Oligopolistic Competition with Choice-Overloaded Consumers

Gerasimou, Georgios and Papi, Mauro

University of St Andrews, University of Aberdeen

18 December 2015

Oligopolistic Competition with Choice-Overloaded Consumers *

Georgios Gerasimou† — University of St Andrews

Mauro Papi‡ — University of Aberdeen

November 15, 2016

Abstract

A large body of empirical work has suggested the existence of a *"choice overload"* effect in consumer decision making: When faced with large menus of alternatives, decision makers often avoid/indefinitely defer choice. A suggested reason for the occurrence of this effect is that the agents try to escape the higher cognitive effort that is associated with making an active choice in large menus. Building on this explanation, we propose and analyse a model of duopolistic competition where firms compete in menu design in the presence of a consumer population with heterogeneous preferences and overload menu-size thresholds. The firms' strategic trade-off is between offering a large menu in order to match the preferences of as many consumers as possible, and offering a small menu in order to avoid losing choice-overloaded consumers to their rival, or driving them out of the market altogether. We study the equilibrium outcomes in this market under a variety of assumptions. We also propose a measure of consumer welfare that applies to this environment and use it alongside our model to provide a critical perspective on regulations that cap the number of products that firms can offer.

Keywords: Choice overload; bounded rationality; oligopolistic competition; behavioural industrial organization

JEL Classifications: D01, D03, D04, D21, D43

*We thank Luke Froeb, Paola Manzini, Patrick Rey, Ariel Rubinstein, Rann Smorodinsky, David Ulph and audiences at Bounded Rationality in Choice (2016, Northwestern), Bounded Rationality & Mechanism Design (2016, Glasgow), St Andrews, Technion and the Athens University of Economics & Business for very useful comments. We also acknowledge beneficial discussions with Yuval Salant, Hugo Sonnenschein, Nikolaos Vettas and Rakesh Vohra. Any errors are our own.

†Corresponding author. Email address: gg26@st-andrews.ac.uk

‡Email address: m.papi@abdn.ac.uk

"As a neophyte shoe salesman, I was told never to show customers more than three pairs of shoes. If they saw more, they would not be able to decide on any of them."

Letter to the editor, New York Times, 26th January 2004[1]

1 Introduction

An important implication of rational choice theory is that increasing the number of consumption alternatives can only make a decision maker better off by enabling him to choose an option that is ranked higher in his preference ordering. Much doubt has been cast on this prediction since the work of Iyengar and Lepper (2000) who reported experimental evidence suggesting that it is significantly more likely for large menus of options to result in the consumer choosing *none* of the market alternatives available to him because of the higher degree of complexity that is associated with making an active choice in such menus. This has come to be known as the *choice overload* or *too-much-choice* effect and has now reached the status of being discussed in leading undergraduate microeconomics textbooks.[2] Importantly, this phenomenon has been observed in experimental as well as real-market environments over important economic decisions such as employee participation in pension savings plans (Iyengar, Huberman, and Jiang, 2004; Iyengar and Kamenica, 2010).[3]

At least partly in response to the large body of empirical work that has followed the original experimental evidence on choice overload,[4] formal recognition of the potentially harmful effects that large menus could have on consumer welfare was recently made by regulatory authorities, such as the UK Office for Gas and Electricity Markets (Ofgem). Specifically, in 2013-2014 Ofgem forced UK energy suppliers to ban complex tariffs and restrict the number of products they could offer to *no more than four*. The justification for this decision was that *"together these changes will make it far easier for consumers to compare deals and find the best tariff for them"* (Ofgem, 2014).

In addition to novel consumer welfare considerations, once it is acknowledged that consumers can become choice-overloaded by large numbers of products, important implications for firm competition also arise. In particular, rather than offering menus with as many products as possible, firms that interact with such consumers have a clear incentive to find a balance between offering menus with sufficiently many products in order to appeal to as many consumers as possible, and offering menus with sufficiently few products in order to not overload consumers and lose them to their competitors or drive them out of the market altogether. This strategic trade-off lies at the very heart of the novel model of duopolistic competition in menu design that we introduce.

The model assumes that consumers are heterogeneous in their preferences as well as in their overload characteristics, and that both firms know how these are distributed in the consumer population. For simplicity, and also to isolate the pure effect of overload on the market outcome, the firms' pricing decisions are suppressed by assuming that each product comes

[1]By Milton Waxman. Originally quoted in Kuksov and Villas-Boas (2010).
[2]See Varian (2014, p. 589), for example.
[3]See also Kamenica (2012).
[4]This literature is carefully surveyed in Chernev et al (2015).

with an exogenously given markup. Hence, to analyze consumer welfare in this setting where consumer surplus is inapplicable we introduce a simple and novel welfare index that is increasing in the number of products that are offered in the market through the two firms' menus, and decreasing in the number of consumers who are overloaded at these menus.

We show that the Hotelling-type maximum-variety/minimum-differentiation symmetric profile where both firms offer all products in the market is always an equilibrium when no consumer is ever overloaded. Moreover, this becomes an equilibrium in strictly dominant strategies when products are equi-profitable. Under the latter assumption we also identify necessary and sufficient conditions for minimum-variety/minimum-differentiation equilibria to exist in which firms offer the same one product. We relate this result to the real-world example of the Apple-Samsung over-specialization in the global market for mobile phones. We also study the symmetric equilibria that arise in the cases of homogeneously distributed overload under heterogeneous markups and identify conditions for uniqueness in this class. We then move on to provide a general condition under which a full product differentiation equilibrium does not exist under any distribution of the consumers' overload characteristics.

As a policy application of our model we re-examine Ofgem's above-mentioned decision to cap the number of energy tariffs that UK providers could offer to four. We first examine the conditions under which it is welfare-maximizing for firms to offer the same menu with four options when overload in the consumer population is distributed normally, uniformly or geometrically. Then, we ask the following question: *Under the conditions that make such a strategy profile welfare-maximizing in each of these three cases, is this profile also an equilibrium?* In each of these three cases, our answer to this question is *"no"*: firms always have a profitable deviation when offering a submenu with three options. Therefore, assuming that the regulator is right in their implicit belief that cognitive costs in the consumer population are such that the social optimum entails firms offering four tariffs, the surprising conclusion of our theoretical analysis in this environment suggests that, for firms to arrive at such a symmetric social optimum through competition, the regulator would need to impose a *lower* bound of four tariffs instead of an upper such bound.

2 The Model

2.1 General Setup

We consider a market with two firms and let $X := \{x_1, x_2, \ldots, x_k\}$ be a finite set of $k \geq 3$ products that can be sold by either of them. Although we do not impose this structure explicitly, we think of each product x_i as being multi-attribute (relevant attributes could be the product's brand name, price, quality, color etc.). As such, we implicitly assume that it is cognitively costly for consumers to make comparisons between the various products. The set $\mathcal{M}$ denotes the collection of all non-empty subsets of X. An element D of $\mathcal{M}$ is a *menu*. Firms are assumed to engage in simultaneous, one-shot competition in menu design, where each firm's pure strategy is a menu in $\mathcal{M}$. Our analysis focuses on pure strategies only. A generic strategy profile is denoted by (A, B), where A and B are the menus offered by the first and second firm, respectively. While mixed-strategy equilibria are guaranteed to exist in our finite game, the existence of pure-strategy equilibria in our model is not immediate. We will show constructively, however, that such equilibria exist.

We abstract from the firms' pricing decisions by associating each product x_i with an exogenous markup $w_i > 0$ that is common across firms. This assumption makes the analysis tractable and allows us to focus on the pure effect that overloaded consumers have on the equilibrium market outcomes. It also allows us to be agnostic on whether complexity or prices are more important for consumers in those borderline cases that would have otherwise arisen where the first menu is marginally less complex and, at the same time, features marginally more expensive jointly offered products. Nevertheless, we acknowledge that our assumption is indeed restrictive and that an extension of our model in the direction of allowing for pricing as well as menu-design decisions would indeed make it more realistic once relevant data becomes available that can provide some guidance in this regard.

We assume that there is a unit mass of consumers who consider each product in X to be desirable and who are not currently endowed with a product from this set. Consistent with these assumptions, we interpret their outside option as a non-market alternative that is inferior to every product in X. However, our consumers are also potentially overloaded in the sense that they face cognitive, time or other types of constraints that may render them unable/unwilling to process menus that exceed a certain complexity threshold. In line with the relevant empirical evidence that was discussed in the introduction, we further assume that the complexity of a menu coincides with its cardinality.

More specifically, we assume that at every strategy profile (A, B), each consumer first prescans each of these two menus sequentially to determine whether it is complex relative to his idiosyncratic complexity threshold or not, and discards any menu(s) that exceed this threshold. For each menu $D \in \{A, B\}$ that he does not find complex, he inspects the products contained in it and identifies his utility-maximizing option, which will be denoted $x^*(D)$. If he finds only one menu D to be non-complex, the consumer purchases $x^*(D)$. If he finds both menus A and B to be non-complex, he compares the best two alternatives $x^*(A)$ and $x^*(B)$ and buys his most preferred one (with ties broken uniform randomly). Finally, if he finds both menus to be complex, the consumer buys nothing.[5]

To prevent confusion, we elaborate with an example that emphasizes the sequential nature of the above decision process. Consider a consumer who is about to relocate and move to a brand new house that he has just bought. Before doing so he must choose which phone & broadband package to buy from one of the two providers that are available in his region

[5]The part of this choice procedure that concerns within-menu decisions is a special case of the "overload-constrained rationality" model that is proposed and axiomatically characterized in Gerasimou (2015). In that model, the consumer behaves as if he had a utility function $u : X \to \mathbb{R}$, a complexity function $\psi : \mathcal{M} \to \mathbb{R}$ that is monotonically increasing with respect to set inclusion, a complexity threshold n (an integer) and a choice correspondence C on $\mathcal{M}$ such that

$$C(D) = \begin{cases} \arg\max_{x \in D} u(x), & \text{if } \psi(D) \leq n \\ \emptyset, & \text{if } \psi(D) > n \end{cases}$$

where, being a choice correspondence on $\mathcal{M}$, C satisfies $C(D) \subseteq D$ for all D and hence also allows for $C(D) = \emptyset$. In our model the complexity function ψ coincides with the cardinality function $|\cdot|$ on $\mathcal{M}$. Complexity measures that are monotonic with respect to set inclusion are also proposed in Tyson (2008) and Frick (2016).

(both of which were previously unknown to him). Suppose the consumer wants to spend no more than five minutes reviewing the products that are available on each provider's website. Suppose also that he spends about a minute to read through each product, and needs an additional minute to find his most preferred one. This translates into a menu-cardinality threshold of four products. Suppose now that both providers offer four-product menus. The consumer can be thought of as pre-scanning and then finding his utility-maximizing option from the first provider in the morning, and as doing the same with respect to the second provider's menu in the evening. Crucially, *even though he is able to process each menu separately, at no point is the consumer assumed to be simultaneously faced with all products that are available in the market.* Therefore, the decision process that we have imposed is internally consistent.[6]

Consumers in our model are assumed to be heterogeneous both in terms of their preferences and in terms of their overload characteristics. To capture preference heterogeneity we assume that the probability of some product x_i being the consumer's most preferred alternative in menu A is given by

$$p_A(x_i) := \frac{1}{|A|}.$$

That is, preferences over the elements in X are uniformly distributed in the consumer population. This assumption too is made in order to keep the analysis simple and also to make it easier to identify the effects of choice overload on the menus offered by firms in equilibrium. One way of motivating it, however, is by thinking of firms as being uncertain about consumers' preferences and of expecting the demand for each product to be the same because of a balancing effect between a product's desirability on the one hand and its affordability on the other. This structure also suggests that, in the absence of overload constraints, the choice process of the average consumer in the population coincides with a special case of the well-known and widely applied Luce (1959) model.

We now turn to the specification of the consumers' overload characteristics. As was outlined above, we assume that consumers have generally distinct overload menu-size thresholds. To capture this heterogeneity, for any strictly positive integer h we let $q(h)$ denote the proportion of consumers who are *not* overloaded at menus with h or fewer elements. This makes q a cumulative density function (cdf) which we will refer to as the *overload cdf*. The support of q is a set $\{1, \ldots, k+n\}$, where n is an integer that may be weakly positive or negative. Given some menu A, we will often abuse notation slightly by writing $q(A) \equiv q(|A|)$. We assume that no consumer is overloaded in menus with just one option, which translates into $q(1) = 1$. Moreover, since q is a cdf, it also holds that $q(A) = q(B)$ whenever $|A| = |B|$ and $q(A) \leq q(B)$ whenever $|A| > |B|$. Given our assumptions on the support of q, the model encompasses all possible cases that lie between the one extreme where no consumer is ever overloaded at any menu and the other extreme where all consumers are overloaded at all menus with more than one option.

[6]We interpret the consumer's complexity threshold at the individual menu level as having been generated by a forward-looking reasoning whereby the consumer correctly anticipates that he will eventually be called to pre-scan and possibly fully process *two* menus with at least one option. This means, for example, that if both firms offer singleton menus, the consumer will be able to sequentially process each item in those menus and make a choice between them even if his complexity threshold at the individual-menu level suggests that he is overloaded when there are two products. Under this interpretation, such a situation would have arisen if the consumer's complexity threshold in the sequential processing of menus was such that the consumer could not consider a *total* of three or more market alternatives.

2.2 Payoffs

Given some menu D, we let $I_D := \{i : x_i \in D\}$ denote the index set of the products in this menu. When the first firm offers menu A and its opponent offers B we let its *baseline payoff* (i.e. the payoff corresponding to the case where no consumer is overloaded at either of these menus) be defined by

$$R_1(A,B) \;=\; \sum_{i \in I_{A \setminus B}} p_{A \cup B}(x_i) w_i + \frac{1}{2} \sum_{j \in I_{A \cap B}} p_{A \cup B}(x_j) w_j \tag{1}$$

Each product x_i that is offered by the first firm only is associated with an expected payoff that is given by its mark-up w_i multiplied by the probability of that product being chosen conditional on menus A and B being available in the market. If x_i is offered by both firms, then the above expected payoff is multiplied by $\frac{1}{2}$ to reflect the assumption that ties are broken uniform randomly. The second firm's baseline payoff is defined symmetrically.

We are now in position to introduce the actual payoff function which also accounts for the possibility that some consumers may be overloaded. When their overload thresholds are distributed according to some cdf q, then the first firm's payoff at profile (A, B) is given by

$$\pi_1(A,B) = \begin{cases} q(A) \cdot R_1(A,B), & \text{if } |A| \geq |B| \\ \\ q(B) \cdot R_1(A,B) + [q(A) - q(B)] \cdot \sum_{i \in I_A} p_A(x_i) w_i, & \text{if } |B| > |A| \end{cases} \tag{2}$$

Consider first the case where the menu A offered by the first firm is weakly more complex than menu B. In this case, the fraction $1 - q(A)$ of consumers who are overloaded at A will discard this menu. At the same time, a fraction $q(A)$ of them will consider A and identify its best alternative $x^*(A)$. These consumers will then compare $x^*(A)$ with the best alternative $x^*(B)$ in B (which is weakly less complex than A) and buy their most preferred of the two, subject to the tie-breaking rule. Hence, firm 1 gets its baseline payoff $R_1(A, B)$ scaled down by the fraction $q(A)$ of consumers that are not overloaded at A.

Now suppose menu B is strictly more complex than A. As in the case above, only a fraction $q(A)$ of consumers will consider A. However, unlike the above case, firm 1 is now able to attract a sub-fraction $q(A) - q(B) \geq 0$ of consumers who are overloaded at B but not at A. For these consumers the entire market consists of the products that are available in A. Hence, they will buy their most preferred product $x^*(A)$ offered by firm 1 regardless of what is offered by firm 2. In addition, for these consumers the relevant choice probabilities of a product x_i in A is $\frac{1}{|A|}$ and *not* $\frac{1}{|A \cup B|}$. On the other hand, consumers who are not overloaded at either A or B will consider both menus and hence follow the procedure described above. Therefore, the expected payoff that is derived from these consumers is simply $q(B) \cdot R_1(A, B)$. Again, the second firm's payoff is defined symmetrically.

The payoff function introduced above captures the key tradeoff between variety and complexity that firms are facing when consumers are potentially overloaded, which is the main motivating force for this paper. Given that consumers' preferences are heterogeneous, firms have an incentive to include in the offered menu as many products as possible in order to maximise the probability that the products contained in their menu are the consumers' most preferred ones. However, by increasing the size of the offered menu, firms also risk overloading consumers and hence losing them to their rival or driving them out of the market altogether.

2.3 Consumer Welfare

Since our model abstracts from prices, traditional welfare measures such as consumer surplus are not applicable in our setting. A key ingredient in our model, however, is that some overloaded consumers may ultimately choose none of the products offered in the market even though, by assumption, choosing something would have been objectively better than choosing nothing. This fact motivates a welfare measure that captures the complexity-variety tradeoff that arises from the consumers' point of view. One such measure could be defined by the function $W : \mathcal{M} \times \mathcal{M} \to \mathbb{R}$ that associates each strategy profile (A, B) with the *weighted-cardinality* welfare value

$$W(A,B) := \begin{cases} q(A) \cdot |A \cup B| + [q(B) - q(A)] \cdot |B|, & \text{if } |A| > |B| \\ q(A) \cdot |A \cup B|, & \text{if } |A| = |B| \\ q(B) \cdot |A \cup B| + [q(A) - q(B)] \cdot |A|, & \text{if } |A| < |B| \end{cases} \tag{3}$$

This welfare index can be motivated intuitively in the following way. If A is more complex than B, then those consumers who are able to consider A will also consider B and hence will benefit from being able to choose after processing all products that are available in the market. On the other hand, those consumers who are overloaded at A but not at B will only benefit from the products that are contained in the latter menu, while consumers who are also overloaded at B will not benefit from either of the two available menus. The welfare measure that is proposed in (3) is defined in a way that accounts for these forces. Specifically, the weighted-cardinality index W weighs the total number of products and the cardinality of the least complex menu by the fractions of consumers who will consider the former and the latter, respectively.

We note that the welfare index is bounded above by k, and this maximum is achieved when $q(X) = 1$ and $A \cup B = X$, i.e. in the case where no consumer is overloaded and the two firms together offer all possible products. We also note that W is bounded below by 0 and attains this minumum when $q(A) = q(B) = 0$, i.e. when all consumers are overloaded at both menus.

If all products are *equi-profitable* in the sense that their markups coincide and equal some $w > 0$, it follows from (2) that

$$\pi_i(A, A) = \frac{q(A)}{2} \cdot w \;\; \text{for } i = 1, 2.$$

Since $q(1) = 1$ and q is weakly decreasing, this implies that, within the class of symmetric profiles (A, A), profits are always maximized when $|A| = 1$. On the other hand, consumer welfare in this class of symmetric profiles is generally not maximized when firms offer the same singleton menu. We will come back to this point later.

3 Analysis

3.1 Maximum-Variety Equilibrium

We start by considering the special case where no consumer is overloaded. The *maximum-variety* symmetric profile (X, X) where each firm offers all products is the benchmark in this environment.

Proposition 1.
If no consumer is overloaded at X, then the maximum-variety profile (X, X) is a strict but generally not unique equilibrium. If, in addition, all products are equi-profitable, then (X, X) is an equilibrium in strictly dominant strategies.

All proofs are in the Appendix. Conditional on a firm's opponent offering menu X, its *unique* best response is to also offer X because the choice probability of each product is $\frac{1}{k}$ regardless of which menu $A \subseteq X$ the firm chooses to offer, while the firm cannot attract any overloaded consumers by offering a less complex menu. By contrast, doing so would result in not receiving an expected payoff equal to $\frac{1}{2k} \cdot w_i > 0$ for each product x_i in X that it does not offer. We note that this conclusion is robust to the presence of some overloaded consumers. In this case, the minimum fraction of non-overloaded consumers at X, i.e. $q(X)$, is increasing in the cardinality of X.

The more interesting statement in the first part of the proposition is that the maximum-variety equilibrium is not unique in general. To illustrate this with an example, suppose $X = \{x_1, x_2, x_3, x_4\}$ and let $w_1 = w_2 = w_3 = 1$ and $w_4 = \frac{1}{10}$. Suppose also that $q(X) = 1$. Let A_{ijk} denote the menu consisting of products x_i, x_j and x_k. It follows from (2) that $\pi_1(A_{123}, A_{123}) = \frac{1}{2}$, $\pi_1(X, A_{123}) = \frac{32}{80}$, $\pi_1(A_{12}, A_{123}) = \frac{1}{3}$ etc. Therefore, (A_{123}, A_{123}) is an equilibrium. More generally, non-maximum-variety equilibria also exist even in the absence of overloaded consumers whenever the least profitable products have markups that are sufficiently lower than the rest, so that when a firm unilaterally deviates by introducing them it is actually hurt due to the associated decrease in all products' choice probabilities outweighing the gains from being the only firm that offers these low-markup products.

Coming to the case where all product markups are equal to a common $w > 0$, offering X becomes a strictly dominant strategy because, as it turns out, even though the introduction of more products in the market lowers all choice probabilities and hence the shared component of the firm's expected payoff, this loss is more than offset by the firm's expected payoff from the products that it offers uniquely. Interestingly, therefore, this special case of our model provides another example where Hotelling's (1929) principle of *minimum product differentiation* applies. Finally, since $W(X, X) = k$, this *maximum-variety/minimum-differentiation* equilibrium (X, X) also corresponds to the global maximizer of consumer welfare.

3.2 Minimum-Variety Equilibria

Since 2010, Apple and Samsung have been the two main competitors at the high end of the global market for mobile phones.At any given time period, both firms have typically been offering a single "flagship" device from their *iPhone* and *Galaxy S* series, respectively. It is probably fair to say that the new features which are introduced every time the latest version of these products is launched by either firm are not always appreciable by prospective buyers without some time spent on the task of exploring these features. It is therefore conceivable that at least part of the reason why these firms do not include more high-end devices in their product lines is to prevent potential buyers from becoming overloaded.

The main idea suggested by this example can be formalized in the context of our model. In particular, if it is assumed that the products' distinct features are sufficiently important for the average consumer to consider the two of them as *imperfect* substitutes, this example would correspond to the case of an asymmetric equilibrium (A_1, B_1) where $A_1 = \{x_i\}$ and $B_1 = \{x_j\}$, $i \neq j$. On the other hand, if it is assumed that these products are perceived as *perfect* substitutes, one could interpret this as a *minimum-variety/minimum differentiation* symmetric equilibrium (A_1, A_1) in which both firms offer the same product.

Under the assumption that the products' profit margins are the same, our next result provides a characterization of such minimum-variety equilibria, and also shows that, within this same framework, single-product asymmetric equilibria also exist.

Proposition 2.
Suppose all products are equi-profitable. The following are equivalent:

1. (A_1, A_1) *is an equilibrium.*
2. *The overload cdf* q *first-order stochastically dominates the cdf* $\widehat{q}$ *defined by* $\widehat{q}(l) := \frac{l}{2l-1}$.
3. (A_h, A_h) *is an equilibrium if and only if* $h = 1$.

Moreover, whenever the above are true, a profile (A_1, B_1) *with* $A_1 \neq B_1$ *is also an equilibrium.*

Proposition 2 shows that for minimum-variety equilibria to exist in an equal-markups environment it is necessary and sufficient that the fraction of consumers who are not overloaded at menus of size $l > 1$ is bounded above by $\frac{l}{2l-1}$. As is also shown in Fig. 1, the target cdf $\widehat{q}$ that is pinned down in Proposition 2 is strictly convex. Thus, it features a very sharp drop in the fraction of non-overloaded consumers when the menu size increases from one to two, with this drop decreasing in subsequent menu-size increments. An example of an overload cdf q that does satisfy this condition is the *geometric cdf* defined by $q(h) := (1-p)^{h-1}$, provided that the parameter $p \geq 0.34$ (see Fig. 1). Interestingly, the value of the target cdf $\widehat{q}$ converges to $\frac{1}{2}$ and not to zero. This suggests, in particular, that, as far as the occurrence of minimum-variety equilibria are concerned, whether more than half of the consumers become overloaded as the menu size increases is irrelevant. The third statement also shows that for minimum-variety equilibria to exist it is necessary and sufficient that symmetric equilibria of higher menu sizes do not exist.

Figure 1: Overload cdf for minimum-variety equilibria

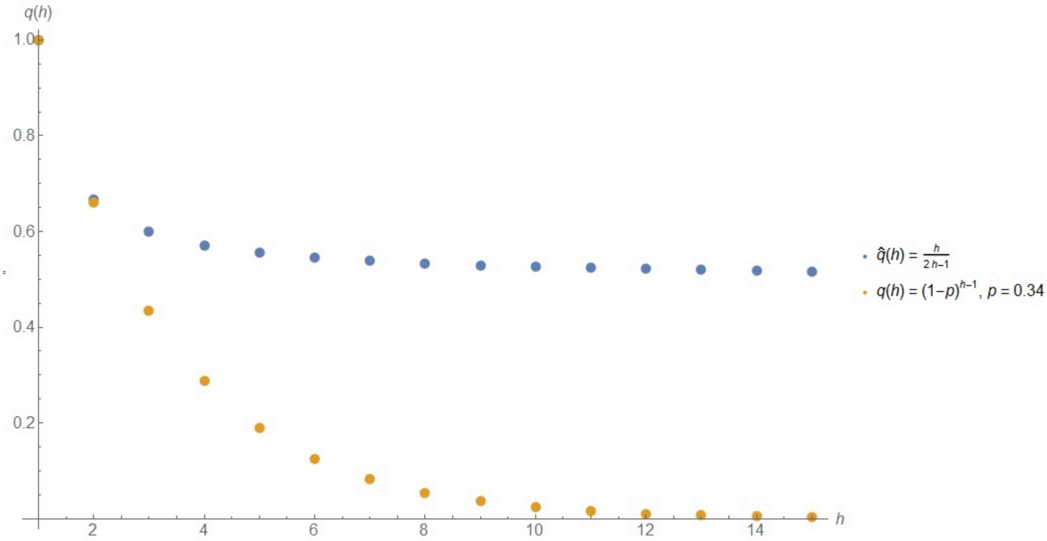

Importantly, Proposition 2 also clarifies that in those environments where minimum-variety equilibria exist, *full product differentiation* profiles (A_1, B_1) of the kind that may be suitable for thinking about the above Apple-Samsung rivalry example are also equilibria that actually *Pareto-dominate* the minimum-variety ones. Indeed, equilibrium profits are the same for the two firms and coincide with half the value of the common markup.[7] However, consumer welfare is twice as high when two distinct products are offered in the market given that, by assumption, both will be considered by *all* consumers.

3.3 Product Differentiation

Our next result identifies a sufficient condition for a full product differentiation equilibrium (A, B) where A and B are disjoint menus to be impossible under *all* overload cumulative distributions.

Proposition 3.
Suppose w_1 and w_2 are the highest and second-highest markups, respectively. If $w_1 \geq 2\frac{k-1}{k}w_2$, then there is no overload cdf q under which a full product differentiation equilibrium exists.

When the condition in the above statement is satisfied, both firms will in fact choose to offer a menu that necessarily includes the most profitable alternative, as all potential gains from full product differentiation are eliminated in this case and the shape of the overload cdf becomes irrelevant. Notably, when there is a total of only three products that firms can choose from, the highest markup need only be one third times higher than the second-highest markup for the above to happen, while this margin converges to twice the second-highest markup as the total number of products becomes large. Intuitively, when there are few products that firms can choose from, the products' choice probabilities are bounded-below by a relatively high margin (e.g. $\frac{1}{3}$ when $k = 3$). This in turn provides more leeway for the most

[7] The reasons for this, however, are different across the two cases. For profile (A_1, A_1) the $\frac{w}{2}$ payoff is due to the tie-breaking rule and the fact that no consumer is overloaded at singletons, whereas in the case of (A_1, B_1) it is due to the latter fact together with the existence of two products in the market, each of which is offered uniquely by each firm.

profitable product to stay relatively close (in terms of profitability) to the next most profitable one without losing its ability to attract both sellers to offering it in equilibrium. As the two firms' product-differentiation possibilities increase when k becomes large, for the most profitable product x_1 to continue to be offered by both of them in *every* equilibrium, its markup must be sufficiently high to offset the lower payoff that is associated with both firms offering this product when many other products can also be made available by them.

Applying the contrapositive of this result in the context of the Apple-Samsung rivalry example of the previous subsection, our model predicts that the equilibrium (A_1, B_1), which features the special kind of full product differentiation that corresponds to the case where the two firms' flagship products are considered imperfect substitutes, can only arise if their markups are sufficiently close to each other in the above sense. However, while this degree of closeness in profitability is a necessary condition for such a profile to be an equilibrium, whether it is also sufficient or not also depends on the shape of the overload cumulative distribution.

3.4 Homogeneous Overload and Heterogeneous Profitability

A question that arises in the context of our model concerns the special case where consumers are homogeneous in terms of their overload thresholds and heterogeneous in terms of their preferences. Under the equivalent interpretation of there being a single consumer in the market over whose preferences and overload threshold the firms are uncertain, this special case corresponds to firms facing uncertainty over the consumer's preferences only. We will now focus on this special case while also allowing for the products' markups to differ. Formally, we will say that *overload is homogeneous at threshold l* if

$$q(h) = \begin{cases} 1, & \text{if } h \leq l \\ 0, & \text{if } h > l \end{cases} \tag{4}$$

which simply says that all consumers have the same overload threshold.

Proposition 4.
Suppose overload is homogeneous at some menu size $l > 1$. Then:

1. A profile (A_l, A_l) is an equilibrium if A_l consists of l most profitable products.

2. The profile (A_l, A_l) is the unique symmetric equilibrium if $w_1 > \ldots > w_l$ and

$$w_{n+1} > \frac{1}{2n} \sum_{i=1}^{n} w_i \quad \textit{for all } n < l. \tag{5}$$

In both cases, the converse implication does not hold in general.

Although it is intuitive that both firms offering the same menu with l most profitable products should be an equilibrium, this conclusion is not immediate. In particular, while it is obvious that no firm will deviate by offering a menu that is contained in A_l or one that has more products than A_l, it is not a priori clear that firms cannot deviate by offering another menu with l products where the l-th most profitable one is replaced by the next most profitable one. One might suspect, for instance, that if w_l and w_{l+1} are equal, then the deviating firm's new expected profit from being the only one that offers x_{l+1} exceeds the expected loss that results from the decrease in all products' choice probabilities from $\frac{1}{l}$ to $\frac{1}{l+1}$. The proof shows precisely that such a deviation is never profitable.

Perhaps more surprising, however, is the fact that symmetric equilibria other than (A_l, A_l) also exist when overload is homogeneous at l. Indeed, the four-product example that was introduced earlier and which shows that (X, X) is generally not the unique equilibrium when no consumer is overloaded in menus with four elements is a case in point. The intuition is, of course, the same: If the least profitable products within the l most profitable ones have a sufficiently low markup, then a firm that introduces them unilaterally receives a lower payoff because the decrease in expected profits from the reduction in choice probabilities exceeds the gain associated with the introduction of new products that will be considered by all consumers.

Condition (5) requires that, for each $i = 1, \ldots, l-1$, the profitability of product x_{i+1} is strictly higher than half the average profitability of the products that are ranked above it on this dimension. Therefore, it restricts the top l markups to be strictly ordered and, in a sense, to also be *non-linearly structured*. In the four-product example where $w_1 = 6$, $w_2 = 5$, $w_3 = 4$ and $w_4 = 3$ it is easy to see that (5) is satisfied. However, if one more product is added whose markup is $w_5 = 2$, then the condition is no longer satisfied for all products. By contrast, if the markups decrease in the "strictly convex" manner whereby $w_1 = 1.9$, $w_2 = 1.5$, $w_3 = 1.3$, $w_4 = 1.15$ and $w_5 = 1.05$, then (5) holds for all markups. We note, however, that (5) also holds in general when all markups are the same. Hence, the additional assumption in the second part of the proposition which which requires the first l markups to be strictly ordered is not redundant.

To see that (5) is generally not necessary for (A_l, A_l) to be the unique symmetric equilibrium, suppose X comprises four products, $l = 4$ and $w_1 = 6$, $w_2 = w_3 = w_4 = 5$. Notice that $w_4 < \frac{1}{2}\frac{w_1+w_2+w_3}{2}$ and therefore (5) fails. From above, (X, X) is an equilibrium. Moreover, it can be easily verified that $\pi_1(A_{123}, A_{123}) < \pi_1(X, A_{123})$, $\pi_1(A_{12}, A_{12}) < \pi_1(A_{123}, A_{12})$ and $\pi_1(A_1, A_1) < \pi_1(A_1, A_{12})$ in this example.

3.5 Policy Case Study: The Nudging Effects of a Cap on the Number of Products

As a policy case study and application of our model we investigate theoretically the recent decision of the UK Office for Gas and Electricity Markets (Ofgem, 2013) to cap the number of energy tariffs that UK providers could offer to *four*, which in large part was based on the argument that *"consumers should face fewer tariff choices to make comparisons between them easier"* (p. 12). This particular *"nudge"*[8] was very recently criticized by the UK Competition & Markets Authority (CMA, 2016) on the grounds that it has led suppliers to withdraw some of the tariffs and discounts they used to offer, *"which may have made some customers worse off"* (p. 41). Under assumptions that we state below, our theoretical analysis complements this

[8] We use the term that was proposed in Sunstein and Thaler (2008) and which has become popular for this kind of policy interventions.

criticism on a different basis.

We first examine the conditions under which it is welfare-maximizing for firms to offer the same menu with four options when overload in the consumer population is distributed *normally*, *uniformly* or *geometrically* (see Fig. 2). We then ask the following question: *Under the conditions that make such a strategy profile welfare-maximizing in each of these three cases, is this profile also an equilibrium?* In other words, could this supposed welfare optimum be achieved as a free market outcome? In all three cases, our answer to this question is *"no"*: firms always have a profitable deviation when offering a submenu with three options.

Proposition 5.
If all products are equi-profitable and overload is distributed uniformly or geometrically and in such a way that a profile (A_4, A_4) *is welfare-maximizing in each of these two cases, then a firm deviates profitably by offering a menu that is a proper subset of* A_4. *This conclusion remains valid when overload is normally distributed and* $(\mu, \sigma^2) \in [4.7, 5] \times [0.8, 1]$.[9]

Figure 2: Welfare maximized at (A_4, A_4) under geometrically, uniformly and normally distributed overload

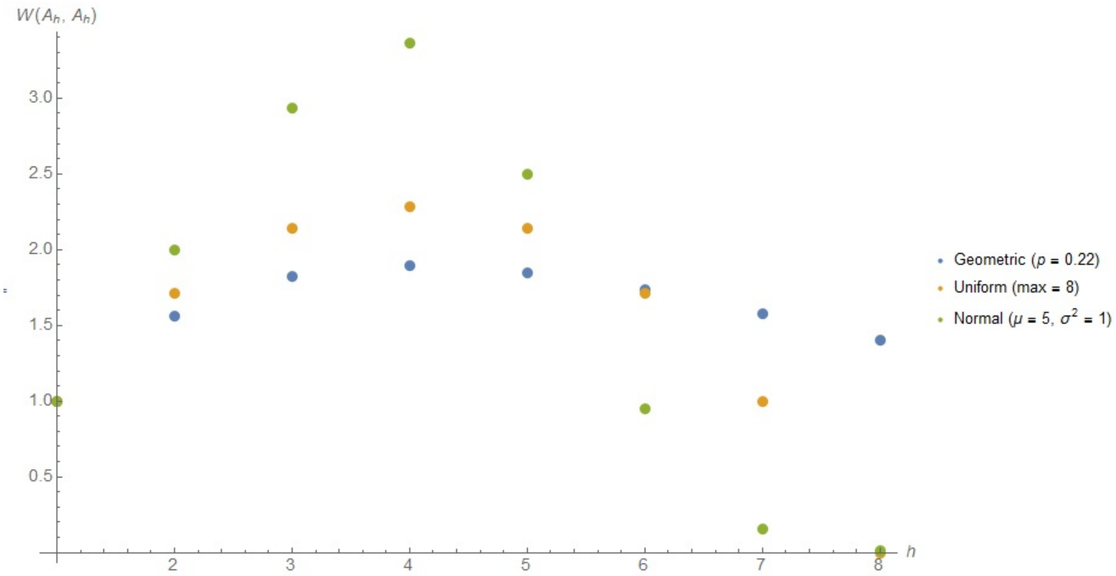

Taken together, our proposed model and welfare index suggest that if the regulating authority in this case are correct in their implicit belief that cognitive costs in the consumer population are such that the social optimum entails firms offering four tariffs,[10] then the regulator would need to impose a *lower bound of four tariffs* instead of an upper such bound in order for firms to arrive at a symmetric social optimum through regulated competition. In other words, in each of the three cases that we consider the welfare-optimal symmetric profile cannot be achieved as an equilibrium due to the overload-driven menu-size reduction pressures that are exercised on each firm. The relative robustness of this conclusion is perhaps

[9] As we show in the proof, this is possible when $p \in [0.2, 0.25]$ if overload is geometrically distributed, and when $h = 8$ or $h = 9$ when it is uniformly distributed. The statement of the proposition also makes clear what are the mean-variance combinations under which the conclusion is also true for normally distributed overload.

[10] Although details were not provided on how the number "four" was arrived at, we note that this number is the average of what cognitive psychologists currently consider to be the humans' short-memory capacity when the latter is measured in terms of chunks of letter/number sequences (see e.g. Cowan, 2001). This literature dates back to Miller (1956) and the "law" that was named after him. Influential arguments for the cognitive limitations that economic agents are faced with go back to Simon (1957).

surprising given that the geometric, uniform and normal cdfs are structurally very different in that they feature strictly convex, linear and strictly concave shapes in the relevant part of the domain, respectively.[11]

4 Related Literature

The broad literature which our paper belongs to attempts to analyze the effects of various forms of consumer bounded rationality on the outcome of firm competition. Examples include consumer *loss aversion* (Heidhues and Köszegi, 2008; Karle and Peitz, 2014; Carbajal and Ely, 2016), *inattention* (Eliaz and Spiegler, 2011; de Clippel, Eliaz, and Rozen, 2014; Bordalo, Gennaioli, and Shleifer, 2016; Manzini and Mariotti, 2016), *bounded-rational expectations* (Gabaix and Laibson, 2006; Spiegler, 2006) and *comparison difficulty* (Piccione and Spiegler, 2012; Bachi and Spiegler, 2014; Papi, 2014, 2015; Fisher and Plan, 2015).[12] In sharp contrast to many of these models where firms add complexity/obfuscation to their price structures and hence manage to sustain positive markups, our analysis highlights the potentially beneficial effect that *menu simplicity* can have in the firms' efforts to increase their market share.

More closely related to our model is Kamenica's (2008), where consumer demand for a monopolist's menu may be decreasing in the number of products contained in it because consumers who are uninformed about their preferences make the "contextual inference" that a smaller menu includes the most popular alternatives, and hence choose one of these due to the higher probability that it will be a good match for them. Our model is different in that small menus may be offered in the market as the equilibrium outcome of duopolistic competition, and with the latter taking place in the presence of cognitively constrained consumers who are fully aware of their preferences. In addition, Bachi and Spiegler (2014) propose a class of duopolistic models where consumers are presented with two-attribute products in Euclidean space and experience difficulties in making trade-offs across these attributes. Their "opt-in" model, in particular, deals with the case where consumers actually defer choice due to such comparison difficulties, for example when none of the feasible alternatives is Pareto dominant in a menu.[13] Although choice overload as analyzed here is not a source of deferral for consumers in that model, that paper studies the effects that indecisiveness-driven deferral has on market outcomes.

Anderson et al (1992, section 7.4) study the equilibrium properties of an oligopolistic market for differentiated products in which consumer demand is determined by the multinomial logit model when an outside option is present. With demand solely determined by this special class of random utility models, the outside option there is chosen only when it is perceived as more attractive than the actual available products.[14] By contrast, consumer demand in our setting is not determined by some random utility model and the outside option

[11]Specifically, the normal cdf in this setting is strictly convex in the relevant range of menus with up to four products, whereas the uniform and geometric cumulative distributions are strictly convex and linear, respectively, throughout their support.

[12]A textbook treatment of this literature is provided in Spiegler (2011), while Spiegler (2015) surveys and synthesizes some more recent developments and trends.

[13]For revealed-preference foundations of choice deferral is caused by the inability of consumers to find a partially or totally dominant option due to preference incompleteness/consumer indecisiveness the reader is referred to Gerasimou (2015, 2016).

[14]See Gerasimou (2015) for the revealed-preference implications of this source of choice deferral.

is chosen only when consumers are overloaded.

Klemperer and Padilla (1997) propose a model in which rational consumers prefer to buy from a single seller due to the presence of shopping costs. In that model firms have an incentive to increase the variety of the offered menu in order to decrease its opponents' market shares. The authors show that this can lead to a socially undesirable outcome, as the reduction in the opponents' profits may outweigh the increase in consumers' surplus. Therefore, unlike the welfare conclusions that stem from our analysis, in that model the occurrence of large equilibrium menus can be socially inefficient not because of the socially sub-optimal number of consumers who can benefit from them but due to possible rival-foreclosing effect that is implied by the consumers' tendency to buy from few sellers.

In addition to the overload-constrained maximization model that is analyzed in Gerasimou (2015) and which our consumers have been assumed to conform with, other decision-theoretic models in which smaller menus are in some sense better for the decision maker are proposed in Billot and Thisse (1999), Mullainathan (2002), Sarver (2008), Ortoleva (2013), Buturak and Evren (2015), Frick (2016) and Dean et al (2016). The reasons for such behavior vary across these models and include regret as well as cognitive costs/attention constraints. Finally, we refer the reader to Chatterjee and Sabourian (2000) for a multi-person bargaining model where players are also assumed to be facing computational costs which, in that context, limit the degree to which their strategies can depend on the game's history.

5 Concluding Remarks

When oligopolistically competitive firms sell their products to consumers who are potentially choice-overloaded in the sense that they defer/avoid choice when they see too many products in a menu, such firms are faced with a novel strategic trade-off. In the context of this trade-off, a firm that offers many products appeals to many consumers' tastes, but at the same time it may also overload these consumers and hence either lose them to its rival or drive them out of the market altogether. This paper proposes the first model in the literature of behavioral industrial organization that aims to provide a framework for thinking in a simple and yet general and precise way about this strategic trade-off.

From a policy point of view, the presence of overloaded consumers calls for re-thinking about conventional measures of welfare such as consumer surplus, as these fail to capture the possibility that certain market outcomes may be harmful not because of high equilibrium prices or markups, but because a potentially significant fraction of the consumer population may not actually benefit from the products that are made available through the equilibrium market outcome due to their cognitive constraints. Our paper proposes a simple welfare measure that may be useful in this regard. Using this measure alongside our market model enabled us to analyze and, perhaps surprisingly, to critically evaluate a recent UK policy intervention that capped the number of products in the energy market in order to increase consumer welfare.

Although our model assumes an "opt-in" formulation and hence that consumers are not already endowed with a default market alternative when they go to the marketplace, a potentially interesting extension would feature an "opt-out" formulation that would allow for such default alternatives to exist and influence the consumers' (and hence the firms') decisions. This extension, in particular, would allow one to study equilibria where consumers do or do not switch away from their default market options and, as such, would be particularly suitable for the theoretical analysis of relevant policy questions.

References

ANDERSON, S. P., A. DE PALMA, AND J.-F. THISSE (1992): *Discrete Choice Theory of Product Differentiation*. Cambridge, MA: MIT Press.

BACHI, B., AND R. SPIEGLER (2014): "Buridanic Competition," *mimeo*.

BILLOT, A., AND J.-F. THISSE (1999): "A Discrete Choice Model When Context Matters," *Journal of Mathematical Psychology*, 43, 518–538.

BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2016): "Competition for Attention," *Review of Economic Studies*, 83, 481–513.

BUTURAK, G., AND O. EVREN (2015): "Choice Overload and Asymmetric Regret," *mimeo*.

CARBAJAL, J. C., AND J. C. ELY (2016): "A Model of Price Discrimination Under Loss Aversion and State-Contingent Reference Points," *Theoretical Economics*, 11, 455–485.

CHATTERJEE, K., AND H. SABOURIAN (2000): "Multiperson Bargaining and Strategic Complexity," *Econometrica*, 68, 1491–1509.

CHERNEV, A., U. BÖCKENHOLT, AND J. GOODMAN (2015): "Choice Overload: A Conceptual Review and Meta-Analysis," *Journal of Consumer Psychology*, 25, 333–358.

CMA (2016): "Energy Market Investigation - Final Report," Competition & Markets Authority, 24 June 2016, https://assets.publishing.service.gov.uk/media/576c23e4ed915d622c000087/Energy-final-report-summary.pdf.

COWAN, N. (2001): "The Magical Number 4 in Short-Term Memory: A Reconsideration of Mental Storage Capacity," *Behavioral and Brain Sciences*, 24, 87–114.

DE CLIPPEL, G., K. ELIAZ, AND K. ROZEN (2014): "Competing for Consumer Inattention," *Journal of Political Economy*, 122, 1203–1234.

DEAN, M., O. KIBRIS, AND Y. MASATLIOGLU (2016): "Limited Attention and Status Quo Bias," *Journal of Economic Theory*, forthcoming.

ELIAZ, K., AND R. SPIEGLER (2011): "Consideration Sets and Competitive Marketing," *Review of Economic Studies*, 78, 235–262.

FISHER, J., AND A. PLAN (2015): "Competition for Imperfectly Informed Consumers," *mimeo*.

FRICK, M. (2016): "Monotone Threshold Representations," *Theoretical Economics*, 11, 757–772.

GABAIX, X., AND D. LAIBSON (2006): “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets,” *Quarterly Journal of Economics*, 121, 505–540.

GERASIMOU, G. (2015): “Indecisiveness, Undesirability and Overload Revealed Through Rational Choice Deferral,” *mimeo*.

——— (2016): “Asymmetric Dominance, Deferral and Status Quo Bias in a Behavioral Model of Choice,” *Theory and Decision*, 80, 295–312.

HEIDHUES, P., AND B. KÖSZEGI (2008): “Competition and Price Variation When Consumers Are Loss Averse,” *American Economic Review*, 98, 1245–1268.

HOTELLING, H. (1929): “Stability in Competition,” *Economic Journal*, 39, 41–53.

IYENGAR, S. S., G. HUBERMAN, AND W. JIANG (2004): “How Much Choice is Too Much? Contributions to 401(k) Retirement Plans,” in *Pension Design and Structure: New Lessons from Behavioral Finance*, ed. by O. S. Mitchell, and S. P. Utkus, pp. 83–96. New York: Oxford University Press.

IYENGAR, S. S., AND E. KAMENICA (2010): “Choice Proliferation, Simplicity Seeking, and Asset Allocation,” *Journal of Public Economics*, 94, 530–539.

IYENGAR, S. S., AND M. R. LEPPER (2000): “When Choice is Demotivating: Can One Desire Too Much of a Good Thing?,” *Journal of Personality and Social Psychology*, 79, 995–1006.

KAMENICA, E. (2008): “Contextual Inference in Markets: On the Informational Content of Product Lines,” *American Economic Review*, 98, 2127–2149.

——— (2012): “Behavioral Economics and the Psychology of Incentives,” *Annual Review of Economics*, 4, 427–452.

KARLE, H., AND M. PEITZ (2014): “Competition under Consumer Loss Aversion,” *RAND Journal of Economics*, 45, 1–31.

KLEMPERER, P., AND A. J. PADILLA (1997): “Do Firms’ Product Lines Include Too Many Varieties?,” *RAND Journal of Economics*, 28, 472–488.

KUKSOV, D., AND J. M. VILLAS-BOAS (2010): “When More Alternatives Lead to Less Choice,” *Marketing Science*, 29, 507–524.

LUCE, R. D. (1959): *Individual Choice Behavior: A Theoretical Analysis*. New York: Wiley.

MANZINI, P., AND M. MARIOTTI (2016): “Competing for Attention: Is the Showiest also the Best?,” *Economic Journal*, forthcoming.

MILLER, G. A. (1956): “The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information,” *Psychological Review*, 63, 81–97.

MULLAINATHAN, S. (2002): “A Memory-Based Model of Bounded Rationality,” *Quarterly Journal of Economics*, 117, 735–774.

OFGEM (2013): "The Retail Market Review - Implementation of Simpler Tariff Choices and Clearer Information," Office for Gas & Electricity Markets, 27 August 2013, https://www.ofgem.gov.uk/sites/default/files/docs/decisions/the_retail_market_review_-_implementation_of_simpler_tariff_choices_and_clearer_information.pdf.

——— (2014): "Finding a Better Deal on Your Energy Is Getting Easier from Today," Office for Gas & Electricity Markets, 2nd January 2014, https://www.ofgem.gov.uk/press-releases/finding-better-deal-your-energy-getting-easier-today.

ORTOLEVA, P. (2013): "The Price of Flexibility: Towards a Theory of Thinking Aversion," *Journal of Economic Theory*, 148, 903–934.

PAPI, M. (2014): "Noncompensatory Consideration and Compensatory Choice: An Application to Stackelberg Competition," *Economic Theory Bulletin*, 2, 53–63.

——— (2015): "Firm Competition with Trade-Off Avoiding Consumers," *mimeo*.

PICCIONE, M., AND R. SPIEGLER (2012): "Price Competition under Limited Comparability," *Quarterly Journal of Economics*, 127, 97–135.

SARVER, T. (2008): "Anticipating Regret: Why Fewer Options May Be Better," *Econometrica*, 76, 263–305.

SIMON, H. (1957): *Models of Man*. New York: Wiley.

SPIEGLER, R. (2006): "Competition Over Agents with Boundedly Rational Expectations," *Theoretical Economics*, 1, 207–231.

——— (2011): *Bounded Rationality and Industrial Organization*. New York: Oxford University Press.

——— (2015): "Choice Complexity and Market Competition," *mimeo*.

SUNSTEIN, C. R., AND R. H. THALER (2008): *Nudge: Improving Decisions About Health, Wealth and Happiness*. New Haven: Yale University Press.

TYSON, C. J. (2008): "Cognitive Constraints, Contraction Consistency, and the Satisficing Criterion," *Journal of Economic Theory*, 138, 51–70.

VARIAN, H. R. (2014): *Intermediate Microeconomics*. New York: Norton, 9th edn.

Appendix: Proofs

We start with the following auxiliary result.

Lemma 1.
Assume that all products are equi-profitable. Let $|A| = a$, $|B| = b$, and $|A \cap B| = c$. Then:

1. *$\pi_1(A, B)$ is constant in c whenever $a = b$.*

2. $\pi_1(A,B)$ *is strictly increasing in c whenever* $a > b$*, and*

3. $\pi_1(A,B)$ *is strictly decreasing in c whenever* $b > a$*.*

Proof of Lemma 1:

Let the common markup w be normalized to $w = 1$. Assume that $a = b$. Then, the profit of a firm offering A when the opponent offers B is $\frac{a-c}{a+b-c} + \frac{1}{2}\frac{c}{a+b-c} = \frac{2a-c}{2(a+b-c)} := \bar{\pi}$. It is clear that this is constant in c when $a = b$. Next, suppose that $a > b$. Then, $\pi_1(A,B) = q(A)\bar{\pi}$. It is easy to see that this is strictly increasing in c when $a > b$.Finally, assume that $b > a$. Then, $\pi_1(A,B) = q(B)\bar{\pi} + q(A) - q(B)$. It is again straightforward to verify that this is strictly decreasing in c whenever $b > a$. ■

Proof of Proposition 1:

The proof of the first part is straightforward and omitted. To prove the second part, without loss of generality normalize the common weighted markup to 1. Consider firm 1 and suppose to the contrary that there are $A \subset X$ and $B \subseteq X$ such that $\pi_1(A,B) \geq \pi_1(X,B)$. We have

$$\begin{aligned}\pi_1(A,B) &= \sum_{i \in I_{A \setminus B}} p_{A\cup B}(x_i) + \frac{1}{2}\sum_{j \in I_{A \cap B}} p_{A\cup B}(x_j) \\ \pi_1(X,B) &= \sum_{i \in I_{X \setminus B}} p_X(x_i) + \frac{1}{2}\sum_{j \in B} p_X(x_j)\end{aligned}$$

We also have

$$\sum_{i \in I_{A \setminus B}} p_{A\cup B}(x_i) + \frac{1}{2}\sum_{j \in I_{A \cap B}} p_{A\cup B}(x_j) \geq \sum_{i \in I_{X \setminus B}} p_X(x_i) + \frac{1}{2}\sum_{j \in I_B} p_X(x_j)$$

Moreover, $p_{E\cup F}(x_i) = p_{E\cup F}(x_j) = p_{E\cup F}$ for all $E, F \in \mathcal{M}$ and all $x_i, x_j \in E \cup F$. Therefore,

$$\begin{aligned}\sum_{i \in I_{A \setminus B}} p_{A\cup B}(x_i) + \frac{1}{2}\sum_{j \in I_{A \cap B}} p_{A\cup B}(x_j) &= |A \setminus B| \cdot p_{A\cup B} + \frac{1}{2}|A \cap B| \cdot p_{A\cup B} \\ \sum_{i \in I_{X \setminus B}} p_X(x_i) + \frac{1}{2}\sum_{j \in I_B} p_X(x_j) &= |X \setminus B| \cdot p_X + \frac{1}{2}|B| \cdot p_X\end{aligned}$$

Suppose $|A| = n$, $|B| = m$ and $A \cap B = \emptyset$, so that $m + n \leq k$. We have $\pi_1(A,B) = \frac{n}{m+n}$ and $\pi_1(X,B) = \frac{k-m}{k} + \frac{1}{2}\frac{m}{k}$. Thus,

$$\pi_1(A,B) \geq \pi_1(X,B) \iff \frac{m+n}{2} \geq k$$

Since $n + m \leq k$, this is clearly false. Therefore, for no such A and B is it true that $\pi_1(A,B) \geq \pi_1(X,B)$.

It remains to be shown that $\frac{m+n}{2} \geq k$ is also false when $A \cap B \neq \emptyset$ and $A \cup B \subset X$. As before, assume $|A| = n$, $|B| = m$ and let $|A \cap B| = l$. It holds that

$$\begin{aligned}\pi_1(A,B) &= \frac{n-l}{m+n-l} + \frac{1}{2}\frac{l}{m+n-l} = \frac{2n-l}{2(m+n-l)} \\ \pi_1(X,B) &= \frac{k-m}{k} + \frac{1}{2}\frac{m}{k} = \frac{2k-m}{2k}\end{aligned}$$

We distinguish three cases.

Case (i): $n > m$. By Lemma 1, $\pi_1(A,B)$ is strictly increasing in l. Hence, it suffices to compare the two payoffs when this attains its maximum value, i.e. at $l = m$. When $l = m$, $\pi_1(A,B) \geq \pi_1(X,B)$ if and only if $n \geq k$, which is obviously false.

Case (ii): $n = m$. By Lemma 1, $\pi_1(A,B)$ is constant in l. Hence, assume without loss of generality that $l = m$. When $l = m$, $\pi_1(A,B) \geq \pi_1(X,B)$ if and only if $n \geq k$, which is false.

Case (iii): $n < m$. By Lemma 1, $\pi_1(A,B)$ is strictly decreasing in l. Hence, let $l = 1$. When $l = 1$, $\pi_1(A,B) \geq \pi_1(X,B)$ if and only if $n \geq m$, which is false.

Therefore, $\pi_1(X,B) > \pi_1(A,B)$ for all $A \neq X \neq B \in \mathcal{M}$. Since the game is symmetric, this proves that (X,X) is a strictly dominant strategy equilibrium. ■

Proof of Proposition 2.

1 ⇔ 2. As above, we normalize the weighted markups to 1. Assume that both firms offer B such that $|B| = 1$. Since $q(1) = 1$ by assumption, (B,B) is an equilibrium if and only if $\pi_1(A,B) \leq \frac{1}{2}$ for every $A \in \mathcal{M}$. Note first that $\pi_1(A,B) = \frac{1}{2} = \pi_1(B,B)$ for all $A \in \mathcal{M}$ such that $|A| = 1$. Suppose a firm deviates to $A \in \mathcal{M}$ such that $|A| \in \{2, \ldots, k\}$. By Lemma 1, profits from this deviation are maximized whenever $|A \cap B|$ is maximized. Hence, $A \supset B$. The deviating firm obtains $q(A)\left[\frac{|A|-1}{|A|} + \frac{1}{2}\frac{1}{|A|}\right]$. The deviation is not profitable if and only if $q(A)\left[\frac{2|A|-1}{2|A|}\right] \leq \frac{1}{2}$ or, equivalently, $q(A) \leq \frac{|A|}{2|A|-1}$ for any $|A| > 1$. Define the cdf $\widehat{q}$ by $\widehat{q}(A) := \frac{|A|}{2|A|-1}$ for all $A \in \mathcal{M}$. This argument therefore shows that (B,B) such that $|B| = 1$ is an equilibrium if and only if q first-order stochastically dominates $\widehat{q}$.

2 ⇔ 3. Without loss of generality, suppose $q = \widehat{q}$, i.e. $q(h) = \frac{h}{2h-1}$ for all $h = 1, \ldots, k-1$. Suppose there exists an equilibrium (A_l, A_l) with $l > 1$. Let A_{l-1} be a subset of A_l that contains $l-1$ products. It holds that

$$\pi_i(A_l, A_l) = \frac{l}{2(2l-1)} \geq \frac{l}{2l-1} \cdot \frac{1}{2} \cdot \frac{l-1}{l} + \frac{l-1}{2l-3} - \frac{l}{2l-1} = \pi_i(A_{l-1}, A_l)$$

This inequality is satisfied if and only if $8l^3 - 20l^2 + 20l - 6 \leq 0$. But this is false for every $l \geq 2$. Therefore, (A_l, A_l) is not an equilibrium for $l > 1$. Moreover, since q coincides with $\widehat{q}$, the argument in the previous part of the proof establishes that (A_l, A_l) is an equilibrium if $l = 1$.

Finally, suppose (A_1, A_1) is an equilibrium. We have $\pi_1(A_1, A_1) = \frac{w}{2}$. Now, for $A_1 \neq B_1$ it holds that $\pi_1(A_1, B_1) = \frac{w}{2}$ as well. Consider a deviation by firm 1 to a profile A, where

$A_1 \neq A \neq B_1$ and $|A| > 1$. Assume that $A \cap B_1 = \emptyset$. In view of the cdf $\widehat{q}$, it follows that, for all such A, $\pi_1(A, B_1) = \frac{|A|}{2|A|-1} \cdot \frac{|A|}{|A|+1} = \frac{|A|^2}{2|A|^2+|A|-1} < \frac{1}{2}$. Now suppose $A \cap B_1 \neq \emptyset$. If follows from Lemma 1 that in this case $\pi_1(A, B_1)$ is maximized when $A = B_1$. Therefore, (A_1, B_1) with $A_1 \neq B_1$ is also an equilibrium. ■

Proof of Proposition 3.

Assume without loss of generality that $|A| \geq |B|$, and $A \cap B = \emptyset$. Let the markups of the products in A be such that $w_1^A \geq w_2^A \geq \cdots \geq w_{|A|}^A$ and those in B such that $w_1^B \geq w_2^B \geq \cdots \geq w_{|B|}^B$. We have

$$\pi(A, B) \;\; = \;\; q(A) \cdot \left(\frac{\sum_{i \in I_A} w_i^A}{|A| + |B|} \right) \tag{6}$$

Assume first that $w_1^B > w_{|A|}^A$. Define A' as the menu of $|A|$ products that is identical to A except that the least profitable product $x_{|A|}^A$ in A is replaced by x_1^B in B. We have

$$\pi(A', B) \;\; = \;\; q(A) \cdot \left(\frac{\sum_{i \in I_A \setminus \{|A|\}} w_i}{|A| + |B| - 1} + \frac{1}{2} \frac{w_1^B}{|A| + |B| - 1} \right) \tag{7}$$

Let $K := \sum_{i \in I_A \setminus \{|A|\}} w_i^A$. Suppose that (A, B) is an equilibrium. It follows from (6) and (7) that

$$q(A) \cdot \left(\frac{K}{|A| + |B|} + \frac{w_{|A|}^A}{|A| + |B|} \right) \;\; \geq \;\; q(A) \cdot \left(\frac{K}{|A| + |B| - 1} + \frac{1}{2} \frac{w_1^B}{|A| + |B| - 1} \right) \tag{8}$$

Since $w_1^B > w_{|A|}^A$, we can write $w_1^B = \alpha \cdot w_{|A|}^A$ for some $\alpha > 1$. Substituting this back into (8) and rearranging, we get

$$w_{|A|}^A \left(|A| + |B| - 1 - \frac{\alpha}{2}|A| - \frac{\alpha}{2}|B| \right) \;\; \geq \;\; K \tag{9}$$

Observe that the assumption on how markpus in A are distributed implies that $K \geq (|A| - 1) \cdot w_{|A|}^A$. Therefore, it follows from (9) that $w_{|A|}^A \left(|A| + |B| - 1 - \frac{\alpha}{2}|A| - \frac{\alpha}{2}|B| \right) \geq (|A| - 1) \cdot w_{|A|}^A$, which is equivalent to $|B| \geq \frac{\alpha}{2}(|A| + |B|)$. Since $|A| \geq |B|$, then in the most favourable case $|A| = |B|$ implying that $1 \geq \alpha$, which leads to a contradiction.

Now assume that $w_1^B \leq w_{|A|}^A$. We have

$$\pi(B, A) \;\; = \;\; q(A) \cdot \left(\frac{\sum_{i \in I_B} w_i^B}{|A| + |B|} \right) + q(B) - q(A) \tag{10}$$

If the most profitable product in X does not belong to either A or B, then it is obvious that both firms have an incentive to deviate by replacing any product they offer with the most profitable product. Hence, assume that the most profitable product belongs to A. Hence, w_1^A is its markup. Define B' as the menu of $|B|$ products that is identical to B except that the least profitable product $x_{|B|}^B$ in B is replaced by x_1^A in A. We have

$$\pi(B', A) \quad = \quad q(A)\left(\frac{\sum_{i\in I_B\setminus\{|B|\}} w_i^B}{|A|+|B|-1} + \frac{1}{2}\frac{w_1^A}{|A|+|B|-1}\right) + q(B) - q(A) \tag{11}$$

Let $L := \sum_{i\in I_B\setminus\{|B|\}} w_i^B$. Suppose that (A, B) is an equilibrium. It follows from (10) and (11) that

$$\begin{gathered} q(A)\left(\frac{L}{|A|+|B|} + \frac{w_{|B|}^B}{|A|+|B|}\right) + q(B) - q(A) \\ \geq \\ q(A)\left(\frac{L}{|A|+|B|-1} + \frac{1}{2}\frac{w_1^A}{|A|+|B|-1}\right) + q(B) - q(A) \end{gathered} \tag{12}$$

By rearranging (12) and solving for w_1^A we get

$$w_1^A \quad \leq \quad \frac{2|A|}{|A|+|B|} w'_{|B|} + G \tag{13}$$

where $G := \frac{(|B|-1)w'_{|B|} - L}{|A|+|B|}$. Since $L \geq (|B|-1)w_{|B|-1}^B$ and $w_{|B|-1}^B \geq w_{|B|}^B$, then $G \leq 0$.

Assume first that $|A| = |B|$. Then, (13) reduces to

$$w_1^A \quad \leq \quad w'_{|B|} + G$$

Since $w_1^A > w'_{|B|}$ and $G \leq 0$, the inequality is false, which is a contradiction.

Next, assume that $|A| > |B|$. In the most favourable case it holds that $|A| = k-1$, $|B| = 1$, and $w'_{|B|} = w_3$, implying that (13) reduces to

$$w_1 \quad \leq \quad \frac{2(k-1)}{k} w_3 + G$$

Since $G \leq 0$ and $w_1 \geq \frac{2(k-1)}{k} w_2$ by assumption, this inequality is false, which leads to a contradiction. ∎

Proof of Proposition 4.

Without loss of generality, suppose $w_1 \geq w_2 \geq \ldots w_k$. Suppose (A_l, A_l) is not an equilibrium, where A_l consists of the l most profitable products. Clearly, a menu A where $A \subset A_l$ or $|A| > l$ cannot be a profitable deviation (details). To establish the result it suffices to show that there is no profitable deviation for firm 1 when A consists of the menu A_l when the l-th most profitable product is replaced by the $(l+1)$-th such product. Suppose that this is not true. Then,

$$\begin{aligned}
\pi_1(A, A_l) > \pi_1(A_l, A_l) &\iff \frac{w_{l+1}}{l+1} + \frac{1}{2(l+1)}\sum_{i=1}^{l-1} w_i > \frac{1}{2l}\sum_{i=1}^{l} w_i \\
&\iff 2lw_{l+1} + l\sum_{i=1}^{l-1} w_i > (l+1)\sum_{i=1}^{l} w_i \\
&\iff 2lw_{l+1} - lw_l > \sum_{i=1}^{l} w_i \\
&\iff 2w_{l+1} - w_l > \frac{1}{l}\sum_{i=1}^{l} w_i
\end{aligned}$$

Since $w_1 \geq \ldots \geq w_k$, the terms on the right- and left-hand side in the last inequality are bounded above and below by w_l, respectively. Hence, a contradiction obtains.

For the second part, suppose to the contrary that (A_n, A_n) is an equilibrium where $n < l$ and, as above A_n, consists of the n most profitable products. This implies that, for all j such that $n < j \leq l$,

$$\pi_1(A_n, A_n) \geq \pi_1(A_j, A_n) \iff \frac{1}{2n}\sum_{i=1}^{n} w_i \geq \frac{1}{j}\sum_{i=n+1}^{j} w_n + \frac{1}{2j}\sum_{i=1}^{n} w_i$$

This, in particular, must also hold for $j = n+1$, i.e.

$$\frac{1}{2n}\sum_{i=1}^{n} w_i \geq \frac{w_{n+1}}{n+1} + \frac{1}{2(n+1)}\sum_{i=1}^{n} w_i$$

which is true if and only if

$$\frac{1}{2n}\sum_{i=1}^{n} w_i \geq w_{n+1}.$$

However, since the top l markups are non-linearly structured, the latter is false. ■

Proof of Proposition 5.

We first establish the claim for the case where overload is geometrically distributed. At a symmetric profile (A_l, A_l) the welfare function is $W(A_l, A_l) \equiv g(l) := (1-p)^{l-1} \cdot l$, where $p \in (0, 1)$ is a parameter. If g is maximized at $l = 4$ when maximization is over the set of all positive integers, it must hold that

$$4(1-p)^3 \geq n(1-p)^{n-1} \text{ for all } n = 1, 2, \ldots$$

It follows, in particular, that $4(1-p)^3 \geq 3(1-p)^2$ and $4(1-p)^3 \geq 5c^4$, from which we eventually obtain

$$0.2 \leq p \leq 0.25.$$

Now let the common markup be $w = 1$ and let $c := 1 - p$. Firm 1's payoff at (A_4, A_4) is

$\pi_1(A_4, A_4) = \frac{1}{2} \cdot c^3$, while its payoff when it deviates to a menu $A_3 \subset A_4$ is $\pi_1(A_3, A_4) = \frac{1}{2} \cdot \frac{3}{4} \cdot c^2 + c^2 - c^3$. We have $\pi_1(A_3, A_4) > \pi_1(A_4, A_4) \Leftrightarrow c < \frac{11}{12}$, which is equivalent to $p > \frac{1}{12}$. The latter obviously holds for all $p \in [0.2, 0.25]$.

Consider now the case where overload is uniformly distributed. At a symmetric profile (A_l, A_l) the welfare function is $W(A_l, A_l) \equiv u(l) := \frac{h-l}{h-1} \cdot l$, where $h > 1$ is an integer. If u is maximized at $l = 4$ when maximization is over the set of all positive integers, it must hold that

$$\frac{4(h-1)}{h-1} \geq \frac{n(h-n)}{h-1} \quad \text{for all } n = 1, 2, \ldots$$

In particular, since this inequality must also hold when $n = 2$ and $n = 5$, it follows that

$$h \geq 8 \quad \text{and} \quad h \leq 9$$

Again, let the common markup be $w = 1$. Firm 1's payoff at (A_4, A_4) is $\pi_1(A_4, A_4) = \frac{1}{2} \cdot \frac{h-4}{h-1}$, while its payoff when it deviates to a menu $A_3 \subset A_4$ is $\pi_1(A_3, A_4) = \frac{1}{2} \cdot \frac{3}{4} \cdot \frac{h-3}{h-1} + \frac{h-3}{h-1} - \frac{h-4}{h-1} = \frac{3h-1}{8h-8}$. We have $\pi_1(A_3, A_4) > \pi_1(A_4, A_4) \Leftrightarrow h < 15$, which is true for both the above values of h.

The reader is referred to the Mathematica file in the Online Appendix for the computational details that establish the validity of the claim under normally distributed overload when $(\mu, \sigma^2) \in [4.7, 5] \times [0.8, 1]$. ■