The impact of price frames on consumer decision making: Experimental evidence

Steffen Huck* Brian Wallace†

October 15, 2015

Abstract

We present a laboratory experiment on the impact of price framing on consumer decision making. Consumer subjects face a search market where two sellers offer a homogeneous good. We examine six different price frames with linear per-unit pricing (that is displayed as such) serving as a benchmark. We find that all frames deviating from the benchmark have some negative impact on consumer decision making. The most striking result concerns drip pricing (where prices are decomposed into three elements and "dripped in" during the purchasing process). While leaving the actual decision problem unchanged, drip pricing wipes out 22% of consumer surplus.

JEL classification numbers: C72, C92, D21, D43

Keywords:

*WZB & UCL

†UCL.

1 Introduction

There has been much attention on behavioral biases in consumer decision making and their possible consequences for competition in the recent literature. This has created a new field, behavioral industrial organization, recently surveyed in a first graduate textbook by Spiegler (2011). There is, however, only little empirical evidence on the actual nature of consumer biases. This paper presents a laboratory experiment on the impact of price framing on consumer decision making (studied in theoretical models by Piccione and Spiegler and Chioveanu and Zhou). Consumer subjects face a search market where two computerized sellers offer a homogenous good. We examine six different price frames with linear per-unit pricing (that is displayed as such) serving as a benchmark.

Two of our other frames leave actual prices untouched: *reference pricing* (which essentially examines the "rhetoric of sales" where subjects see an irrelevant higher previous price) and *drip pricing* (where prices are decomposed into three elements and "dripped in" during the purchasing process). Three further frames change actual pricing and impact on the optimal search rule for which we control: *time-limited offers* (where sellers do not guarantee the same price if the customer wants to search elsewhere before buying), *complex (non-linear) pricing* (in form of a "3 for 2" offer), and *baiting* (our only treatment with price advertising which, however, does not imply commitment as the advertised price is subject to availability).

We find that all frames deviating from the benchmark have some negative impact on consumer decision making. The most striking result concerns drip pricing . While leaving the actual decision problem unchanged, drip pricing wipes out 22% of consumer surplus. The result is surprising as drip pricing occurs in a rather benign form. Compared to the baseline subjects only need two additional mouse clicks to see the final relevant price. Our experiment has informed the UK Office of Fair Trading (OFT) which has subsequently recommended that government should take action against drip pricing. In a first response, the UK coalition government has announced that it will ban "excessive" credit card surcharges.

All other frames also impact negatively on the quality of consumer decision making. Our paper provides the first experimental evidence on the impact of such frames in a unified setting. While prices are assumed to take on a crucial role for economic behaviour, their presentation, or framing, has historically received little focus in economic research because how prices are presented has been deemed irrelevant. Yet, sellers use many different ways of presenting or framing prices in the marketplace and they often change these frames. As changing the presentation of prices is costly, there appeared to be a "price framing paradox" where standard economics suggested on the one hand, that price framing is irrelevant for consumer behaviour but must, on the other hand, have some impact on demand because otherwise firms would not spend money on manipulating them.

Our experimental environment retains all the crucial features of real consumer choice problems: goods are on offer at multiple shops, consumers might want to buy single or multiple units and they can search among different shops. Actual prices are drawn from a uniform distribution and (with the exception of the baiting treatment with price advertising) the two sellers are ex ante identical. We derive the optimal search rules and find that subjects in the baseline treatment come very close to optimality. The main source of error in the baseline is to be found in a pattern of *over-search*, that is, subjects who should buy at the first shop they visit have a tendency to continue their search too often — a somewhat predictable deviation in an environment like this where "demand effects" will tempt subjects to interact with their environment as fully as possible. Under drip pricing, over-search does not only disappear but is actually reversed in a pattern of *under-search*. We suggest that this is due to a shifting reference point where abandoning the purchasing process due to additional drip would be perceived as a loss which increases consumers' willingness to pay.

In Section 2 we discuss some of the existing empirical evidence on price framing and offer a brief review of the theory literature that assumes that price frames impact on consumer demand. In Section 3 we describe the experimental design in more detail and derive the optimal seach rules. Section 4 contains the results and we conclude in Section 5.

2 Price framing in the literature

2.1 Empirical findings

The literature on price framing and its impact on consumer behaviour is surprisingly patchy. Even within the marketing literature the effects of different price frames have not been systematically explored and many studies focus on hypothetical choice or simply rely on perception or recall data. Here we will focus on choice-based studies.

One price frame which has received some attention is *partitioned pricing* where prices are decomposed into several elements such as base price and a shipping or handling charge (but where in contrast to drip pricing all these elements are shown simultaneously). Morwitz, Greenleaf and Johnson (1998) show how partitioned pricing lowers consumers' recalled price for the good and increases demand. For example, in a fully incentivised auction experiment in which participants bid for a jar for pennies, one group of participants/buyers was given a bid form which told them that they must pay 15% in addition to their bid if they win the auction. A second group just had to pay winning bids. *Morwitz et al.* found that the partitioned pricing pushed up subjects' willingness to pay.

Hossain and Morgan (2006) present a field experiment on partitioned pricing. Buyers in an eBay auction made bids for CDs and Xbox games. In the auctions there was a reserve price and separate (fixed) shipping and handling costs. Hossain and Morgan observed that when the reserve price is low as compared to the retail price of the good, and the shipping and handling costs are high, then the auction always results in a higher sale price than in a situation in which the reserve price is high (relative to retail) and the shipping and handling is low. In a follow-up paper Brown, Hossain and Morgan (mimeo 2007) sold different iPod models on Yahoo Taiwan and eBay Ireland. They find that shrouding low shipping charges is a money-losing strategy for the seller but that raising shrouded shipping charges does increase revenue. In this instance shipping charges were either shown on the search page for each auction or they were shrouded because buyers had to read each individual auction page to learn the shipping costs.

Wansink, Kent, and Hoch (1998) report results from a field experiment to examine *multi-unit price frames* which they compare to *single-unit offer frames.* The multiple-unit frames typically used the format "buy two for x" where one unit would have been available for $x/2$, that is, both frames relied on simple linear pricing. Utilizing a data set from 86 stores that are randomly assigned to one of the two offer frames they find that, for a wide range of products, sales under the multiple-unit frame were, on average, 32% higher than under the single-unit frame.

Finally, Ellison and Ellison (2005) analyse a market characterized by the presence of price search engines. Price search engines are designed to make consumer search and price comparison easier, and as such, reduce friction in the market such that the price of identical goods should be identical. However, retailers may seek to put friction back into the market by making price search more difficult and thereby less of a threat to profitability. Ellison and Ellison, using field data from an actual price comparison site in the US observed a highly successful version of *baiting* where sellers advertise a low-quality product at a very low price as a bait and then try to convince consumers to switch to better quality more expensive products once they are in their on-line store (despite the cheap product being available).

2.2 Theory

The rather sluggish accumulation of empirical evidence on consumer biases stands in sharp contrast with the the rapid growth in theoretical models that examine how competition unfolds in the presence of biased consumers. Behavioural IO has been one of the fastest growing fields in the recent (applied) theory literature. Huck and Zhou (2011) present a survey and Spiegler (2011) offers a textbook treatment. Here we focus on models where consumers are explicitly assumed to react to price frames.

Gabaix and Laibson (2006) analyse markets where sellers can shroud the prices of expensive add-ons. Such examples include the sale of printers and the add-ons for ink cartridges, or consumer credit products with additional costs for late minimum payments. Biased consumers ignore shrouded add-on prices when making choice about the main product, rational consumers understand that shrouded prices are likely to be high. This creates incentives for sellers to sell the main product below marginal costs and impose high monopolistic margins for add-ons. Both biased and sophisticated consumers will prefer to buy from the shrouding firm rather than from a firm that prices according to marginal costs — biased consumers because they overlook the add-ons and sophisticated consumers because they can choose substitues for the add-ons.

Chioveanu and Zhou (2009) study oligopoly markets where identical sellers of an identical (homogenous) product choose both, prices and price frames. The authors assume that price frames have two effects on consumer behaviour. Consumers may fail to compare prices correctly because of the complexity of a price frame and/or because of the difficulty comparing different frames. Remarkably, these behavioural biases can completely overturn standard intuition about the functioning of markets. Indeed it is shown that an increase in the number of firms can increase industry profits and harm consumers (while standard theory would assume that new entrants reduce industry profits, thereby reducing prices and benefiting consumers). This arises because the framing acts as a form of price differentiation meaning consumers cannot compare prices for identical goods. In related work, Piccione and Spiegler (2009) analyse duopoly markets where consumers make price comparisons only with a certain probability that is assumed to depend on the price frames chosen by firms. Broadly in line with the results by Chioveanu and Zhou, they argue that product differentiation can also be viewed as a means to reduce consumers' ability of comparing prices effectively.

Finally, Armstrong and Zhou (2010) study how time-limited offers might arise endogenously. If firms can track consumers they have an incentive to offer buy-now discounts which, in equilibrium, helps to raise market prices.

3 The model & experiment

In order to analyse all of the five price frames (plus a baseline with straight unit pricing) using the same design, we need, as a minimum, an environment with multiple shops in which multiple units of at least one good can be purchased. Moreover, we need scope for an advertising stage such that the experimental environment needs to mirror not only the shops but also the consumer's home where he might be reached by some advertising before actually going to a shop.

We have opted to include all these facets in the simplest manner in order to minimize the noise in decision making. The less noise that originates from the complexity of the basic environment, the sharper will be the results of the price-frame comparison.

3.1 Baseline environment model

The consumer environment is designed in the following way. There are two shops, both of which sell the good that the consumer wishes to buy. At the start, the consumer is at “home”. The consumer can go back and forth between his home and the two shops as often as he likes and buy units of the good at the shops (and incurs search/travel costs doing so). He does not know the price of the good at either shop until he visits it (the price of the good does not change between visits). Consumers get utility from buying units of the good and incur costs of buying and travelling to the shops (i.e. search costs).

The payoff function of the consumer is thus:

$$\pi(n_1, n_2, t) = u(n_1 + n_2) - n_1 p_1 - n_2 p_2 - ct$$

where n_i is the number of units bought at shop i, p_i is the price at Shop i, $u(n)$ is the utility of buying n units in total, t is the total number of visits to shops and c is the search cost per visit.

Goods at both shops are of the same quality. This allows us to focus on the pure effect of the frames. In many real-life markets, prices (current or former) may serve as a quality signal which renders the decision problems much more complicated and confounds the issue of price framing.[1]

In the baseline treatment, each shop draws a unit price from the uniform distribution $U[\frac{1}{2}, 1]$ (the price interval is known to the consumer).[2]

The consumer has a utility function with decreasing marginal utility; marginal utilities of the first four units purchased are 1, $\frac{2}{3}$, $\frac{1}{6}$, $\frac{1}{12}$ respectively and the marginal utility of further units is zero. Notice that with the straight unit pricing implemented here, the consumer will never buy more than two units, as the marginal utility of the third unit is smaller than the lowest possible price[3].

[1]Some of the literature on reference pricing, for example, Simonson (1999), suggests that the reference price can encourage consumers to switch from low-quality to high-quality options (Chapter 3).

[2]Exogenous prices help us to compare treatments, but clearly, in an environment where firms can adjust prices optimally the effects of any suboptimal consumer behaviour would presumably be more pronounced. Hence, with exogenous prices we will underestimate the consumer detriment (if any) because firms do not respond to consumer behaviour.

[3]However, these values enable us to study the “3 for 2” price frame in a meaningful way: there may

In this baseline treatment, consumers see the per-unit price a shop charges once they visit this shop. This price stays constant. That is, if the consumer returns to the shop he will still get the same price. This holds for both shops.

3.2 The implementation of the price frames

The experiment studies five different price frames and compares these to the baseline of straight, per-unit prices described previously.

For the implementation of the five different price frames, we have opted for "typical" implementations, informed by field choices. For example, we have opted for two price drips which we found to be common.[4] Of course, our comparisons between different price frames are a function of our design choices. For example, complex pricing could be much more complex than "3 for 2" (we could imagine highly non-linear tariffs) to an extent where consumers completely fail to understand the marginal prices of extra units. Our design choice has been to opt for simple generic versions of all price frames in order to maximize possibilities for comparability. We design the experiment so that both shops always employ the same type of price framing.

The setup for the *reference pricing* frame is almost identical to the baseline. The selling prices are determined in precisely the same way as for the baseline and the only difference to the baseline is that consumers are additionally informed about a "former price". Specifically, they are told that the former price is a price that is chosen randomly between the actual selling price and the maximum possible price. They are also told about the consequent "discount" that this represents in percentage terms. As such, it should be easy for the subjects to see that the sales information is actually entirely meaningless. Even though subjects are not told how the former price is generated, it is apparent that the good has not been sold previously at this price. Of course, in an experimental setting with many rounds, it might have been to some other subjects in some earlier round but for the subject faced with the decision at hand this is irrelevant information and should easily be identified as such.

Drip pricing is also virtually identical to the baseline. Again, actual selling prices are

be cases where it is optimal to only buy one unit. They also enable consumers to see cases where the marginal utility is positive but less than the price.

[4]In more complicated cases, we did find instances where there were up to 4 compulsory price drips.

determined in precisely the same manner. This time, however, consumers learn about the selling price only in drips. Once they visit a shop, they see a base price (with no mention of additional charges). Once they decide to buy one or more units, they see a first drip and need to click ok to proceed. If they do so, they are informed about a second drip and need to ok this as well. Finally, they see the total price (and its decomposition into the three components) and they need to click one more time to confirm their purchase. The only difference to the baseline is that subjects need to click twice more to learn the actual selling price; actual selling prices are determined in precisely the same way as in the baseline (being drawn from $U[\frac{1}{2}, 1]$). The actual selling price is decomposed into the base, the first drip and the second drip. The first drip is randomly chosen to be between 5% and 15% of the selling price, and the second drip is randomly chosen to be between 10% and 20% of the selling price. The base price is the remainder. In the experiment, the drips were labelled "shipping" and "handling"; again, these clearly have no meaning in an experimental context, so should have been identified as irrelevant.

Under *complex pricing (3 for 2)*, the unit prices are again determined in the same way as in the baseline but consumers are not charged for a third unit if they wish to buy three. They are informed of this offer once they enter the shop and the offer remains in place for all visits to both shops[5].

In *time-limited offers*, consumers are told when they enter the first shop that the price they are confronted with now is only on offer at this visit. If the consumer comes back at a later stage, the price will have changed. Specifically, the shop simply draws a new price from the same distribution; that is, sometimes consumers will actually get a price upon return that is below the time-limited offer price. The other shop engages in the same tactics (that is if the consumer had visited the other shop first, he would also have been confronted with a time-limited offer). However, once the time limit passes both shops offer a price that then remains fixed. This means that the consumer never experiences a time-limited offer at the second shop (simply because the offer there was also just valid for the first shop visit).

Finally, under *baiting*, both stores advertise prices. That is, in contrast to all other treatments, consumers have some price information at their "home screen". These

[5]The complex pricing treatment is one reason for the choice of utilities of the units - we needed to choose them such that, in some cases, there was the possibility that buying only 1 unit was optimal, but that also offered the opportunity for subjects to overbuy.

prices are under a generic "while stocks last" qualification that is printed next to the price information that subjects can see at the home screen. Specifically, if the selling price is in the interval $[0.5, 0.6]$ (recall the price interval is $[0.5, 1]$), the advertised price is equal to the selling price. If the selling price is greater than 0.6, then with probability 0.5 the advertised price will be the selling price and with probability 0.5 the advertised price will be some price randomly drawn from the interval $[0.5, 0.6]$ (i.e. a "bait"). When a consumer visits their first shop, they will see whether the price is real or a bait. The selling price at the other shop is now randomly drawn from the interval $[0.5, 1]$ (as if any baits had already been sold). If the consumer returns to the home screen, the advertising has gone.

3.3 Optimal strategy in the baseline

The first decision is to choose one of the two shops to visit first. As there is no history and no information about the two shops, this is inevitably a random (and, hence, meaningless) decision. (Of course, this changes in the presence of some advertising, as in the baiting treatment.) So, without loss of generality, we can call the shop the consumer chooses first "Shop 1" and the other shop "Shop 2".

Once the consumer is at Shop 1, he can see the unit price that the shop charges. He can then either buy as many units as he desires (up to a maximum of 4) or he can also decide not to buy and return to the home screen empty-handed. He can then travel to Shop 2 (or indeed go back to Shop 1 but that would, of course, not be reasonable as no new information would be revealed), again paying the search costs, c. At Shop 2 the same rules apply. That is, the consumer learns the unit price charged at Shop 2 and can buy as many units as he desires. Again, he can also return empty-handed and, if he desires, return to Shop 1. He can, of course, also return to shop even if he has bought some units already, either at Shop 2 or even at Shop 1, but that would also not be rational.

We can analytically derive the optimal consumer search strategy for a fully rational risk neutral consumer in the continuous price space. Here, we summarize the solution.[6]

Notice first that the marginal utilities and prices used in this experiment mean that it is only ever optimal to buy one or two units in total. The marginal utility of the third

[6]The theory is available in an Appendix. In the experiment, prices were discretized.

unit is always less than the price.

At Shop 1, there is a cut-off price, p^*, which depends on c. p^* is the cut-off buying at Shop 1 or searching more by visiting Shop 2.

Let p_1 denote the price at Shop 1. If $p_1 < p^*$, the consumer will do all their shopping at Shop 1. If $p_1 < \frac{2}{3}$, the consumer will buy two units at Shop 1 and otherwise only one.

If $p_1 > p^*$, then the consumer will not buy at Shop 1, but rather go to Shop 2 and see what the price is there. The consumer may return to Shop 1 later, depending on the price at Shop 2 and the prevailing search costs[7].

Thus, there are four different possibilities that can then arise after going to Shop 2: (i) He can buy two units at Shop 2 and return home. (ii) He can buy one unit at Shop 2 and return home. (iii) He can decide not buy at Shop 2 and return to shop 1 to buy two units there. (iv) He can decide not to buy at Shop 2, return to Shop 1 and buy one unit there.

Which of these four options is optimal depends on the search costs and the two prices p_1 and p_2. The higher the search costs, the higher the cut-off p^*.

3.4 Optimal strategy in the price frames

Two of the five price frames leave the optimal strategy from the baseline completely untouched because the true unit price does not change and is clearly discernible to the unbiased consumer. In the *reference pricing* frame, the consumer directly sees the true unit price and just receives some completely irrelevant additional information (the former price)[8] and under *drip pricing* the true price gets clearly revealed, albeit only in drips. So in both of these frames, a fully rational consumer will simply adopt the same strategy as that we have derived as optimal above.

Optimal search is, however, slightly different under the other three frames. Under *complex pricing* ("3 for 2"), however, the derivation of the optimal strategy strategy

[7] It is never optimal to buy a unit at Shop 1 and then continue to Shop 2

[8] In practice, former prices might often serve as a signal of quality. By abstracting from such quality issues we can measure the pure effect of framing a price as a sales price.

is essentially identical. We can revisit the baseline analysis and simply change the marginal utility of the second unit. Specifically, we replace it by the sum of the marginal utility of the second and third unit ($\frac{2}{3} + \frac{1}{6} = \frac{5}{6}$). The logic is simple: as the third unit provides some positive utility you will always take it if you do buy two units. In other words, you will never refuse the offer and buy just two units. However, you might still sometimes prefer just to buy one unit (i.e. if $p > \frac{5}{6}$). This changes the cut-off p^*. Table 1 shows the relevant cut-offs for complex pricing for each search cost level.

For the remaining two frames, the structure of the analysis also changes slightly. First, let us consider *time-limited offers*. If a time-limited offer is not taken and the consumer returns home, both shops draw new prices. Consequently, if a consumer does not take the time-limited offer and returns to his home, he faces a situation that is identical to the original straight pricing baseline because, from now on, the shops do simply stick to one price and refrain from any further such offers. This implies that, when the consumer sees the time-limited offer at Shop 1, he has the choice between buying now or playing the original game[9]. This means the consumer can compute the utility he would receive from buying optimally one or two units now (for the time-limited offer price) and compare this with the ex ante expected utility he receives from playing the baseline game, assuming, of course, he plays this optimally. This generates the different cut-offs shown in Table 1.

Finally, let us consider *baiting*. Under baiting, the consumer needs to take into account that low prices in the range from $[0.5, 0.6]$ might turn out to be baits, while with higher prices, he can rest assured that they will turn out to be true. This solution necessitates knowing the precise rule the firms employ which, in the experiment, is initially not the case but, through repetition, may be able to be learnt approximately. As the baiting offer is only available initially, the rational consumer essentially has to take just one additional decision that goes beyond the optimal baseline search strategy. At the home screen, he has to decide which shop to visit first. Notice that this is indeed the only frame where the initial choice of shop is meaningful. Once he is at the first shop and the true price gets revealed, the consumer is back to the original problem of the baseline. It is, of course, this elegance of the rational benchmark solution that has inspired the specific design here; that is, the decision to have the baits only initially and revert to a new price draw after that (for both firms).

The optimal choice of shops is not simply determined by the lowest advertised price.

[9] It is never optimal to buy one unit and return home

Search cost	Baseline	Complex	Time-limited offer
c = 1/60	0.621	0.621	0.660
c = 1/20	0.719	0.699	0.715
c = 1/10	0.815	0.766	0.790

Table 1: p^* cut-offs

As low prices may be baits, but high prices tend to be honest, the consumer needs to work out average expected payoffs based on the offered prices. For prices that look like potential baits, but that are not extremely low, the consumer might actually be better off to seek out a shop that advertises a higher but honest price.

3.5 The role of search costs

As described in the previous section, search costs will affect the optimal strategies for subjects by changing the search cutoffs. We chose to vary search costs; this both enables us to assess the effects of search costs on behaviour (and any interaction between them and the price frames) and focus subject attention in the experiment, viewing each repetition as truly different.

Specifically, search costs varied with c chosen at random from $\{\frac{1}{60}, \frac{1}{20}, \frac{1}{10}\}$. Table 1 shows the cut-offs for the search costs we have implemented.

3.6 Experimental procedures

A total of 166 subjects participated in the experiment; all subjects were students and they were recruited using the ORSEE system. Each subject was confronted with the baseline and two of the five frames. Subjects played for thirty rounds, ten for each type of price frame. The sequence in which they faced the different frames was partly randomized[10]. This repetition enables subjects to learn about the environment and adapt their behaviour. This enables us to determine whether any potential effects of the price frames can be overcome through experience.

In order to enhance attention (and make sure that each round was viewed as a truly

[10]Specifically, in each tenth of the rounds, subjects encountered each price frame once

Product	c (search costs)	0 units	1 unit	2 units	3 units	4 units	Price Range
GREEN	1,3,6	0	60	100	110	115	30 to 60
ORANGE	1,3,6	0	80	140	170	195	50 to 80
BLUE	2,6,12	0	110	180	190	190	50 to 110
RED	2,6,12	0	120	200	220	230	60 to 120

Table 2: Parameters for different goods

new round), we scaled payoffs in four different ways. Specifically, subjects faced four different goods, labelled GREEN, ORANGE, BLUE and RED. Utilities and prices for each good were obtained from the model above through different ways of scaling up. This ensures that the decision problem structure is always identical, regardless of the specific goods subjects could buy. Note that two goods, BLUE and RED, are scaled up by a factor of 2 from the other goods (so incentives in the rounds where these goods were on offer was essentially doubled).

The actual payoff, price and search costs are given in Table 2. The model was discretized so that prices and costs were always given in integers[11]. Subjects were provided with written instructions which were read aloud by the experimenter.

There are 10 combinations of two out of five price frames and we implemented all of them. That is, we studied ten different groups of subjects where each group of subjects is characterized by a combination of two price frames subjects are faced with in addition to the baseline that every subject experiences. This allows us both within and between subject comparisons.

In addition to the 30 rounds of the experiments, each subject undertook:

1. Pre-experiment test to ensure that they understood the experimental instructions

2. An incentivized 12 question IQ test (using variants of Raven's matrices)

3. 15 question personality test[12]

[11]This changes the search cutoffs slightly from the continuous model; we calculated the cutoffs for the discretized model numerically. At most, the cutoff moved by one point.

[12]We found no correlation between personality and behaviour in the experiment.

4. Feedback questionnaire about the experiment

Experimental sessions lasted on average 145 minutes and average earnings were approximately £20 which included a £5 show-up fee.

4 Experimental results

In our analysis, we proceed as follows. We first study whether price frames have any effect on consumer welfare. To make treatments comparable we define losses in consumer welfare relative to what consumers could have achieved under optimal behaviour.[13] In a second step, we analyse errors in behaviour, distinguishing errors in search and errors in purchasing. Finally, we will zoom in more closely on search patterns and purchasing behaviour, in order to understand the root causes of poor performance and how they relate to known behavioural phenomena.

4.1 Consumer welfare

We now turn to the first fundamental question this research addresses: do price frames matter for consumer welfare? Despite our extremely simple environment (and using a rather smart subject pool), we find that price frames do matter.

Since the amount of consumer welfare obtainable under optimal behaviour differs between the different frames and depends on the prices drawn, we need to normalize achieved payoffs in an appropriate manner. In order to do this, we take, for each of the 4895 observations[14] we have, the difference between the actual achieved payoff and the payoff that would have resulted from following the optimal decision rule. We call this variable the consumer's loss. If a consumer could have achieved a payoff of 0.87 under optimal behaviour but only achieved a payoff of 0.69 then his loss is $0.87 - 0.69 = 0.18$. We look at the mean loss a consumer has made in a particular price frame which is simply calculated as the arithmetic mean of all the losses in all rounds in which they encountered that frame.

[13] The optimal strategy is the one a subject that knows the experimental environment and is fully rational would employ.

[14] the number of observations from 166 subjects should be 30 * 166 = 4980; however, in one session of 17 subjects, the software terminated after 25 periods.

Additionally, we compute a further welfare indicator, the extra loss relative to the baseline loss. This is computed as follows. For each subject we have three average loss variables, one for the baseline, and two for the two price frames encountered.[15] The extra loss a subject incurred under a price frame is then simply defined as the difference between the mean loss in this price frame and the mean loss in the baseline. We can then also compute the extra loss made on average by all subjects under a particular price frame. This difference-in-difference approach controls for both different earning potentials under different frames and subject-specific differences in performance levels.

Table 5.1 shows both loss and extra loss for all price frames and the baseline. The magnitude of losses as shown in Table 5.1 refers to the normalized model that we sketched above. For the price frames that are equivalent to the baseline, the mean payoff under optimal decision making us .29, so the baseline welfare loss translates to about 13% of possible earnings and the worst of the frames 22%, a substantial fraction of subjects' earnings (and 68% more than in the baseline).

Table 5.1: Welfare losses under the different price frames

Treatment	Frame	Welfare Loss	Extra Loss
1	Baseline	0.038	0.000
2	Complex pricing	0.052	0.005
3	Drip pricing	0.064	0.026***
4	Baiting	0.048	0.016**
5	Reference pricing	0.055	0.019
6	Time-limited offers	0.064	0.018**

Note: Stars indicate significant differences to baseline: * 10%, ** 5%, *** 1%

For each price frame, we test whether the average performance is significantly different from the baseline performance (using a wilcoxon test on the welfare loss by subject). We see that there is a significant difference for three of the prices frames (the remaining two are significant at just outside 10%).

Two main results emerge from Table 5.1.

Result 1 Price framing is indeed detrimental for consumer welfare for all frames, but

[15]Recall that each subject did three treatments including the baseline within a session.

the amount of detriment and significance varies between frames.

Result 2 Drip pricing emerges as the worst culprit with the biggest average loss and the biggest average extra loss.

4.2 Errors in consumer behaviour

There are generally two types of errors subjects can make in the experiment: errors in their search activity and errors in purchasing behaviour. Errors in search activity occur where a consumer makes more or fewer visits to the shops than is optimal. Specifically, there are two types of search errors. A consumer makes a search error if he buys at the present shop but should optimally have continued his search. Or, vice versa, if he continues his search but should have optimally bought at the present shop. Errors in purchasing behaviour occur where subjects do not buy the optimal amount of the good. For example, they buy one unit when it would be optimal to buy two units given prices and marginal utilities of consumption. Notice that search and purchasing errors can also occur together. For example, when a consumer buys one unit at the "Shop 1" while according to the optimal strategy he should have bought two units at the second. While determining the optimal number of units is a rather simple task (it just requires a very basic understanding of the pay-off table), the decision about search is more demanding.

There are some arguments for why in the context of this study it might be more important to focus on errors rather than overall performance as the performance measures above are sensitive to the precise parameters chosen in the experiment. For example, the losses would have been much bigger if the value of the goods had been higher.

Table 5.2: Error rates between treatments

Treatment	Frame	Error rate	Search error rate	Purchase error rate
1	Baseline	0.282	0.216	0.089
2	Complex pricing	0.306	0.244	0.094
3	Drip pricing	0.390	0.313	0.141
4	Baiting	0.362	0.288	0.117
5	Reference pricing	0.326	0.262	0.089
6	Time-limited offers	0.467	0.410	0.086

In order to get a first grip on errors in decision-making concerning search activity and purchasing behaviour, we examine each of the 4895 rounds played by our 166 subjects. Whenever the observed search behaviour in a given round departs from the optimal number of visits, we classify the round as a round with a search error. Analogously, whenever the observed purchasing behaviour in a given round departs from the optimal number of units in one of the shops (regardless of whether the subject has searched optimally), we classify the round as a round with a purchasing error[16].

Firstly, we tabulate the error rates in Table 5.2. We can immediately see that error rates are much higher in several of the treatments than the baseline (broadly corresponding with the welfare findings in the previous section).

We then regress both types of errors on prices, search costs, the scaling factor, together with variables that capture learning and IQ (we discuss the role of these in a future section) and the different price frames (where each frame is captured by a binary dummy with the baseline serving as the reference)[17][18] . We use probit regressions and cluster the standard errors on the subject level in order to account for dependencies resulting from repeated measurement. Tables 5.3 shows the estimated marginal effects.

All detailed regression results are contained in the technical appendix where we also include estimations from linear probability models that throughout confirm the robustness of our results.

[16] There are a small number of cases where subjects buy from both Shops or don't buy at all. The inclusion or exclusion doesn't significantly affect the results

[17] The treatment effects are the same even if the decision specific parameters (i.e. p_1, p_2, c, multiplier) are omitted

[18] If we include a dummy for having made a search error decision in the purchase decision, it does come out significant but doesn't affect the other coefficients

The way to read Table 5.3 is the following: the estimated coefficients show if subjects in the experiment either searched 'too much' or 'too little' than was *optimal* in the relevant price frame as compared to the search errors that subjects made in the baseline treatment. The second column does likewise for purchasing errors. The third column captures any error (regardless of whether a purchasing or search error).

Table 5.3: Probit estimation of errors in search and purchasing behaviour behaviour

	search error	**purchase error**	**any error**
variable	coefficient (standard error)	coefficient (standard error)	coefficient (standard error)
p1 (price at 1^{st} shop visited)	0.201 *** (0.065)	0.096*** (0.031)	0.279*** (0.065)
p2 (price at 2^{nd} shop visited)	0.209*** (0.039)	0.037 (0.027)	0.224*** (0.045)
c (search costs)	-0.419** (0.192)	0.123 (0.140)	-0.267 (0.209)
Highvalue good (Blue or Red goods)	0.007 (0.013)	-0.001 (0.009)	0.013 (0.014)
d2 (Complex pricing)	0.033 (0.025)	0.010 (0.016)	0.030 (0.024)
d3 (Drip pricing)	0.106*** (0.027)	0.052*** (0.019)	0.116*** (0.026)
d4 (Baiting)	0.087*** (0.025)	0.037** (0.017)	0.099*** (0.026)
d5 (Reference pricing)	0.054** (0.023)	-0.001 (0.015)	0.050** (0.022)
d6 (Time-limited offers)	0.202*** (0.026)	-0.001 (0.015)	0.192*** (0.024)
Period	-0.002** (0.001)	-0.004*** (0.001)	-0.004*** (0.001)
IQ (1 to 12)	-0.004 (0.004)	-0.012*** (0.002)	-0.012*** (0.005)

Looking at the third column, with the exception of complex pricing, we observe significantly more erroneous behaviour (relative to the optimal behaviour benchmark) in all price frames. The impact of the price frames regarding the order of the strength of effects is not the same for both types of errors: the treatment effects are stronger on search than purchasing errors with almost all frames being affected by an increase in search errors, but only drip pricing and baiting having significantly more purchasing errors.

The highest marginal effect on error rates are observed under time-limited offers. This pricing frame increases the rates of search errors by 20 percentage points compared to the baseline. Drip pricing has the second highest marginal effect with 11 percentage points more errors. Baiting generates around 9 percentage points more. Even the reference pricing frame increases the rate of errors by almost 5 percentage points for search errors. These higher error rate induced by the simple "was X is now Y" statement appears striking, as the reference price is quite obviously meaningless in the experimental environment.

The regression results are broadly consistent with the welfare analysis, although the large lead in terms of errors for time limited offers that is not reflected in worse welfare outcomes (compared to drip pricing). There are a number of reasons behind this. First of all, we do not control for the size of errors in the regressions (precisely to offer an alternative view on consumers' performance that is less dependent on our parameter choices). Second, we do not account for multiple errors in the regressions. For example, subjects who search excessively (and through these multiple errors lose substantial amounts of money) are classified for the purpose of the regressions in the same way as a subject who makes a single mistake.

With respect to the variable "high value good", there is no significant effect of scaling for neither type of error. Hence, our results are robust to the incentives. In other words, higher stakes do not reduce errors.

Lower prices at the first shop are associated with lower errors. For the search errors, as we will see in more detail below, consumers exhibit in several treatments a tendency "to buy anyway" (instead of optimally continuing their search). This behaviour is obviously only erroneous when prices are high. A similar argument may explain the effect of low prices on purchasing errors. Consumers tend to buy too much, which is harder to do when buying two units of the good is optimal due to the low price.[19]

[19]Notice the difference between price and value. High value refers to a higher multiplier for both, con-

The effect of the price at the second shop is significant for both types of errors, but the effect goes into opposite directions. Regarding the error rates in search behaviour, the effect (and the reasoning) of low prices is the same as for the price at the first shop. But for errors in purchasing behaviour, subjects make more errors if the price at the second shop is comparatively lower.

Search costs have an ambiguous effect on the quality of consumer's choice behaviour. As would be expected, they play no role in purchasing errors, but higher search costs reduce, on average, search errors. In order to better understand the effect of search costs we need to investigate their differential impact on behaviour across the different treatments. Hence, we re-run our regressions including interactions between search costs and treatment dummies. For search errors, we find a rather mixed output as the coefficient on search costs changes sign and some treatment dummies become more significant. In terms of interaction effects, search costs do show up negatively if interacted with drip pricing (-1.74), baiting (-2.25), and time-limited offers (-2.52). In other words, we find that higher search costs reduce search errors only under these three price frames and not in any of the others.

In order to determine whether high or low search costs are more detrimental on consumer behaviour in markets, we extend the third column regression on all errors to include interactions between search costs and treatments. Unsurprisingly, given the lack of effect of search costs on purchasing errors, the picture for general mistakes in decision making is virtually identical to what we have seen in the search error estimations. Higher search costs increase errors but in the price frames the reverse is true with varying degrees of strength. In all, we conclude:

Result 3 Price frames have a stronger effect on consumer behaviour in markets with lower search costs.

We will revisit the causes for the effects of the price frames further below when we investigate the how search patterns change from the baseline to the frames and discuss the behavioural biases that may drive these deviations from optimal consumer choice.

Search errors including interactions:

```
                                (Std. Err. adjusted for 166 clusters in subject)
```

sumer payoff and price. Changing this multiplier leaves the decision-relevant (relative) price completely unaffected.

s~sear~r	dF/dx	Robust Std. Err.	z	P>|z|	x-bar	[95%	C.I.]
p1	.1975332	.0653773	3.04	0.002	.746793	.069396	.32567
p2	.2108458	.0389828	5.45	0.000	.74507	.134441	.287251
highva~d*	.0118086	.012901	0.92	0.360	.504188	-.013477	.037094
c	.7435583	.3419767	2.17	0.030	.055751	.073296	1.41382
d2*	.0855523	.0436557	2.04	0.041	.143003	-.000011	.171116
d3*	.2181708	.0474424	4.84	0.000	.121553	.125185	.311156
d4*	.2436679	.0436406	5.89	0.000	.134831	.158134	.329202
d5*	.095863	.0424076	2.37	0.018	.132789	.012746	.17898
d6*	.3672871	.0489156	7.58	0.000	.134627	.271414	.46316
cd2	-.858145	.6280573	-1.37	0.172	.007943	-2.08911	.372825
cd3	-1.743011	.5780775	-3.00	0.003	.006687	-2.87602	-.61
cd4	-2.522612	.6120776	-4.13	0.000	.007361	-3.72226	-1.32296
cd5	-.6908013	.5507969	-1.25	0.210	.007545	-1.77034	.388741
cd6	-2.536856	.59026	-4.30	0.000	.007606	-3.69374	-1.37997
period	-.0016963	.0007885	-2.14	0.032	15.2829	-.003242	-.000151
aptitude	-.0044842	.0040033	-1.12	0.263	9.77017	-.01233	.003362
obs. P	.2737487						
pred. P	.2679906	(at x-bar)					

All errors including interactions

(Std. Err. adjusted for 166 clusters in subject)

error	dF/dx	Robust Std. Err.	z	P>|z|	x-bar	[95%	C.I.]
p1	.2775902	.0644307	4.34	0.000	.746793	.151308	.403872
p2	.2254784	.0453379	4.98	0.000	.74507	.136618	.314339
highva~d*	.0130762	.0140832	0.93	0.353	.504188	-.014526	.040679
c	.3947292	.3745265	1.05	0.293	.055751	-.339329	1.12879
d2*	.068712	.0440653	1.59	0.112	.143003	-.017654	.155078
d3*	.1887509	.0464953	4.14	0.000	.121553	.097622	.27988
d4*	.1695438	.0429796	4.04	0.000	.134831	.085305	.253782
d5*	.0516707	.0436784	1.20	0.229	.132789	-.033937	.137279
d6*	.2948716	.0480427	6.07	0.000	.134627	.20071	.389033
cd2	-.6607607	.6434637	-1.03	0.305	.007943	-1.92193	.600405
cd3	-1.228126	.6068582	-2.02	0.044	.006687	-2.41755	-.038706
cd4	-1.1892	.6106281	-1.95	0.052	.007361	-2.38601	.007609

```
     cd5 |  -.0414783   .6226765    -0.07   0.947   .007545   -1.2619  1.17895
     cd6 |  -1.719253   .6801258    -2.53   0.012   .007606  -3.05227 -.386231
  period |  -.0040434   .0008181    -4.92   0.000   15.2829  -.005647  -.00244
aptitude |  -.0132016   .0044718    -2.96   0.003   9.77017  -.021966 -.004437
---------+--------------------------------------------------------------------
  obs. P |    .340143
 pred. P |   .3354834  (at x-bar)
------------------------------------------------------------------------------
```

4.3 Zooming in

We have seen evidence that identifies drip pricing as the price frame that is most detrimental to consumers, both in terms of errors of both kinds and welfare loss consequences. Time-limited offers come second (with similar welfare losses but even more errors than drip pricing in search), which is surprising given that the literature so far has either ignored or exonerated time-limited offers. Consumers have the least problems with complex prices and sales frames where the former has no effect on errors and the latter no effects on welfare. They are, however, both not completely unproblematic. In the middle is baiting with systematically more errors and a significant welfare loss.

In order to better understand the behavioural forces that are the root cause of inferior consumer decision-making stems, we will now study search and purchasing behaviour in much finer detail. We have already seen that error rates (especially in search behaviour) differ drastically from the baseline to the price frames. We look at these decisions more closely now.

Before we actually turn to a detailed analysis of consumer behaviour under the different price frames, we examine the pattern of choices in the baseline and drop pricing at Shop 1.

Table 5.4 shows all observations we have for our subjects at the first shop in the baseline. The rows indicate what would have been optimal, the columns what has actually been chosen. So, the first row ("0") in the table indicates all situations where the optimal strategy prescribes further search (the purchase of 0 units at the first shop). The second row contains all cases where consumers should have bought 1 unit and the third all cases where consumers should have bought 2 units. The columns indicate the actual number of units purchased.

While the table shows that subjects do generally very well in this situation (80.0% of all choices buy the optimal number of units) it also reveals two interesting asymmetries. First of all, it shows that errors are typically errors in search while purchasing the wrong number of units is much rarer. In fact, 92.0% of all errors are search errors. The second asymmetry occurs within the class of search errors. While subjects do not buy in 87.3% of all cases when it is optimal to continue to search, they buy optimally only in 70.4% of all cases where it is optimal to buy. In other words, there is a clear tendency to over-search. This is particularly dramatic when subjects should optimally buy just one unit. In this case only 54.6% of choices are correct and subjects over-search 40.2

Table 5.4: Optimal vs actual choices at the first shop visited in the baseline

		Actual choice					
		0	1	2	3	4	Total
Optimal choice	0	809	99	18	0	1	927
	1	78	106	8	1	1	194
	2	104	14	390	0	2	510
	Total	991	219	416	1	4	1631

Table 5.5 shows a similar table for choices made at Shop 1 in the Drip pricing frame. This time, only 73% of all choices buy the optimal number of units and whilst approximately the same number of errors are search errors (91.4%), in this case the pattern of search errors is dramatically different with 25.6% under-search compared with 12.6% in the baseline and over-search about the same at 17.3%.

Table 5.5: Optimal vs actual choices at the first shop visited in Drip pricing

		Actual choice					
		0	1	2	3	4	Total
Optimal choice	0	268	66	26	0	0	360
	1	22	34	6	0	0	62
	2	34	6	131	0	2	173
	Total	324	106	163	0	2	595

This pattern of search error reversal is present in all of the price frames. Tables similar to 5.4 and 5.5 for the other treatments can be found in the appendix where there are also corresponding tables for behaviour at Shop 2; Shop 2 behaviour is more similar across treatments and is consistent with behaviour at Shop 1: those subjects that go to Shop 2 optimally tend to act optimally with no bias (to over- or under-search) whereas those that go to Shop 2 through over-searching tend to over-search there also. The exception is time-limited offers where 30% of subjects under-search at Shop 2.

For all pricing frames, table 5.5 summarizes the percentage of optimal (unit) choices for both shops, as well as the share of consumption choices that can be classified as over-search and under-search. A subject engaged in over-search whenever she visited the shops more often than she optimally should have. Under-search occurs whenever a subject visits fewer shops than she optimally should have.

Table 5.5: Optimal choices and suboptimal search patterns under the different price frames

	Shop 1	Shop 2		
	Optimal choices (%)	Optimal choices (%)	Over-search (%)	Under-search (%)
Baseline	80.0	76.8	12.2	9.4
Complex pricing	77.4	77.7	5.4	19.0
Drip pricing	72.8	65.3	13.4	17.8
Baiting	70.0	76.7	6.1	22.7
Reference pricing	76.5	71.6	13.2	12.9
Time-limited offers	76.9	54.4	8.3	32.6

At a first glance, all of the price frames make is harder for the consumers to buy the optimal number of units at both shops. Even more strikingly, the pattern of relative over-search that is observed in the baseline treatment, is replaced by a pattern of relative under-search under all price frames but the reference price frame. In other words, while many consumers who should have bought at the first shop in the baseline decided to check out the second shop, this is reversed under all frames but the reference pricing frame. Consumers who should continue to search are lured into buying.

Let us now examine the effect of drip pricing. Comparing the numbers with the baseline frame reveals how dramatic the effect of drip pricing really is. The rate of optimal decisions falls to 72.8% for Shop 1 and to 65.3% for Shop 2. But even more strikingly, the phenomenon of relative over-search that is so prevalent in the baseline is still present and complemented by a stark pattern of under-search. In 25.6% of all cases where consumers should not have bought from the first shop they do so now. This compares to just 13.0% in the baseline.

These findings appear all the more remarkable as the difference between the two environments is really rather small. In essence, it just requires two extra clicks to see the true full price under drip pricing. Everything else is completely identical. Moreover, the environment is very simple and the subject population is highly selected and presumably much more capable of sophisticated behaviour than the average consumer.

Under time-limited offers, baiting, and complex pricing over-search is mostly eradicated and replaced buy a very strong pattern of under-search. Hence, consumers tend to buy when optimal behaviour suggests they should not. Somewhat surprisingly, for baiting there is no indication of consumers' "punishing" sellers that lured them into their shops with low offers that then turn out to be unavailable.

Finally, let us look at the reference price frame, in many ways the weakest treatment as it should be easy for subjects to understand that the former price that is mentioned is entirely meaningless. While the welfare losses that subjects incurred under the sales frame were statistically not significant we did find significantly higher error rates in our regressions. Interestingly, even the sales frame destroys the strong pattern of over-search that we detected in the baseline. However, this is not reversed into an under-search pattern. Rather over- and under-search are roughly equally frequent. We summarize this as:

Result 4: A pattern of over-search in the baseline is replaced by a pattern of under-search in the price frames.

Table 5.7: Suboptimal purchasing patterns under the different price frames

Price frame	Over buy rate	Under buy rate
	Total (Shop 1, Shop 2)	Total (Shop 1, Shop 2)
Baseline	0.059 (0.036, 0.055)	0.029 (0.021, 0.038)
Complex pricing	0.054 (0.045, 0.058)	0.04 (0.024, 0.061)
Drip pricing	0.119 (0.081, 0.117)	0.022 (0.026, 0.017)
Baiting	0.077 (0.098, 0.027)	0.039 (0.039, 0.043)
Reference pricing	0.058 (0.033, 0.064)	0.031 (0.033, 0.037)
Time-limited offers	0.036 (0.048, 0.031)	0.050 (0.048, 0.035)

In Table 5.7, we examine patterns of purchase errors. We have already seen in the Table 5.2 and Table 5.3 that there is a significant proportion of purchasing errors (defined as buying the wrong number, given that you buy and ignoring whether the correct search decision has already been made). Purchasing the correct amount would not seem to be a difficult decision and the prices frames do not add significant difficulty to this, yet we still see in Table 5.7 significant differences between the baseline and drip pricing and baiting – with more overbuying in these two treatments and less in time limited offers[20]. Two facts are striking from this table: firstly, that in drip pricing, subjects significantly over buy at both shops. Subjects see a low price, decide to buy the "correct" and then find the drips added. Notice two things: firstly, they can cancel at any time and secondly that they see the final total price before they click to confirm the purchase.

The second striking feature is that there is over buying in baiting, but only at Shop 1. An explanation is that they see the price advertized on the home screen, and try to click quickly through the buying process without noticing that the advertized price is no longer available.

Table 5.6: Achieved sales potential

[20]The reason the total may be larger or smaller than both Shop 1 and Shop 2 is we have not considered behaviour upon return to Shop 1

Frame	% sales potential Shop 1	% sales potential Shop 2
Baseline	88.1	107.1 (88.3)
Complex pricing	130.2	94.1 (88.5)
Drip pricing	107.8	90.9 (82.6)
Baiting	120.6	93.0 (89.3)
Reference pricing	89.7	98.6 (86.2)
Time-limited offers	101.2	168.5 (130.5)

It is interesting to examine the shops' sales volume, relative to what they would sell under optimal consumer behaviour (that is, the total number of units sold). Table 5.6 compares all treatments with regards to their achieved sales potential. Notice that this table is not equivalent to actual total sales as it does not include sales through return visits or sales to consumers who strayed from the optimal strategy at a previous stage (the number in brackets includes those that strayed from the optimal strategy to visit Shop 2). Therefore, the table is more indicative of consumer behaviour than of total profitability for firms which will analyse separately below.

In the baseline, the first shop sold 0.66 units per customer which compares to 0.74 units that are predicted under optimal consumer behaviour.[21] In other words, in the baseline the first shop reaches only 88.1% of its sales potential. In contrast, the second store achieves 107.1% of its sales potential.

As a consequence of the reversed search patterns under drip pricing, baiting, and complex pricing, many more units are sold at Shop 1. The first store now sells more units as predicted, and not fewer as in the baseline. However,the second shop does not benefit to the same extent as the first shop, even conditional on being visited. In fact, it only reaches about 92% of its sales potential under these three price frames

An exception with regards to the effect of the price frames on sales potential is time-limited offers. Even though the phenomenon of over-search is reversed just as under drip pricing, baiting, and complex pricing, there is a marked difference of consumer behaviour concerning sales. With time-limited offers the achieved sales volume is slightly smaller than under drip pricing (albeit still much better than under the baseline) but qualitatively similar. While the second shop could not benefit from drip pricing (if

[21] Notice that these numbers can be directly computed from the tables 5.5 showing optimal vs. actual choice.

anything consumers reversed to over-search when facing drips at the second shop), it does benefit substantially from the time-limited offers consumers faced at the first shop. It (the second shop) reaches an enormous 168.5% of its sales potential. From that it becomes apparent that consumers fear to face substantially higher prices upon return to the first store. In that sense, time-limited offers do the trick, they do confuse consumers substantially. But, somewhat ironically, the competition benefits more from this than the shop first visited by the consumer.

As the reference price frame has the weakest impact on the search pattern compared to the baseline, the sales potential is not different from the baseline.

The overall picture from this more detailed analysis is that price frames are effective because they mainly change behaviour at the first shop. To some extent they reduce inefficient over-search. However, in particular under drip pricing, baiting, and complex pricing, the picture actually reverses and consumers buy too often at the first shop. Relative to the baseline this implies even worse outcomes for consumers.

Given the general tendency of experimental subjects to explore environments they are faced with rather more fully than they may in the field, our results may underestimate the consumer detriment caused by these practices simply because an experiment like this will always bias subjects towards some over-search. And clearly, with less search, the baseline would perform better and the different price frames would perform even worse.

4.4 Learning and IQ

Given that the decision environment in this experiment is not completely trivial and given that subjects take the same kind of decisions repeatedly, there is, just as in real life, ample scope for learning. In this subsection we discuss whether there is any evidence for learning and whether learning is different across the different price frames.

In a first step we refer to the regressions on errors in subjects' behaviour (errors are documented above in Tables 5.2, 5.3, and 5.4) which control for a linear time trend and a measure of IQ (logical reasoning ability). The linear time trend variable turns out to be highly significant and important in terms of its impact. The estimate is that per period the error rate falls by 0.40 percentage points. This may initially look like a small number but remember that there are thirty periods altogether. A linear approximation would, thus, suggest that subjects' error rate is more than 12 percentage points lower at the end of the experiment than at the beginning. Allowing for non-linear effects (through a quadratic term) suggests that learning is levelling out towards the end of the experiment. Comparing the size of the learning effects with the overall error rates that we have examined above we can conclude that learning substantially reduces but does not eliminate erroneous behaviour.

Decomposing the errors into search and purchasing errors, similar pictures emerge although the speed of learning is slower when it comes to search errors. These are estimated to decline by 0.2 percentage points per period while purchasing errors decline by 0.4 percentage points.

We also see the effect of IQ. We conducted a short test at the end of the experiment. The aptitude score is between 0 and 12 and the marginal effect of one extra point is estimated to be 1.2 percentage points decrease for overall errors, 0.4 percentage points for search errors, and 1.2 percentage points for purchasing errors. The latter two are highly significant.

In a further stage, we add interaction terms between price frames and the linear time trend in order to examine whether learning speeds differ between different time trends. For the overall error rate we find no differences in learning speed between the baseline and the different price frames with the exception of time-limited offers. For time-limited offers learning is significantly slower (and very significantly so). In fact, the estimate is so big that it completely wipes out the general learning trend. On balance, we find that there is no learning at all under time-limited offers. Consumers appear to believe that the first shop will always charge a higher price upon return and therefore typically do not return. Hence, they can never learn that their belief is wrong. To what extent this holds in real-life markets will crucially depend on the precise informational structure. For example, for shops located along a high street, consumers might notice if time-limited offers advertised in a window display turn out to be repeated over time or are replaced by even lower prices. In other markets where such information does not come for free it would be similarly hard to learn as in our experiment.

The same findings hold for the decomposed search and purchasing errors: there is no differential speed of learning between the baseline and the other price frames with the exception of time-limited offers where there is simply no learning at all. Neither search, nor purchasing errors are reduced over time in the presence of time-limited offers. In fact, search errors are even slightly increasing over time for time-limited offers.

Clearly, this examination of learning renders our finding that consumers have trouble with time-limited offers substantially more worrying. Not only is it difficult for subjects to optimally adjust their search strategy in the presence of time-limited offers they also are not able to improve their performance over time.

As we have evidence on declining error rates that are attributed to learning we can also examine whether consumer welfare increases over time. For that purpose we conduct simple non-parametric tests to compare subject welfare in the first half of the experiment with subject welfare in the second half of the experiment.

Welfare is found to be significantly increasing in the second half of the experiment under all price frames (including the baseline) with one exception: time-limited offers. This is shown in the following table we report welfare loss in the 1^{st} and 2^{nd} half of the experiment sessions by price frame. Note that whilst the gap between the welfare between the price frames and the baseline closes, the regression analysis from Table 5.3 remains very similar even when restricted only to the second half of the experiment.

Table 5.18: Welfare loss across 1^{st} and 2^{nd} half of the experiment sessions

Frame	1^{st} half	2^{nd} half	Average
Baseline	0.058	0.017	0.038
Complex pricing	0.068	0.035	0.052
Drip pricing	0.091	0.034	0.064
Baiting	0.067	0.029	0.048
Reference pricing	0.084	0.028	0.055
Time-limited offers	0.073	0.055	0.064

5 Behavioural forces at work

We conclude our empirical analysis of consumer behaviour by reflecting on what the data tell us about behavioural forces at work. We briefly discuss each of the five pricing practices under investigation.

Drip pricing: As we have seen drip pricing turns consumers who tend to search too much into consumers who tend search too little. This effect is particularly striking as in our experiment the drips are revealed through just two mouse clicks and consumers can see the total price very clearly before they make their final purchasing decision. The objective costs of going through the drips are very close to zero (just a few seconds that pass for the two clicks). Accordingly, it is completely implausible to attribute the change in behaviour to increased costs of search and sunk costs. In principle, consumers might *rationally* decide to accept higher prices after being led through complicated drips if they were to expect equally costly practices elsewhere. Accepting a higher price at the first outlet would after all avoid the costs of clicking through a labyrinth at a competing outlet. Notice that this is an entirely rational response and has nothing to do with the so-called sunk cost fallacy where consumers ignore that they cannot recover the costs of (failed) activities and, consequently, may "throw good money after bad". Here the extra search costs are so tiny that any explanation along the sunk costs line is simply not justified. Rather the data suggest that consumers who see a low base price and do not yet know that the effective price will go up through "shipping and handling" charges experience an increase in their willingness to pay for the good which is in line with loss aversion and the so-called endowment effect.[22] Consumers who decide to buy the product at the low price experience a shift in their reference point as they already imagine departing with the good. Changing the initial decision, that is, giving up the good that is already in the virtual basket would be perceived as a loss. This loss can be avoided by purchasing the product despite an increased price.

Time-limited offers: Time-limited offers eliminate over-search at the first store but have an even more dramatic effect on the sales at the second store. Consumers who have rejected the time-limited offer at the first store simply tend not to return to it even if the price at the second store is comparatively high. The underlying problem in consumer behaviour is obvious: As the consumers believe the store (that is, as they

[22]This is similar to the findings of previous researchers such as Morwitz, Greenleaf and Johnson (1998), and Hossain and Morgan (2006) discussed in chapter 3.

erroneously believe that prices will go up) they have a) a tendency to buy more now at the first store, and b) if they don't buy at the first store, they buy at the second shop if this appears at all profitable. This false preconception can then never be revised, simply because consumers do not learn that time-limited offers are not real offers with true discounts but that prices can go up and down once they expire. Consumers simply do not understand the real mechanism because of naive beliefs, confusion and too little exploration. In all, this effect appears to be due to purely cognitive problems.

Baiting: As under drip pricing consumers buy too much at the first shop under baiting while their behaviour at the second shop is largely optimal. As baiting generates expectations about good deals, it appears once again that it is loss aversion and the endowment effect that is at work. Consumers pick a store from which they expect a good deal and this act in itself raises their willingness to pay for the good in comparison to the baseline. In contrast, to standard experiments on the endowment effect it is, once again, anticipated or imagined ownership that causes the positive shift in consumers' valuation of the good.

Complex pricing: This frame is different from the other price frames in that it actually lowers the prices. Controlling for that, we find that consumers make slightly higher welfare losses under complex pricing than under the baseline. Error rates are not significantly higher than in the baseline but the type of error is very different. Instead of over-search consumers under-search and buy too often at the first shop. Notice that they do not simply buy too many units of the good (which would follow from an endowment-effect like shift in valuations) but that they do buy the offer. This suggests that the offer has an attraction beyond the mere reduced price or that consumers are cognitively limited.

Sales frame: Finally, we find that somewhat surprisingly even the meaningless by-line "was X" where X is a higher price than the current price eradicated over-search. However, the sale frame is not able to trigger loss aversion or the endowment effect to the point where consumers start to under-search. Rather it generates evenly distributed errors suggesting mainly cognitive problems with processing information that optimally should be discarded.

In light of the behavioural forces that we have identified it is worthwhile to revisit the issue of how search costs impact on consumer choice. As we have seen earlier, search errors are falling in search costs under drip pricing, baiting and time-limited offers –

precisely those treatments where we now conclude that search errors are driven by loss aversion. So why and how do search costs matter in these treatment? The nexus is simple. Loss aversion increases consumers willingness to pay at the first shop (once they imagine themselves as owner of the good). Now we simply need to observe that the larger the search costs, the more likely it is that the decision to buy at a comparatively high price is actually optimal! In other words, in environments with large search costs it matters less if consumers want to buy straight away because they would experience not buying as a loss.

We have also argued that purchasing errors are either due to cognitive failure or to the sunk cost fallacy and we have seen earlier that they increase with higher search costs under all treatments with the exception of time-limited offers and drip pricing (but including the baseline). The basic intuition for the impact of search costs on the sunk costs fallacy is obvious. The higher the search costs consumers have spent, the bigger their urge to justify them through purchasing too many units. It is less clear why this interacts with the different treatments. One possibility is that, insofar purchasing errors occur after prolonged search, that those who are prone to loss aversion are less likely to make them (simply because they tend to buy straight away at the first shop). Furthermore, if these biases are correlated (as recent research by Burks et al. (2008) suggests) then this implies that under those treatments that trigger loss aversion there will be fewer biased consumers who do long searches. In other words, the heterogeneous treatment effect on purchasing errors is quite plausibly simply due to a selection effect. This would then suggest that the sunk cost fallacy occurs indeed independent of price framing.

6 Sellers and total welfare

In our analysis we have so far very much focussed on the effect of price frames on consumer behaviour and welfare although we had some results that pointed towards the effect on firms and, generally, the two are of course closely linked.

While shops were completely computerized it is still interesting to ask how their performance is affected by price frames. Price frames that perform well for the shops are not necessarily those that are detrimental for consumers. Some price frames could hurt both buyers and sellers. Of course, one would expect that in a market environment mainly those price frames are used that actually improve sellers' performance.

Table 5.18 shows average number of units sold by the two shops under the different price frames as well as average turnover. The table also contains units sold and turnover for the entire industry. These industry performance indicators are perhaps the most important ones as, a specific shop, is in our experiment equally likely to become the first or the second shop.

Table 5.18: Average number of units sold by the two shops

Frame	Units Shop 1	Units Shop 2	All units (Industry indicator)	Revenue Shop 1	Revenue Shop 2	Total revenue (Industry indicator)
Baseline	.93	.62	1.56	0.60	0.42	1.02
Complex pricing	1.65	1.05	2.71	1.10	0.72	1.82
Drip pricing	.98	.60	1.58	0.64	0.39	1.03
Baiting	1.22	.38	1.59	0.77	0.24	1.02
Reference pricing	.94	.64	1.58	0.61	0.42	1.03
Time-limited offers	0.88+0.13	.50	1.52	0.57+0.09	0.35	1.00

The table shows impressively that, contrary to what we might have expected, the shops gain nothing from employing the different price frames. In fact, under time-limited offers, which is one of the two most detrimental practices for consumers, shops experience even reduced sales.

Sales for complex pricing are, of course, substantially higher but whether this would translate into higher profits depends entirely on unit costs which remain unmodelled. If unit costs are equal to the average market price (as suggested by a Bertrand model) then complex pricing would not be profitable. If there is a higher average mark-up (say if costs were equal to the lowest price ever charged) then complex pricing would be profitable for firms.

While the other price frames have essentially no effect on industry performance there are dramatic shifts in the distribution of sales, generally to the advantage of the shop that is first visited (which is in line with what we have said above about the reversal from over-search to under-search). While in our experiment, the sequence in which shops are visited is essentially random (with the exception of baiting), firms can in many markets influence the order in which consumers search, for example through advertising or sponsored links in search engines. Our results suggest that the incentives to engage in such activities are substantially increased through elaborate price framing. We would thus predict that, across different markets, the occurrence of price framing is positively correlated with the amounts firms spend on enticing consumers to search their offers first. This is a prediction that could be easily tested using data from auctions for sponsored search.

Of course, such activities to change consumers' search order would reduce total welfare even further. In our setup, price framing reduces total welfare already in any case as it harms consumers and does not benefit firms. However, contests for being searched first would render the picture even bleaker.

7 Discussion & and implications for non-Laboratory markets

External validity for studies of this nature is typically highly asymmetric. Both simplification and stylization of the decision problems and the highly selected subject pool imply that we are more likely to observe good or even optimal performance in the laboratory than in the field under real-life conditions. This implies that whenever we find close to perfect performance external validity is severely limited. If student subjects do well in a simple task it is hard to conclude from that the general population would do well in a more complicated task. However, the other way round things look much brighter. If a highly selected student sample does badly in a simple decision environment that also offers scope for repetition and learning, it would be very surprising if the general population did much better in more complicated situations.

Given this asymmetry, the design choices we took were risky. We designed a very simple search environment that was, given our requirement (multiple units, two stages) pretty much minimal. It would have been difficult to come up with a simpler environment. Similarly, the implementation of the price frames was typically simple and subjects could experience them repeatedly in almost identical manner. The risk of these design choices was that we might have found close to optimal performance in all treatments in which case we would have learned very little from this study.

The alternative strategy to design a more complicated environment would have entailed different risks. In more complicated environments decision errors and noise will invariably go up and accordingly it will be more difficult to detect differences between treatments for given sample sizes.

However, as we have documented we have observed substantial difference between treatments with surprisingly poor performance in some. The results on drip pricing stand out. Being "just two clicks away" from the baseline its effects on search patterns and performance are dramatic. Outside the laboratory where drip mechanics are more elaborate and it is more time consuming to reach a stage where full prices are clearly visible it is likely that the effects that stem from loss aversion or the endowment effect will be even stronger. On top of that, more elaborate drips will also increase the true costs of searching for the price which will enhance the effects from loss aversion. Similarly, we have not tried to optimize the drip sizes and it would be very bewildering if we had, by accident, stumbled across the most effective drips. Once again, this suggests that, if anything, we might still underestimate the true consumer detriment resulting from drip pricing.

Given our basic design choices and selection of subjects the observation holds, of course, more generally: there is a built-in tendency to underestimate consumer detriment for all price frames. This also implies that the sales frames which receives almost a clean bill of health in this study could potentially be more harmful than we detect. As soon as we are in environment where former prices contain some hard information there is much more scope for consumers to process this information in a less than adequate way. On the other hand, references to former (higher) prices could also serve as a useful signal of quality.

In this context it is important to notice that, with the exception of complex (3 for 2) pricing, we isolate the pure effect of price frames and not of offers. Of course, real offers with lower prices might benefit consumers even if their presentation confuses them or triggers some behavioural biases. The net effect of lower prices and the adverse consequences we measure here might still be positive. What this study shows, however, is that also real offers could benefit consumers more if presented in a straight way as in our baseline treatment.

In our experiment, both price frames and prices themselves are exogenously fixed while in real markets they are, of course, chosen. Given our consumer data it appears clear that certain price frames will allow firms to charge higher prices (in particular those that trigger loss aversion and the endowment effect as these effects are akin to increased willingness to pay or an outward shift of the demand curve). Consumers would then suffer doubly, from their direct negative consequences we measure in this experiment *and* from the higher prices.

One aspect of endogenous choice of price frames that we have not studied at all is that sellers in the same market might choose different price frames which makes price comparisons much harder. As Chioveanu and Zhou (2009) show these effects can even overturn standard intuition on how the number of firms in a market relates to consumer welfare. With added confusion from a greater variety of price frames, consumers might actually suffer from the entry of additional sellers.

On the other hand firms may elect to not use price frames that annoy customers. Firms may seek to establish a reputation for not using annoying practices, such a drip pricing.

While we can neither validate nor reject these theoretical predictions we can say a little about how our findings on relevant behavioural biases would impact on markets. Clearly, the strongest force that causes consumer detriment in our experiment is the endowment effect or loss aversion. Consumers' imagination of owning a good shifts their willingness to pay. We observe strong evidence on this in both drip pricing and baiting. In the field there will be many other practices of hot selling that play on these effects. If the consumer tries out a product in a shop it will give him some objective information about how the product handles but it also makes envisaging ownership easier and what we have seen here is that envisaging ownership is all that is needed to increase willingness to pay.

There are, of course, many institutional and physical details that will matter for the effect of these practices in non-laboratory markets. For example, it might be easier to encourage the imagination of ownership for some goods than for others. By thinking about the product characteristics that make imagination of ownership easier, we could then derive comparative static predictions about in which markets we would expect to observe certain price frames more frequently.

Similarly, there might be particular characteristics of sellers that tinge the decision problems in real-life markets. For example, for closing-down sales (where, say, the consumer can see that a building is about to be torn down) the time-limitedness of offers might be more credible than for other "mid season" sales.

For one important aspect of real-life markets we do have some indication in our data, the role of search costs. Our analysis suggests that the detrimental effects of loss aversion increase with lower search costs. This is intuitive: Buying too early too often, tends to coincide with optimal behaviour when search costs are high. On the other hand, we have found that the sunk cost fallacy (that drives some of the purchasing errors) increases with search costs. Again this seems plausible for real-world applications. The more time and money consumers have spent on search, the more desperate they might be to justify these high search costs through making (too many) purchases. In our experiment, this effect is much smaller than the effect of search errors which can be traced to loss aversion. However, this may well be a consequence of our parameter choice. For example, in markets with very high search costs (for example, because of location in just one city in a larger geographical area, say, some sort of fair) it is plausible that consumers might be very frustrated to leave empty-handed and would thus be tempted to buy more than they would have bought had the market taken place at their doorstep.

Summarising, let us stress however again that we have good reason to believe in the general external validity of our results – that these practices do cause consumer detriment and that what we identify in the lab is probably rather the tip of the ice berg as there are many aspects of real-life markets that will accentuate the problems we document here.[23]

8 Policy implications & recommendations

From a policy point of view, this experimental study has two broad implications. We identify price frames that clearly cause consumer detriment. As they do so in a comparatively simple environment and with a comparatively sophisticated subject pool, it appears clear that these practices are also harmful outside the laboratory.

[23] Although inevitably there are also likely to be some factors in real life which will mitigate concerns about practices (even when frames are false offers) such as the desire for firms to build reputation, and consumers to learn about honest firms, as discussed in the next section.

Our findings suggest that firms will be tempted to confuse consumers through drips, baits or time-limited offers. Clearly, the occurrence of any of these three practices which do not represent genuine offers might provide reason to worry.

While drip pricing has previously been known to be problematic, time-limited offers have never before been identified as a source of consumer confusion and detriment. As we have shown, consumer errors under time-limited offers are particularly severe in that they are not reduced through learning (at least in this setting). Thus, even in markets for goods that are purchased frequently, one would expect that time-limited offers are used and indeed problematic.

Consequently, one may be particularly worried about drip pricing in markets for not particularly frequently purchased goods or time-limited offers in markets for goods that are purchased with high frequency.[24]

We also have clear indication that in this experimental setting these effects are particularly pronounced in environments with low search costs for consumers. However, this result would require further analysis before clear policy conclusions could be drawn from it. Not least because, with higher search costs consumers become prone to the sunk cost fallacy and this effect could become much stronger if search costs are substantially higher.

If the main effect of price frames is that they shift demand from one firm to another, enforcement (which helps to promote a level playing field) may be welcomed not only by consumers but also by firms. But this needs more research into environments with endogenous prices. This could be done in the same experimental framework.

If consumers are annoyed by some price frames, there is some scope for the self-healing powers of the market. Firms may gain a reputation for not using such practices. This is more likely to work for practices that are indeed perceived as an annoyance such as drip pricing. In contrast, time-limited offers might, in fact, be perceived as something positive by consumers who will not even be aware of their struggle with understanding the true nature of such offers. Accordingly, one would have less hope that the market can overcome the occurrence of such positively perceived practices. These issues could be investigated in a similar experimental study where firms are no longer simple static

[24] Clearly, as a seller one would rather employ frames that are robust to learning for frequently purchased items.

computers.

9 Appendix

Experiment instructions

Welcome to our experiment! In the course of this experiment you can earn a substantial amount of money. The precise amount will depend on your choices and some luck. We kindly ask you to remain silent throughout the entire experiment. Do not attempt to communicate with your neighbours and do not try to look at their screens. If you have any questions, please, raise your hand and we will come and answer it in private.

This experimental session will consist of several on-screen stages:

1. **Quiz** (to ensure you understand the instructions)
2. **Experiment** (described in the instructions below)
3. **Multiple choice quiz**
4. **Questionnaire about yourself**
5. **Feedback questionnaire about the experiment**

Your payment for this session will consist of the amounts you earn in Stages 2, 3 together with the £5 show up fee. We will pay you in cash. You will need to sign a receipt, which we will supply.

In this experiment, we simulate the idea of shopping for different products. The experiment will have 30 periods and in each period, you can buy zero or more units of a product that is available to buy at two different shops (the same product is available at both shops). By buying units of the product, you will get some *points*.

At the end of the experiment, the total sum of all the points you earned over the 30 periods will be converted into pounds, at a rate of £1 for 50 points.

Let us describe the experiment. You can buy four different products. They are called GREEN, ORANGE, BLUE, and RED. For each of these products, you can see in the attached table (on page 3) how many points you will get depending on the number of units you buy in a period. For example, if you buy a total of two units of GREEN in a period you get 100 points. Or for a total of three units of RED you get 220 points.

The prices for these products will also vary from period to period and shop to shop. The table shows you, for each good, the price range *in points* you can usually expect. This is the price is *per unit.* The price will be randomly chosen each period from the price range and the price for each shop will be determined separately.

At the beginning of each period, you will see your *home screen.* This will inform you of which product is available to buy in this period (only one product will be sold in any one period and it is chosen at random) so that you know which row of the table is relevant. The product available will also be identified by the colour of the screen surround. On your home screen, there are two buttons, representing the two shops where you can buy the product. They will be labelled "Go To Shop 1" and "Go To Shop 2". There is also a button labelled "I'm done" which ends the period and brings up the results.

If you want to go to a shop to buy something (or simply to check the price), you can do so by clicking on this shop's button. Every time you visit a shop, you will incur some costs (think of time it takes to get there, petrol, etc). The cost (denoted *in points*) is displayed on the home screen (the cost is the cost to make one return trip to a shop).

Once you are at the shop, you will see the price of the good that is available to buy in the period and, if you want to buy, you can enter the number of units you wish to buy (the maximum total number of units you can buy in any one period is 4). A box will appear where you need to confirm your purchase. Once your purchase is confirmed (or if you choose to cancel), you will be returned to the Shop screen (not to your home screen). If you make a purchase, a box will briefly appear informing you of the number of units you have successfully bought.

If you don't want to buy anything (because you find the prices too high or you first want to check out the other shop) or once you have finished shopping, you can click the *Go home* button to return to the home screen.

You can go back and forth between your home screen and the shops as often as you like but notice that you are charged for every trip you make. Once you are done with your shopping for a period, you can click the *I'm done* button at the bottom right corner of your home screen and proceed to the results screen, where you are informed of your decisions and earnings in the period.

Your earnings in each period depend on the number of units of the product you bought as shown in the attached table. From these points, we take away your travel costs and the amount of points you spent on buying the product.

For example, suppose in a period the GREEN product is available to buy and that travel costs are 3 and you take the following actions:

1. Go to Shop 2 and observe a price of 48 per unit for GREEN.
2. Go back to the home screen.
3. Go to Shop 1 and observe a price of 31 per unit for GREEN.
4. Buy 3 units of GREEN at Shop 1. [You are still at Shop 1 after the purchase]
5. Go to home screen.
6. Click "I'm done"

Then your earnings for the period would be:

Points for buying 3 units of GREEN: 110

Minus travel costs (two visits to the shops): 2 * 3 = 6

Minus cost of buying 3 units in Shop 1: 3 * 31 = 93

Your earnings for the period would be: 110 – 6 – 93 = 11 points

The prices shown above are purely for example and the strategy of the shopper in this case may or may not be good. The real prices for GREEN that you will encounter will be randomly drawn each period between 30 and 60.

Before the experiment, there will be a short quiz on the screen to check that you understand these instructions. You will need to get all the answers right before you can proceed to the experiment, but you can have a second, third, etc. go if you get an answer wrong.

After the experiment, there will be another on-screen quiz consisting of 12 multiple choice questions. Instructions will appear on the screen. If you get more than 10 correct, we will pay you an extra £2 or if you get more than 8 correct we will pay an extra £1. You have a total of 8 minutes to complete the quiz – it will time out after this point.

Following this quiz, there will be two short questionnaires where we ask you to give us your thoughts on the experiment, your strategy and about yourself.

At the end of the experiment, you will be asked to fill in a receipt. We will call you up to the front, one by one, to pay you in cash. After this, you are free to go.

Product	0 units	1 unit	2 units	3 units	4 units	Price Range
GREEN	0	60	100	110	115	30 to 60
ORANGE	0	80	140	170	195	50 to 80
BLUE	0	110	180	190	190	50 to 110
RED	0	120	200	220	230	60 to 120

10 References

Armstrong.M and Zou.J (2010) Conditioning prices on search behaviour, University of Munich Library, MPRA Paper.

Biswas, A, Wilson, EJ and Licata, JW (1993) Reference pricing studies in marketing: a synthesis of research results, Journal of Business Research, 27, 239–256.

Alford.B.L., and A Biswas (2002) The effects of discount level, price consciousness and sale proneness on consumers' price perception and behavioral intention, Journal of Business Research, 55, 775-783.

Bakos, Y. (2001) The Emerging Landscape for Retail E-commerce Journal of Economic Perspectives, Vol. 15, No.1 pp 69-80, 2001

Baye, M., Morgan, J., and Scholten, P. Price Dispersion in the Small and in the Large: Evidence from an Internet Price Comparison Site, Journal of Industrial Economics 52 (**??**), pp. 463-496, 2004

Baye, M.. Baye.M.R., Gatti. R.R and Morgan.J (2006) Did the Euro Foster Online Price Competition? Evidence from an International Price Comparison Site, Economic Inquiry, 44(**??**), 265-279.

Blair, EA and Landon, EL, (1981) The effects of reference prices in retail advertisements. Journal of Marketing, 45, 61–69.

Brehm, S. S. & Brehm, J. W. (1981). Psychological reactance, New York: Academic Press.

Brown, J. and Goolsbee, A. (2002) Does the Internet Make Markets More Competitive? Journal of Political Economy, 110 (**??**), pp.481-507., 2002

Brown.J., Hossain.T., and Morgan.J., (2007) Shrouded Attributes and Information Suppression: Evidence from the Field. Mimeo, University of California (*faculty.haas.berkeley.edu/rj* ***mo***

Burks, SV, Carpenter, JP, Gotte, L, and Rustichini, A (2008) Cognitive skills explain economic preferences, strategic behaviour, and job attachment, IZA Discussion Paper No. 3609.

Cialdini. R. (2009) Influence: Science and Practice, 5th Edition, Boston: Pearson.

Cason.T. and Friedman.D., (2003) Buyer search and price dispersion: a laboratory study, Journal of Economic Theory

Chandrashekaran, R. & Grewal, D. (2006). Anchoring effects of advertised reference price and sale price: The moderating role of saving presentation format. Journal of Business Research, 59, 1063–1071

Charness.G. and Chen.K. (2001). Minimum advertised price policy rules and retailer behaviour: An experiment by Hewlett-Packard.

Chioveanu, I, and Zhou, J (2009) Price competition with consumer confusion, MPRA Paper No. 17430. http://mpra.ub.uni-muechen.de/17340/

Davis.D and Holt,C (1996) Consumer search costs and market performance, Economic Inquiry, 34, 133-51.

DellaVigna, S (2009) Psychology and economics: Evidence from the field, Journal of Economic Literature 47, 315-372.

DellaVigna, S, and Malmendier, U (2006) Paying not to go to the gym, American Economic Review 96, 694-719.

DelVecchio,D, Krishnan,HS and Smith, DC (2007), Cents or Percent? The Effects of Promotion Framing on Price Expectations and Choice, Journal of Marketing, 71 (**??**), 158-170.

Devlin.J, Ennew.C, McKechnie.S, and Smith.A, (2007) A study of time limited price promotions Journal of Product Brand and Management, 280 – 285.

Dickson, PR, and Sawyer, AG (1990) The price knowledge and search of supermarket shoppers, Journal of Marketing 54, 42-53.

Ellison, G, and Ellison, SF (2005) Search, obfuscation, and price elasticities on the internet, mimeo, February 2005.

Gerstner, E, and Hess, JD (1990) Can bait and switch benefit consumers?, Marketing Science 9, 114-124.

Gabaix, X, and Laibson, D (2006) Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets, Quarterly Journal of Economics, 121, 505-540.

Hertwig, R, and Ortmann, A (2010) Experimental practices in economics: A methodological challenge for psychologists?, Behavioural and Brain Sciences 24, 383-403.

Hess, JD, and Gerstner, E (1998) Yes, "bait and switch" really benefits consumers, Marketing Science 17, 283-289.

Hossain.T and Morgan.J (2006) Plus Shipping and Handling: Revenue (Non) equivalence in Field Experiment on eBay, Advances in Economic Analysis and Policy, 6(**??**), 1-26.

Lynch.J. and Ariely.D. (2000) Wine online: Search cost affect, competition on price, quality and distribution, Marketing Science, vol.19(**??**)

Lynn, M. (1989). Scarcity effect on value: Mediated by assumed expensiveness. Journal of Economic Psychology, 10, 257-274.

Morton, FS, Zettelmayer, F, and Silva-Risso, J (2001) Internet car retailing, Journal of Industrial Economics 49, 501-519.

Morwitz, VG, Greenleaf, EA, and Johnson, EJ (1998) Divide and prosper: Consumers' reactions to partitioned prices, Journal of Marketing Research 35, 453-463

Ofcom (2010) Using experiments in consumer research, available at: http://www.ofcom.org.uk/research/telecoms/reports/experiments/

OFT (2009) Road testing consumer remedies, available at: http://www.oft.gov.uk/shared_oft/economic

Piccione.M. and Spiegler.R. (2009) Farming Competition, available at http://else.econ.ucl.ac.uk/newweb

Simonson, I. (1999). The Effect of Product Assortment on Buyer Preferences. Journal of Retailing, 75, 347–370

Wansink B., Kent, R.J., and Hoch S.J. (1998). An Anchoring and Adjustment Model of Purchase Quantity Decisions. Journal of Marketing Research, 65, 71-81

Wilkie, WL, Mela, CF, and Gundlach, GT (1998) Does "bait and switch" really benefit consumers? Marketing Science 17, 273-282.