay increases. Because of the deterioration of the distribution over time, there is an additional cost to the seller: that of changing the price. We do not include it in our model, because we are focusing in the transaction costs associated to the singular transaction, but the inclusion of menu costs as an additional constraint would make the frequency of price changes a decision variable for the seller.

Keeping the above in mind, if the conditions in the market change, it is to be expected that

[28]The buyer indexed $i = 0$ has the lower inspection cost and therefore the higher threshold of acceptance. No buyer will byu if the price is above their threshold.

some non viable markets should become viable when it becomes possible to:

- prevent picking and choosing (while keeping buyer's confidence on the randomness of their own pick).
- when the inspection cost for the low inspection cost people increases.
- when the inspection cost for the high inspection cost people decreases.
- when the uniformity of the goods increases.

4.7 Some Extensions of the Model

4.7.1 Purchase of More Than One Unit

Extend, then divide up C

4.7.2 Heterogeneous Valuations

Extend model so not every buyer has the same utility for the thing

Some non viable markets should become viable when it becomes possible to:

- prevent picking and choosing (while keeping buyer's confidence on the randomness of their own pick)
- when the inspection cost for the low inspection cost people increases
- when the inspection cost for the high inspection cost people decreases

5 Conclusions

We propose the inclusion of transaction costs in a model with non uniform commodities as a source of greater realism in the market. In our model, transaction costs arise when buyers spend resources -time and effort- collecting information about the merchandise in order to obtain the best units when they are all priced the same.

Although heterogeneous goods are considered in models of market segmentation and monopolistic competition, they are treated as a strategic decision and a source of market power. We instead forcus on involuntary nonuniformity, that which results from diversity in the natural world and imprecisions in the production process, is seldom studied.

On the other hand, the optimal pricing strategies for non-homogeneous goods has been explored in the auction literature. But individual pricing or auctioning is not always an option. Menu costs, sorting costs and metering costs can make it inefficient for a seller to price units according to their individual characteristics. Consequently, uniform pricing is sometimes used even when it is acknowledged that not all specimens are identical, and this is the focus of our study.

The setting in our model leaves three broad options for selling variable commodities: The first alternative is to make the selection for the buyer, which is what happens when the buyer does not inspect and the seller picks a unit, which the buyer then takes. The buyer will not accept this scheme if she does not trust the seller to give her a truly random unit. This is how online grocery commerce takes place. The second alternative is for the seller to force random choosing, like when packaging makes inspection very costly, as is the case with cereal boxes which are picked from the shelves with almost no inspection. Again, this requires the buyer to have some faith in the contents of the box. But when there is no trust, sellers might not have any option but to implement the third alternative: buyers will be allowed to pick and choose. We derive several implications for this latter case. price dynamics emerge that are not predicted by traditional economic models. We predict that, if inspection is allowed, the price of a nonuniform commodity cannot remain in place until the full batch is sold. Thus for such items, we observe either disposal of the remainder or a secondary market where units of lower quality are sold.

This potential picking-and-choosing creates transaction costs, because any effort spent selecting is a cost to the buyer which does not translate into a gain for the seller. Our model predicts that in such circumstances the good will be necessarily sold at a price above the average value, with buyers still participating in the market if they can inspect until they find items from the upper end of the distribution. Because the best units are taken and the worst units rejected, the distribution will decay over time, and the seller will be forced to lower the price or dispose of the remainder units.

The costs associated with the consequences described above are greater when more picking-and-choosing is induced. This will happen when inspecting is easier for the buyers or when the distribution is more disperse (more heterogeneity between units). Thus it is to be expected that sellers will spend resources limiting the ability of buyers to inspect individual units and increase their ability to scan whole distributions. Improving the uniformity of the distribution is also desirable to the seller to mitigate the problems described above.

In addition to heterogeneous commodities, we consider the implications of including heterogeneous buyers, that is, buyers with different costs for inspecting units. This could be simply the result of variations in the opportunity cost of time. When this is the case, buyers with lower inspection costs will obtain the very best units, and the seller then must choose decide on a price to capture as much as possible of the value of the merchandise, with the consideration that increases in price can drive buyers away if their inspection costs are too high to pick and choose, but the price is too high to pick at random. This problem would be exacerbated with more and more heterogeneity between buyers. This is a reason for sellers to try to have a customer base that has approximately the same inspection costs.

We believe that allowing explicitly for inspection costs is a suitable choice when modeling non uniform commodities, and there remain many aspects of these transaction to be studied. There are several mechanisms used by sellers to avoid the problems of picking and choosing, such as guarantees, reputation, third party certification, and so on. Some of those mechanisms could probably be included in an extension of our model that permits long term relationships between buyers and sellers.

6 Appendix

6.1 Acceptance Threshold and Consumer Surplus Illustration

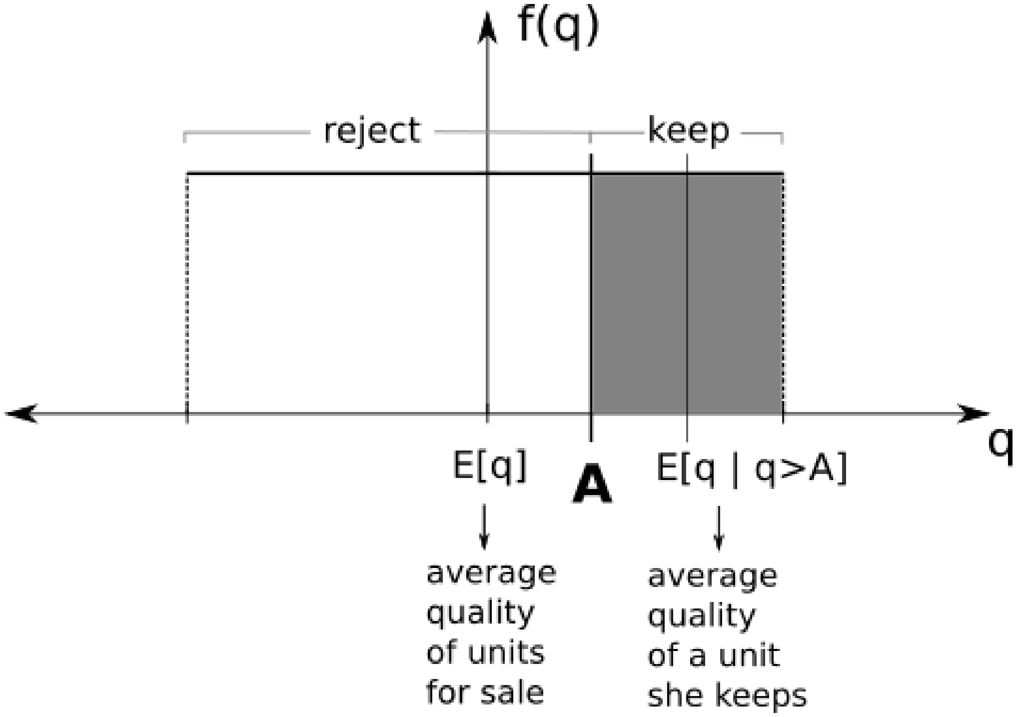

6.2 Derivation of Equation ??

We have that the expected Consumer Surplus for a buyer with inspection costs s and threshold of acceptance A is given by Equation **??**, reproduced below.

$$E(CS) = E(q|q > A) - p - \frac{s}{1 - F(A)} \tag{??}$$

If we substitute the expression for $E(q|q > A)$ given in Equation **??**, we get:

$$E(CS) = \frac{\int_A^\infty tf(t)dt}{1 - F(A)} - p - \frac{s}{1 - F(A)}$$

Using the First Fundamental Theorem of Calculus, we know that:

$$\frac{d}{dA}\int_A^\infty tf(t)dt = -Af(A)$$

So the First Order Conditions for Equation **??** are:

$$\frac{dE(CS)}{dA} = \frac{-Af(A)(1-F(A)) - (-f(A))\int_A^\infty tf(t)dt}{(1-F(A))^2} - s\frac{(-f(A))}{(1-F(A))^2} = 0$$

Which simplifies to:

$$-A(1-F(A)) + \int_A^\infty tf(t)dt - s = 0$$

Which can be rearranged into Equation **??**

6.3 Proof of Theorem ??

To prove that the optimal threshold depends negatively on inspection costs, we can look at the equation that determines A^*, which is derived in the previous section of this Appendix.

$$(1-F(A^*))\left[E\left(q \mid q \geq A^*\right) - A^*\right] = s$$

We can use the Envelope Theorem to analyze the derivative $\frac{dA^*}{ds}$ in the expression above. So, differentiating with respect to s, and using the First Fundamental Theorem of Calculus again:

$$\frac{dA^*}{ds}\left(-A^*f(A^*) - (1-F(A^*)) + A^*f(A^*)\right) = 1$$

which simplifies to:

$$\frac{dA^*}{ds} = -\frac{1}{(1-F(A^*))}$$

Which is always negative because $F(A^*) \leq 1$.

To prove the second part of Theorem 2. it's convenient to look again at Equation **??**:

$$(1-F(A^*))\left[E\left(q \mid q \geq A^*\right) - A^*\right] = s \tag{??}$$

A larger dispersion on the right tail of $f(q)$ means that $\left[E\left(q \mid q \geq A^*\right) - A^*\right]$ is larger. For a given s, the term $(1-F(A^*))$ would need to be smaller, which means that A^* should be larger.

References

[1] Akerlof, G. A. (1970). The market for" lemons": Quality uncertainty and the market mechanism. The quarterly journal of economics, 488-500.

[2] Barzel, Yoram, (1982). Measurement Cost and the Organization of Markets. The Journal of Law & Economics, 25(1), 27-48.

[3] D. Bertsekas, (1976). Dynamic Programming and Stochastic Control

[4] Buzby, J.; Bentley, J.; Padera, B.; Ammon, C. and Campuzano, J.(2015). Estimated Fresh Produce Shrink and Food Loss in U.S. Supermarkets. Agriculture, 5(3), p. 1-23.

[5] Gustavsson, J., Cederberg, C., Sonesson, U., Van Otterdijk, R., Meybeck, A. (2011). Global food losses and food waste (pp. 1-38). Rome: FAO.

[6] Hotelling, Harold (1929). Stability in Competition. Economic Journal. 39 (153), 41–57

[7] Laffont, Jean-Jacques and Maskin, Eric, (1987). Monopoly with asymmetric information about quality: Behavior and regulation. European Economic Review, 31, issue 1-2, 483-489.

[8] Rothschild, Michael, (1973). Models of Market Organization with Imperfect Information: A Survey, Journal of Political Economy, 81(6), 1283-1308.

[9] Salop, S., and Stiglitz, J. (1977). Bargains and ripoffs: A model of monopolistically competitive price dispersion. The Review of Economic Studies, 493-510

[10] Stigler, George, (1961). The Economics of Information. Journal of Political Economy, 69, 213.

[11] Waterson, M., (1989). Models of Product Differentiation. Bulletin of Economic Research, 41, pp. 1-27.