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

December 18, 2015

Abstract

A large body of experimental work has suggested the existence of a *"choice overload"* effect in consumer decision making: Faced with large menus of choice options, decision makers often defer or avoid choice. A suggested reason for the occurrence of this effect is that the agents attempt to avoid the cognitive effort that is associated with choosing from larger 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. Assuming uniformly distributed preferences, we focus on symmetric pure-strategy equilibria under various assumptions on the overload distribution and product markups. We also propose and analyse a measure of consumer welfare that applies to this environment. Among other things, we provide conditions for *"maximum-"* and *"minimum-variety" equilibria* to be possible, whereby both firms either offer the entire set of available products or the same one product, respectively.

*We would like to thank Patrick Rey, Ariel Rubinstein, Rann Smorodinsky, David Ulph and seminar audiences at St Andrews, Technion-Israel Institute of Technology and Athens University of Economics & Business for very useful comments. Any errors are our own.

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

‡Email address: m.papi@abdn.ac.uk

1 Introduction

An important implication of the theory of rational choice 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 weakly higher on 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 alternatives available to him due to the complexities associated with choosing from such menus. This has come to be known in the literature as the *choice overload* or *overchoice* effect. The presence of this effect has been further established in decisions that involve retail products (e.g. chocolates in Berger et al 2007) as well as significantly more important options such as pension plans (Iyengar et al 2007).

At least partly in response to the large body of work in the marketing and consumer psychology literature that followed Iyengar and Lepper (2000) (which is surveyed in the excellent meta-analysis of Chernev et al 2015), formal recognition of the harmful effects that choice overload could have on consumer welfare was recently made by leading governmental authorities such as Ofgem, the UK regulator of the electricity and gas markets. In 2014 Ofgem forced energy supplies to ban complex tariffs and also to restrict the number of tariffs 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).

Once it is acknowledged that consumers can become choice-overloaded by large numbers of products that are made available to them, the implications for firm competition immediately arise. In particular, rather than offering menus with as many products as possible, firms facing 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 not to overload consumers and lose them to their competitors. This strategic trade-off lies at the very heart of the duopolistic competition model that we introduce in this paper.

Our model builds on the assumption that consumers are faced with limitations in the cognitive resources they have available for making decisions. Specifically, when given the opportunity to choose *from* a menu they are portrayed as pre-scanning it to determine how many options are included in it, and if this number exceeds their idiosyncratic complexity threshold they avoid/defer choice from that menu. By contrast, if the menu is not complex, they choose like standard utility maximizers.[1] A consequence of this model is that if a consumer's complexity threshold is higher than that of all possible menus, then he always behaves as a utility maximizer. Additionally, we assume that in decisions *between* menus consumers choose in a lexicographic manner, first comparing menus according to the number of options included in them, and only resorting to their preferences in the event of a cardinality tie.

These assumptions motivate analogous restrictions in the firms' payoff functions in the duopolistic model we later introduce. First, we model these firms as simultaneously choosing a menu to offer, given a fixed, common and finite grand set of available products. Firms are facing a population of consumers who are generally heterogeneous in their preferences as well as in their complexity/overload thresholds. Consumers may be of two types, depending on whether they defer choice altogether if both menus in the market are complex for them or whether they choose a product from the relatively less complex menu. To make the analysis more tractable, and also to isolate the pure effect of overload on the market outcome, we abstract from the firms' pricing decisions. Instead, we assume that all products come with an exogenously determined profit markup that is the same for both firms for any given product. To analyze consumer welfare in this setting where consumer surplus is inapplicable we introduce a novel welfare index that is effectively a weighted average of the total number of products in the market and those in the least complex menu, with the weights determined by the fractions of consumers who can consider the former and the latter, respectively.

Our first result shows that the *maximum-variety* profile, whereby both firms offer all possible products, is always

[1] Minimization of cognitive effort of the kind that is featured in this model has been recognized as one of the four main drivers of choice overload in the above-mentioned meta-analysis of Chernev et al (2015).

an equilibrium when no consumer is overloaded at this grand menu, and in strictly dominant strategies when all products are equi-profitable. This profile is the unambiguous welfare-optimum for consumers even if they are overloaded, provided that they are of the non-deferring type. However, this is not necessarily the case for consumers who defer when overloaded.

We then introduce overload by means of a cumulative density function that is strictly monotonic in menu size, and provide conditions for this profile to remain an equilibrium also in the case of such consumers. Interestingly, although a sufficiently large fraction of consumers must not be overloaded at the grand menu in order for this profile to be an equilibrium, this threshold is decreasing in the total number of products for deferring consumers (up to a certain lower bound), while it is increasing for non-deferring consumers. An intuitive discussion of this fact is offered.

Similarly, we provide conditions for the other extreme case to be sustained as an equilibrium, namely a *minimum-variety* menu profile whereby both firms offer a menu that contains the same single product. While this is the unambiguous welfare worst for both types of consumers, it turns out that, at least in the case where all products have the same markup, it can only be sustained in extreme situations where, for instance, a very large fraction of consumers is overloaded in menus with only two products.

Both the above results employ the assumption that all products are equiprofitable. Our last result examines the case where this assumption is relaxed on the one hand, and overload is assumed to be homogeneous in the population on the other. It is shown that under these conditions a symmetric profile where the common menu of the two firms is the most profitable and largest that can be offered without overloading the consumers is an equilibrium. Moreover, it turns out that this is the unique symmetric equilibrium in the case where all consumers defer when overloaded.

2 The Individual Consumer

We assume throughout that there is a finite set X of choice alternatives. A *menu* is a nonempty subset of X. The collection of all menus is denoted by $\mathcal{M}$, while $\mathcal{M}^* := \mathcal{M} \cup \{\varnothing\}$. A consumer's decisions across menus is captured by a choice correspondence $C : \mathcal{M} \to \mathcal{M}^*$. This is a possibly multi- and empty-valued mapping that satisfies $C(A) \subseteq A$ for all A in $\mathcal{M}$. Whenever $C(A)$ contains more than one alternative we interpret them as those which the decision maker might choose from menu A, and consider each of them to be equally likely to be chosen. On the other hand, when $C(A) = \varnothing$ we understand that the decision maker *avoids* or *defers* choice at menu A. Choice avoidance and choice deferral are generally not identical concepts, for the latter suggests that choice from a menu may be delayed whereas the former leaves open the possibility that choice from that menu will never be made. However, for the purposes of this paper the distinction is not important as we will not formally incorporate a temporal dimension in the consumer's decision making process to model behavior after choice of no alternative has been made in a menu.

2.1 Within-Menu Decisions

Our consumers are assumed to face cognitive costs during their decision-making process which, when high enough, induce choice deferral. In particular, we assume that their behavior is described by the model of *overload-constrained utility maximization* that is developed and axiomatically characterized in Gerasimou (2015). In this model the consumer's within-menu behavior is guided by a *utility function* $u : X \to \mathbb{R}$, a *complexity function*

$\psi : \mathcal{M} \to \mathbb{R}$ and a *complexity threshold* n (an integer), and is such that for all A, B in $\mathcal{M}$ and all x, y, z in X,

$$C(A) = \begin{cases} \emptyset, & \textit{if and only if } \psi(A) > n \\ \\ \arg\max_{x \in A} u(x), & \textit{otherwise} \end{cases} \tag{1a}$$

$$\psi(\{x,y\}) \leq n \;\; \& \;\; \psi(\{y,z\}) \leq n \implies \psi(\{x,z\}) \leq n \tag{1b}$$

$$\psi(\{x\}) \leq n \tag{1c}$$

$$B \supset A \implies \psi(B) \geq \psi(A) \tag{1d}$$

The interpretation of part (1a) in this model is that the agent defers choice at menu A if and only if the cognitive cost that this menu would impose on him if considered carefully (which is captured by the complexity measure $\psi(A)$) exceeds his idiosyncratic complexity threshold n. Yet, if the cognitive cost is below that threshold, he is portrayed as a standard utility maximizer. The complexity function ψ may stand, for example, for the number of alternatives in a menu, the number of relevant product attributes, or the time required for the consumer to make a choice at a menu. The interpretation of part (1b) is that the consumer's complexity criterion satisfies a binary-menu transitivity property in the sense that if menus $\{x, y\}$ and $\{y, z\}$ are not considered to be complex, then neither should menu $\{x, z\}$ be. In the special case that we consider below where menu complexity coincides with menu cardinality this condition is clearly satisfied. Parts (1c) and (1d) simply state that singleton menus are never complex and that complexity is weakly increasing with respect to menu inclusion, respectively.

With regard to the related choice-theoretic literature, Frick (2015) proposes another model that features a monotone threshold representation similar to (1), which predicts increasingly inconsistent behaviour as the menu size increases but does not allow for choice deferral. Dean (2008) and Dean et al (2014) propose models that predict a positive effect of choice overload on *status quo bias*, a phenomenon that is related but distinct from choice deferral.[2] Sarver (2008) and Ortoleva (2013) provide preference representation theorems in the context of choice over menus of risky options where individuals are modelled as having a preference for smaller menus due to regret anticipation and thinking aversion, respectively. For a choice-theoretic foundation of Sarver's model that also features choice deferral the reader is referred to Buturak and Evren (2015).

2.2 Between-Menu Decisions

The previous subsection describes the consumer's behavior when he is facing a menu of choice alternatives. In anticipation of the duopolistic model that is presented below we will now introduce a procedural model that describes our consumer's decisions when he is faced with the problem of *choosing between two menus*. Intuitively, such a problem comes about at the decision stage that precedes the one in which overload-constrained utility maximization in the sense of (1) takes place.

As has already been mentioned we will now assume that the function ψ takes the form

$$\psi(A) = |A|$$

so that menu complexity coincides with menu cardinality. This assumption is consistent with the marketing and consumer psychology literature on choice overload that motivates this paper. Moreover, we define $\mathcal{D} := \mathcal{M} \times \mathcal{M} = \{(A, B)\} : A, B \in \mathcal{M}\}$ to be the collection of all pairs of menus in $\mathcal{M}$. The consumer's choice or deferral in

[2] Gerasimou 2016 provides a formal distinction of the two concepts and a related conceptual discussion. For a comparative analysis that brings these two phenomena under the common umbrella of *"decision avoidance"* see Anderson 2003.

this class of decision problems is captured by the correspondence $S : \mathcal{D} \to \mathcal{M}^* \cup \mathcal{D}$ that is defined by

$$S(A,B) = \begin{cases} A, & \text{if} \quad |A| \leq n < |B|, \quad \textit{or} \\ & \qquad |A|, |B| \leq n \ \ \& \ \arg\max_{x \in A \cup B} u(x) \in A \setminus B \\ \\ B, & \text{if} \quad |B| \leq n < |A|, \quad \textit{or} \\ & \qquad |A|, |B| \leq n \ \ \& \ \arg\max_{x \in A \cup B} u(x) \in B \setminus A \\ \\ \{A,B\}, & \text{if} \quad |A|, |B| \leq n \ \ \& \ \arg\max_{x \in A \cup B} u(x) \in A \cap B \\ \\ \emptyset, & \text{if} \quad |A|, \ |B| > n \end{cases} \tag{2}$$

if the consumer is allowed to defer choice between menus A and B, and by

$$S(A,B) = \begin{cases} & \qquad |A| \leq n < |B|, \quad \textit{or} \\ A, & \text{if} \quad n < |A| < |B|, \quad \textit{or} \\ & \qquad |A|, |B| \leq n \ \ \& \ \arg\max_{x \in A \cup B} u(x) \in A \setminus B \\ \\ & \qquad |B| \leq n < |A|, \text{or} \\ B, & \text{if} \quad n < |B| < |A|, \quad \textit{or} \\ & \qquad |A|, |B| \leq n \ \ \& \ \arg\max_{x \in A \cup B} u(x) \in B \setminus A \\ \\ \{A,B\}, & \text{if} \quad |A| = |B|, \quad \textit{or} \\ & \qquad |A|, |B| \leq n \ \ \& \ \arg\max_{x \in A \cup B} u(x) \in A \cap B \end{cases} \tag{3}$$

if he must necessarily choose a menu when faced with this decision problem. In words, (2) and (3) suggest that irrespective of whether the consumer can defer or not, he chooses menu A over B either if B is complex and A is not, or if both A and B are non-complex but the agent's most preferred option is in A and not in B. If the consumer cannot defer, however, he may also choose A over B if A is relatively less complex than B even if A itself is complex. The case where B is chosen over A is symmetric. The consumer is indifferent between A and B and may choose one of them at random if both menus are non-complex and contain his most preferred option from those available. In the case of non-deferring consumers indifference also occurs when A and B are equally complex. Finally, if the consumer can defer, he does so when both A and B are complex.

Effectively this specification assumes that the consumer uses a *lexicographic rule* in decisions between menus, whereby he first compares them according to complexity and only resorts to his preferences when both menus are non-complex. This is consistent with our general approach of modelling consumers as trying to make decisions by minimizing cognitive effort. However, caution should be exercised in the interpretation of the consumer's postulated decision process when his preferences come into play. In particular, when faced with the pair of menus (A, B) our consumer may be thought of as pre-scanning A first according to the cardinality criterion and if it is considered to be non-complex, then and only then does he employ his preferences. He is then assumed to do the same with menu B. If both menus are non-complex (or equally complex in the case of non-deferring consumers), he compares the two best options in each of them and chooses his menu accordingly. Therefore, the consumer is not assumed to consider $A \cup B$ as a single menu, which may go against the nature of our model if this "menu" $A \cup B$ is complex according to his criterion even though its individual components A and B are not.

It is worth noting that our model of between-menu decision making assumes that the consumer correctly anticipates that complex menus will be cognitively taxing once chosen. Interestingly, however, the empirical evidence on between-menu decisions suggests that the choice-overload effect is weaker and sometimes even reversed when consumers are faced with decisions of this type (Chernev, 2006; Chernev and Hamilton, 2009; Goodman and Malkoc, 2012). That is, consumers have been shown to have a weak "preference for flexibility" (Kreps, 1979). Chernev et al (2015, p. 353) interpret this finding as suggesting that *"when choosing between assortments, consumers are likely to underestimate the choice overload they will experience when making their final choice from the assortment selected"*. While incorporating such biased expectations into the analysis is certainly interesting, in this paper we focus on the benchmark theoretical case where consumers have correct expectations, leaving such extensions for future work.

3 A Duopolistic Model

3.1 General Setup

We consider a duopolistic market with a fixed set $X = \{x_1, \ldots, x_k\}$ of $k \geq 3$ products. Any one product is homogeneously produced by the two firms. We assume that there is a continuum of consumers with mass one who have heterogeneous preferences over the elements of X. One may, of course, also interpret this specification as there being a single consumer over whose preferences the firms are uncertain.

The two firms are assumed to compete simultaneously in menu design, i.e. they choose (possibly in a probabilistic manner) a menu in $\mathcal{M}$ to offer. We abstract from the firms' pricing decisions in order to keep the analysis simple and to focus exclusively on the effects of this kind of menu-design competition in an environment where consumers may be facing cognitive costs of the type that were discussed above. We note, however, that even though prices are often the most important variables in competition between real firms, there are also important examples such as TV/radio channels, online search engines, web browsers, free newspapers etc. where prices are completely absent. In the latter example, for instance, menu design of the kind we consider here may correspond to the number of pages and/or articles that each free newspaper may choose to offer.

To this end, we will let each product x_i be associated with an exogenously given markup $w_i > 0$ that applies to both firms. In the free newspapers example, this may be thought of as arising from a common marginal cost of including an additional page and/or article as well as a common expected return from the introduction of such an entry. Another way to interpret this assumption would be to think of each firm as having a different belief about the fraction of consumers who can afford the given product. In this case, even if the product's price and marginal costs are different across the two firms, the non-coinciding beliefs might be assumed to offset the resulting difference. Specifically, if p_i^f and c_i^f denote firm f's price and marginal cost of producing x_i, and $r_i^f \in (0,1)$ denotes firm f's subjective probability about the fraction of consumers who can afford x_i, then $w_i^1 = w_i^2 = w_i$ if $r_i^1(p_i^1 - c_i^1) = r_i^2(p_i^2 - c_i^2)$. In this case, w_i can be thought of as firm i's *expected* or *weighted* markup for product x_i.

3.2 Choice Probabilities

For menu $A = \{x_j, \ldots, x_m\}$ we define the index set I_A by $I_A = \{j, \ldots, m\}$. For any two menus A and B, let $A \setminus B := \{x_j \in A : x_j \notin B\}$. To ease notation, we will write AB interchangeably with $A \cup B$. The probability of x_i being chosen when all elements of X are available in the market is generally denoted by $p_X(x_i)$. This may be thought of as the *unconditional choice probability* of x_i. On the other hand, the probability that x_i is chosen *conditional on menus A and B being offered* by the two firms is denoted by $p_{AB}(x_i)$.

Assumption 1:
For every product x_i and all menus A and B,

$$p_{AB}(x_i) = \frac{p_X(x_i)}{\sum\limits_{j \in I_{AB}} p_X(x_j)} = \frac{1}{|A \cup B|}$$

This first part of this assumption, captured by the first equation, effectively imposes the "Luce axiom" (Luce, 1959) on the products' choice probabilities. An important and well-known implication of this axiom is that the relative choice probability of two products is independent of what other products are available in the menu. The second part effectively postulates that the consumers' preferences are uniformly heterogeneous so that any two products' choice probabilities are always the same. Assuming the uniform distribution obviously makes the analysis more tractable than it would have been otherwise, but it may also be suitable as a modelling approach if there is genuine uncertainty over the distribution of consumers' preferences, in which case one may invoke the well-known "principle of insufficient reason". Although the second part implies the first, we have chosen to include both as this suggests a way in which the assumption may be relaxed.

3.3 Overload Probabilities

A crucial element of our duopolistic model is that consumers are faced with cognitive costs that translate into choice overload of the kind that was described above in the previous section. Specifically, for each menu $A \in \mathcal{M}$ with $|A| \geq 1$ we assume that there is a fraction $q(A) \equiv q(|A|)$ of consumers who are *not* choice-overloaded at menu A. The function q is a cumulative density function (cdf) on the set of strictly positive integers, capturing the distribution of choice overload in the consumer population.

Assumption 2:
The (no-) overload cumulative density function (cdf) q satisfies the following:
a) $|A| = 1$ implies $q(A) = 1$.
b) $|A| = |B|$ implies $q(A) = q(B)$.
c) $|A| > |B|$ implies $q(A) \leq q(B)$.

Consistent with the models of individual decision making that were proposed in (1), (2) and (3), this assumption suggests that consumers may be choice-overloaded in the sense that there is a threshold menu size for each of them so that whenever a menu exceeds this threshold the consumer who is overloaded at that menu does not even look at its contents and hence does not buy from the firm that is offering it. Part a) requires no consumer to be overloaded at singleton menus, consistent with condition (1c) of the underlying choice model. Part b) states that menu size is the only source of overload by requiring the same fraction of consumers to be (non-)overloaded at menus of equal size. This is consistent with the between-menu decision model as captured in both (2) and (3). Finally, part c) requires the no-overload cdf q to be weakly monotonic in menu size in the sense that at larger menus there are weakly less non-overloaded consumers than at smaller ones.

3.4 Payoffs

We are now ready to specify the firms' payoff functions. We start by defining the baseline payoffs $R_1(A, B)$ and $R_2(A, B)$ at pure-strategy profile (A, B) as follows:

$$\begin{aligned} R_1(A, B) &= \sum_{i \in I_{A \setminus B}} p_{AB}(x_i) w_i + \frac{1}{2} \sum_{j \in I_{A \cap B}} p_{AB}(x_j) w_j \\ R_2(A, B) &= \sum_{i \in I_{B \setminus A}} p_{AB}(x_i) w_i + \frac{1}{2} \sum_{j \in I_{A \cap B}} p_{AB}(x_j) w_j \end{aligned}$$

The first equation defines the first firm's payoff (i.e. the one offering menu A) at profile (A, B) when overload considerations are absent. The first term on the right hand side of this equation is the expected contribution to that firm's payoff from those products that it offers uniquely, i.e. those that are in A but not in B. The second term captures the expected contribution to its payoff from those products that are offered by both firms. Consistent with the indifference condition in (2) and (3) in the between-menus decision model, consumers are assumed to randomize uniformly between the two firms when both menus are considered and both include their most preferred option. To capture this aspect of that choice procedure, the second expected contribution is scaled down by the probability $\frac{1}{2}$. Baseline payoff $R_2(A,B)$ is similarly interpreted.

With this notation at hand we can now go ahead and introduce the firms' payoff functions taking into account also our assumption that some consumers may be overloaded.

Assumption 3:
Suppose that there is a fraction $\mu \in [0,1]$ of consumers who do not defer when overloaded. Then, firm 1's payoff at profile (A,B) is given by

$$\pi_1(A,B) = \begin{cases} [\mu + (1-\mu) \cdot q(A)] \cdot R_1(A,B), & \text{if } |A| = |B| \\ \\ q(A) \cdot R_1(A,B), & \text{if } |A| > |B| \\ \\ q(B) \cdot R_1(A,B) + \Big[\mu \cdot [1 - q(A)] + (1-\mu) \cdot [q(A) - q(B)]\Big] \cdot \sum_{i \in A} p_A(x_i) w_i, & \text{if } |B| > |A| \end{cases}$$

Firm 2's payoff is defined symmetrically.

We will explain this payoff function in words by considering the two extreme cases where $\mu = 1$ and $\mu = 0$. In the former case no consumer ever defers regardless of whether he is overloaded or not. Individual decisions in this case are captured by (3). Consistent with (3) the payoff function now specifies the baseline payoff $R_1(A,B)$ for firm 1 whenever A and B are equally complex, so as to be compatible with the assumption that "there is no easy way out" for the consumers in this scenario, as they go through both A and B before choosing. If A is more complex than B, however, it follows from (3) that consumers who are overloaded at A will opt for B. This is reflected in the payoff function by multiplying the baseline payoff $R_1(A,B)$ by the fraction $q(A)$ of consumers who are not overloaded at A and will therefore consider both A and B. Finally, if the rival menu B is more complex than A, then a fraction $q(B)$ of consumers are still able to examine both but, consistent with (3), the remaining fraction $1 - q(B)$ of overloaded consumers only considers menu A, so that firm 1's expected payoff from these consumers is as if menu A was the only one in the market.

Let us now consider the case where $\mu = 0$ and all overloaded consumers defer. Consistent with (2), only the consumers who are not overloaded at A will consider both menus if these are equally complex. As with the case where $\mu = 1$ this remains true when A is more complex than B. Thus, the baseline payoff $R_1(A,B)$ is scaled down by $q(A)$ whenever $|A| \geq |B|$. However, when the rival menu B is more complex, firm 1 still has its menu considered by the consumers who are able to examine both A and B (hence the first term in its payoff is identical to that where $\mu = 1$), but now this firm will not attract all of the remaining consumers to menu A but only those who are overloaded at B and not in A. This is reflected in the second part of its payoff when $|B| > |A|$.

This payoff specification highlights the key tension that the firms in our model are assumed to be facing. On the one hand, firms aspire to appeal to as many consumers as possible, which puts an upward pressure on the size of the menu they will offer. At the same time, however, firms are mindful of the fact that different consumers are overloaded at different menu sizes, and that the larger the menu they offer, the smaller will generally be the fraction of consumers who choose from it, with the remaining ones either choosing from the rival menu or avoiding choice altogether. This strategic trade-off is at the heart of our model and makes the analysis both interesting and non-trivial.

3.5 Consumer Welfare

Because our model abstracts from the firms' pricing decisions, it also lacks a standard notion of welfare such as consumer surplus. A new such concept is therefore introduced which helps to analyse whether consumers are better- or worse-off across different equilibria, and also motivates a relevant comparative-statics analysis.

Assumption 4:
When a fraction $\mu \in [0,1]$ of consumers do not defer when overloaded, consumer welfare is captured by the ***weighted cardinality index*** $W : \mathcal{M} \times \mathcal{M} \to \mathbb{R}$ *defined by*

$$W(A,B) = \begin{cases} [\mu + (1-\mu)\cdot q(A)]\cdot |A \cup B|, & \text{if } |A| = |B| \\ \mu \cdot \Big[q(A)\cdot |A\cup B| + [1-q(A)]\cdot |B|\Big] + (1-\mu)\cdot \left[\dfrac{q(A)\cdot |A\cup B| + [q(B)-q(A)]\cdot |B|}{q(A)+q(B)}\right], & \text{if } |A| > |B| \\ \mu \cdot \Big[q(B)\cdot |A\cup B| + [1-q(B)]\cdot |A|\Big] + (1-\mu)\cdot \left[\dfrac{q(B)\cdot |A\cup B| + [q(A)-q(B)]\cdot |A|}{q(A)+q(B)}\right], & \text{if } |B| > |A| \end{cases}$$

Similar to the strategic trade-off faced by the firms, the proposed index is motivated by the consumers' own trade-off that is generated by having many versus few alternatives in the market. Specifically, the general idea underlying it is that consumers benefit by having a large number of products made available in the market by the two firms because this increases the chances that more consumers will find their most preferred product. However, at the same time consumers are harmed by the fact that some of these products are not actively considered by some of them due to choice overload. Thus, both at the extreme case where $\mu = 1$ or $\mu = 0$ as well as in the cases in between these extremes, the index may be thought of as the average of the total number of products in the market weighted by the fraction of consumers who can consider all these products before choosing and the number of products in the least complex menu weighted by the fraction of consumers who consider only those products. The definition of this weighted average differs between the two extreme cases $\mu = 1$ and $\mu = 0$ because in the latter case the two relevant consumer fractions may not add up to one. For μ strictly between 0 and 1 the welfare index is the corresponding convex combination of the two weighted averages.

Observation 1. *When no overloaded consumer defers (i.e. $\mu = 1$), the consumer-optimal symmetric equilibrium is attained at the maximum-variety profile (X,X):*

$$W(X,X) = |X| \geq W(A,B) \quad \textit{for all } A,B \in \mathcal{M} \qquad\qquad \textit{if } \mu = 1.$$

Due to the symmetry of profile (X,X) and of the fact that consumers cannot escape choice even though their cognitive costs may be high at X, they are forced to go the trouble of considering all options from either firm and to make a choice as if they were facing no cognitive costs. This effort on their part in turn translates into a higher level of welfare.

As the next observation clarifies, however, at the other extreme where all overloaded consumers defer, the maximum-variety equilibrium may in principle be Pareto-dominated by another symmetric equilibrium that offers lower variety.

Observation 2. *If (A,A) and (X,X) are equilibria and $A \subset X$, then*

$$W(A,A) > W(X,X) \iff \frac{q(A)}{q(X)} > \frac{|X|}{|A|} \qquad\qquad \textit{if } \mu = 0.$$

In words, a lower-variety symmetric equilibrium (A,A) Pareto-dominates the maximum variety equilibrium (X,X) if and only if the fraction of overloaded consumers when moving from A to X grows at a higher rate than the rate at which the menu itself grows in the transition from A to X.

3.6 Symmetric Equilibrium Profits

Much of our analysis revolves around symmetric pure-strategy equilibria, which in a certain sense arise naturally in our model as we explain below. It will be useful to take note of the following fact for this class of equilibria.

Observation 3. *Suppose all products are equi-profitable in the sense that their common markup is w. If (A, A) is a symmetric equilibrium, then*

$$\pi(A,A) = \begin{cases} \dfrac{w}{2}, & \textit{if no consumer defers} \quad (\mu = 1) \\[2ex] q(A) \cdot \dfrac{w}{2}, & \textit{if all overloaded consumers defer} \quad (\mu = 0) \end{cases}$$

The two firms are therefore indifferent across all symmetric equilibria when overloaded consumers do not defer, while they weakly prefer lower-variety equilibria over higher-variety ones when all overloaded consumers do defer (they strictly prefer such equilibria whenever q is strictly decreasing). Thus, given that consumer welfare is assumed to coincide with the weighted-cardinality index that was introduced above, the unambiguous Pareto optimum from the aggregate-welfare point of view in the case where no consumer defers due to overload is the maximum-variety equilibrium (X, X). When some overloaded consumers do defer, however, such a generally-applicable ranking of symmetric equilibria is not possible, as is also implied by Observation 2. We note that Observation 3 generalizes to the case of symmetric mixed-strategy equilibria.

4 Maximum- and Minimum- Variety Equilibria

4.1 The No-Overload Benchmark

Our first result concerns the benchmark case where no consumer is overloaded at the grand menu X. The proofs of this and all other results that follow appear in the Appendix.

Proposition 1. *If no consumer is overloaded at X, then the maximum-variety profile (X, X) is a strict equilibrium. If, in addition, all products are equi-profitable, then (X, X) is an equilibrium in strictly dominant strategies.*

Indeed, conditional on a firm's opponent offering menu X, the firm's *unique* best response is to offer X as well. This is so because the choice probability of each product is $\frac{1}{k}$ regardless of which menu $A \subseteq X$ the firm chooses to offer. Moreover, since no consumer is overloaded at X and hence $q(X) = 1$, 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.

Coming to the case where all product markups are equal to a common $w > 0$, offering X in this case is 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, however, this is not generally the case when products differ in their profitability. As the following example illustrates, this is so because when the difference between the least profitable product and the one just above it is big enough, introducing the former product may be unprofitable due to a possibly offsetting role of the reduction in all products' choice probabilities.

Example 1. *Suppose $X = \{x_1, x_2, x_3, x_4\}$, $w_1 = w_2 = w_3 = 1$, $w_4 = \frac{1}{10}$ and $q(X) = 1$. Let $A_{ijk} := \{x_i, x_j, x_k\}$. We have*

$$\begin{aligned} \pi_1(X, X) &= \frac{1}{2}\left(\frac{3}{4} \cdot 1 + \frac{1}{4} \cdot \frac{1}{10}\right) = \frac{31}{80} \\ \pi_1(A_{123}, X) &= \frac{1}{2} \cdot \frac{3}{4} \cdot 1 = \frac{30}{80} \\ \pi_1(A_{12}, X) &= \frac{1}{2} \cdot \frac{2}{4} \cdot 1 = \frac{20}{80} \\ \pi_1(A_1, X) &= \frac{1}{2} \cdot \frac{1}{4} \cdot 1 = \frac{10}{80} \end{aligned}$$

which shows that (X, X) is an equilibrium. Moreover, we also have

$$\begin{aligned}
\pi_1(A_{123}, A_{123}) &= \frac{1}{2} \\
\pi_1(X, A_{123}) &= \frac{1}{2} \cdot \frac{3}{4} \cdot 1 + \frac{1}{4} \cdot \frac{1}{10} = \frac{32}{80} \\
\pi_1(A_{12}, A_{123}) &= \frac{1}{2} \cdot \frac{2}{3} \cdot 1 = \frac{1}{3} \\
\pi_1(A_{14}, A_{123}) &= \frac{1}{2} \cdot \frac{1}{4} \cdot 1 + \frac{1}{4} \cdot \frac{1}{10} = \frac{3}{20} \\
\pi_1(A_{1}, A_{123}) &= \frac{1}{2} \cdot \frac{1}{3} \cdot 1 = \frac{1}{6}
\end{aligned}$$

which shows that (A_{123}, A_{123}) is an equilibrium too.

4.2 Overloaded Consumers and Maximum-Variety Equilibrium

We now move on to study the implications that our consumers' postulated cognitive costs and the ensuing choice overload have in this duopolistic market. Our next result identifies conditions under which the maximum-variety profile (X, X) remains an equilibrium despite such overload.

Proposition 2. *Suppose all products are equi-profitable and q satisfies the condition*

$$\frac{q(h-1)}{q(h)} \geq \frac{3k-h+1}{3k-h} \qquad \text{for all } h \leq k-1. \tag{4}$$

Then, (X, X) is an equilibrium if and only if

$$q(X) \geq \begin{cases} \dfrac{k}{k+1}, & \text{if no consumer defers} \quad (\mu = 1) \\[2ex] \dfrac{2k}{3k-1}, & \text{if all overloaded consumers defer} \quad (\mu = 0) \end{cases} \tag{5}$$

Condition (4) implies that q is strictly decreasing. In addition, it imposes a lower bound on how rapidly the consumer population becomes overloaded when the menu size increases marginally. An example where condition (4) is always satisfied with equality is the one where $k = 33$, the support of q is the set $S := \{1, \ldots, 98\}$ and q coincides with the uniform distribution on this set, i.e. $q(h) = \frac{100-h}{99}$. In this example, the condition is satisfied with strict inequality if the support of q is a proper subset of S, while it is violated if it properly contains S.

What Proposition 2 then says is that under this structural condition on q and the assumption that all products are equi-profitable, the maximum-variety profile (X, X) is an equilibrium if and only if sufficiently many consumers are not overloaded at the grand menu X that contains all k products. While this threshold size is a function of k both when consumers do and do not defer when overloaded ($\mu = 0$ and $\mu = 1$, respectively), an interesting feature is that this threshold is strictly increasing in k for non-deferring consumers and strictly decreasing (up to the lower bound of $\frac{2}{3}$) for deferring ones.

Suppose consumers force themselves to choose despite being overloaded. Suppose also that, due to technological progress or some other reason, the number k of all available products increases to $k' > k$. If the fraction of consumers who are not overloaded at k is the same as that at k' so that $q(k) = q(k')$, then offering a strictly smaller menu A when the opponent offers all k' products becomes more attractive than if $q(k') > q(k)$ was true instead. This is so because the firm can offer a menu A with k products and attract all consumers who are overloaded at the new X while at the same time minimizing the forgone expected profit from the products that are in $X \setminus A$. If the fraction of overloaded consumers is large enough, this trade-off may be such that a deviation is profitable. When $q(k') > q(k)$, however, the probability that offering such a menu is indeed a profitable deviation decreases because the fraction of overloaded consumers has also decreased. Therefore, the threshold cumulative probability $q(X) \equiv q(k)$ is strictly increasing in k in order to offset the forces that might render small menus profitable in this case where consumers do not defer. We note that although this comparative-statics exercise shows that $q(k)$ must be increasing in k for (X, X) to be sustained as an equilibrium as k varies, the density q itself for a given and fixed k need not satisfy condition (4).

Figure 1: Threshold conditions for the maximum-variety equilibrium (X, X)

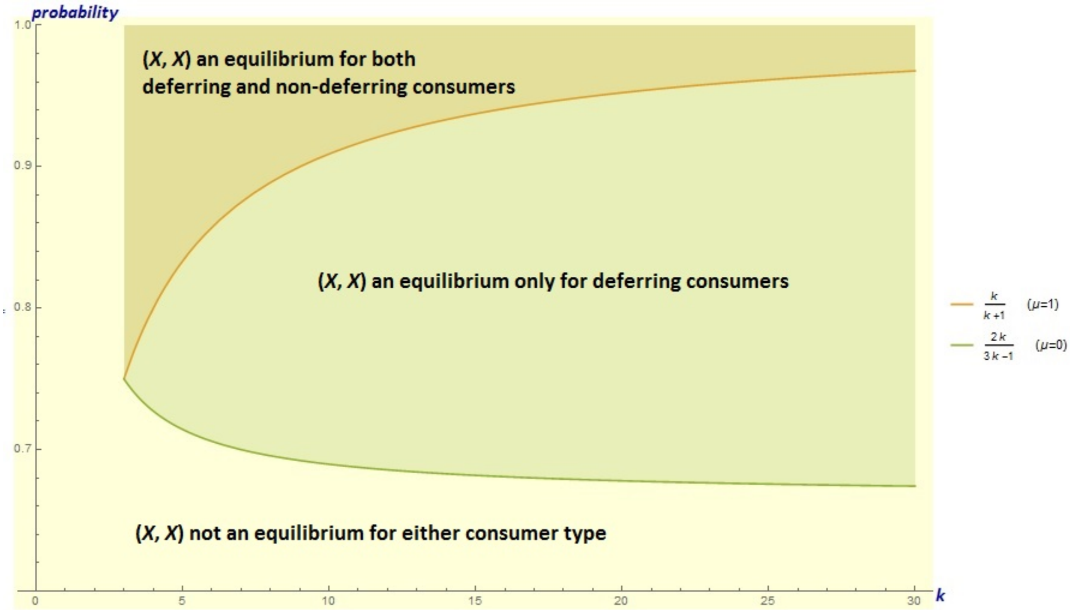

The situation is different, however, when overloaded consumers actually defer when overloaded instead of choosing the relatively less complex menu. In particular, deviating to a smaller menu A in this case is not as profitable as before because a fraction $1 - q(A)$ of the consumers remain overloaded and hence defer at A. Consequently there is less pressure on $q(X)$ to be as high as in the case of non-deferring consumers. Moreover, suppose now that the total number of products *de*creases from k to $k'' < k$. Suppose also that $q(k) = q(k'')$. Since q is assumed to satisfy (4) which in turn implies that q is strictly decreasing, the higher probability mass to the left of k'' is distributed to fewer cardinality values than those lying to the left of k. Thus, the fraction of consumers who are not overloaded at menus with two alternatives, for instance, is strictly closer to 1 under a total of k'' products than under k, rendering those menus candidates for profitable deviations from X. To offset this effect, $q(X)$ must be strictly decreasing in k.

Example 2. *Suppose $X = \{x_1, \ldots, x_4\}$, $w_i = w_j = w = 1$ for $i \neq j$, and let $q(A) = [1 - \frac{1}{20}(|A| - 1)]$. Since $q(X) = 0.85 > 0.8 = \frac{k}{k+1}$, (X, X) is an equilibrium for both deferring and non-deferring consumers.*

In particular, in the case of consumers who defer when overloaded ($\mu = 0$), we have

$$
\begin{aligned}
\pi_1(X, X) &= 0.85 \cdot \frac{1}{2} = 0.425 \\
\pi_1(A_3, X) &= 0.9 \cdot \left(0.85 \cdot \frac{1}{2} \cdot \frac{3}{4} + 0.15\right) = 0.422 \\
\pi_1(A_2, X) &= 0.95 \cdot \left(0.85 \cdot \frac{1}{2} \cdot \frac{2}{4} + 0.15\right) = 0.344 \\
\pi_1(A_1, X) &= 1 \cdot \left(0.85 \cdot \frac{1}{2} \cdot \frac{1}{4} + 0.15\right) = 0.256
\end{aligned}
$$

where A_1, A_2 and A_3 denote generic menus with 1, 2 and 3 products, respectively. Moreover, we also have

$$\pi_1(\{x_1, x_2\}, \{x_3, x_4\}) = 0.95 \cdot \frac{1}{2} < \pi_1(X, \{x_3, x_4\}) = 0.85 \cdot \left(\frac{1}{2} \cdot \frac{2}{4} + \frac{1}{2}\right) = 0.6375$$

which shows why it may be hard for asymmetric profiles (A, B) to be equilibria in general: Once the market is split because none of the available products is offered by both firms, each firm has a profitable deviation to offer the products of its rival in addition to the ones of its own.

4.3 Overloaded Consumers and Minimum-Variety Equilibrium

We proceed next to analysing the case that corresponds to the unambiguously welfare-worst outcome for both deferring and non-deferring consumers, namely symmetric equilibria (A, A) where menu A is a singleton. We will refer to such profiles (and equilibria) as *minimum-variety* ones. In particular, our next result provides a general characterization of such equilibria under the assumption of equi-profitability.

Proposition 3. *Suppose all products are equi-profitable. Then, for all $\mu \in [0, 1]$, a minimum-variety profile (A, A) with $|A| = 1$ is an equilibrium if and only if the cdf q first-order stochastically dominates the cdf $\widehat{q}$ that is defined by*

$$\widehat{q}(h) = \frac{h}{2h - 1}.$$

Had menu size been assumed to be a continuous variable, then the target cdf $\widehat{q}$ in Proposition 3 would have been generated by the pdf $\widehat{p}$ defined by $\widehat{p}(h) = \frac{1}{(2h-1)^2}$. As shown in Figure 2, the main feature of this distribution is the very sharp drop in the fraction of overloaded consumers when menu size moves from one to two. In addition to this, however, what the characterization given above requires of q for minimum-variety profiles to be equilibria is that for each menu size $h \geq 2$, the fraction of consumers who are not overloaded at h is weakly smaller under q than under $\widehat{q}$, and strictly so for some h.

Figure 2: The target cdf $\widehat{q}$ for minimum-variety equilibria and the associated pdf

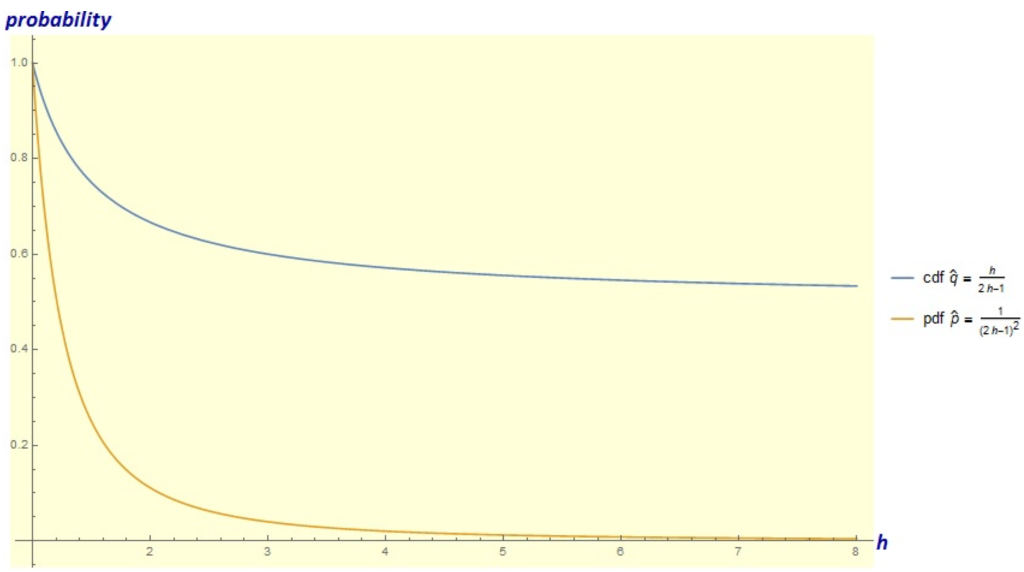

The natural question, of course, is what types of cumulative density functions (and under what additional assumptions) result in a cdf q that first-order stochastically dominates $\widehat{q}$. It can be easily demonstrated graphically that neither the uniform nor the (normal approximation of the) binomial cdf dominates $\widehat{q}$ in this way under reasonable specifications. Yet, it can also be verified that the geometric cdf $\widetilde{q}$ defined by $\widetilde{q}(h) = (1-p)^h$ does dominate

$\widehat{q}$ in this sense when the "probability of success in each Bernoulli trial" p is such that $p \gtrapprox 0.14$. However, the appropriateness of this cdf as a modelling tool for the present purposes is questionable. In any case, the conclusion one may draw from Proposition 3 and this discussion is that minimum-variety profiles are attainable as equilibria only under extreme conditions.

4.4 Special Case: Uniformly Distributed Overload

As already mentioned, the uniform density may be a useful modelling tool when limited information is available on the behaviour of some relevant random variable such as the consumers' most preferred products or overload thresholds. To this end, we now consider the special case of Propositions 1 and 3 when overload is assumed to be uniformly distributed over the interval $\{1, \ldots, k+n\}$, where n is a (possibly negative) integer that is implicitly defined by $q(k+n) > 0$ and $q(k+n+1) = 0$. That is, $k+n$ denotes the menu size threshold of the *least overloaded consumer* in the population. We note that the formula for the overload cdf is $q(h) = \frac{k+n-h}{k+n-1}$ for $h \leq k+n$.

Corollary 4. *Suppose all products are equi-profitable and choice overload is uniformly distributed over* $\{1, \ldots, k+n\}$. *Then:*
(I) (A, A) *with* $|A| = 1$ *is an equilibrium if and only if* $k+n \leq 2k$.
(II) (X, X) *is an equilibrium if and only if*

$$k+n \geq \begin{cases} k^2, & \textit{if no consumer defers} \quad (\mu = 1) \\ 3k, & \textit{if all overloaded consumers defer} \quad (\mu = 0) \textit{ and } (4) \textit{ is satisfied} \end{cases}$$

Corollary 4 provides conditions that relate the highest overload threshold in the consumer population with the total number of products which ensure that minimum- and maximum-variety equilibria are possible, respectively.

Figure 3: Maximum- and minimum-variety equilibria under uniformly distributed overload

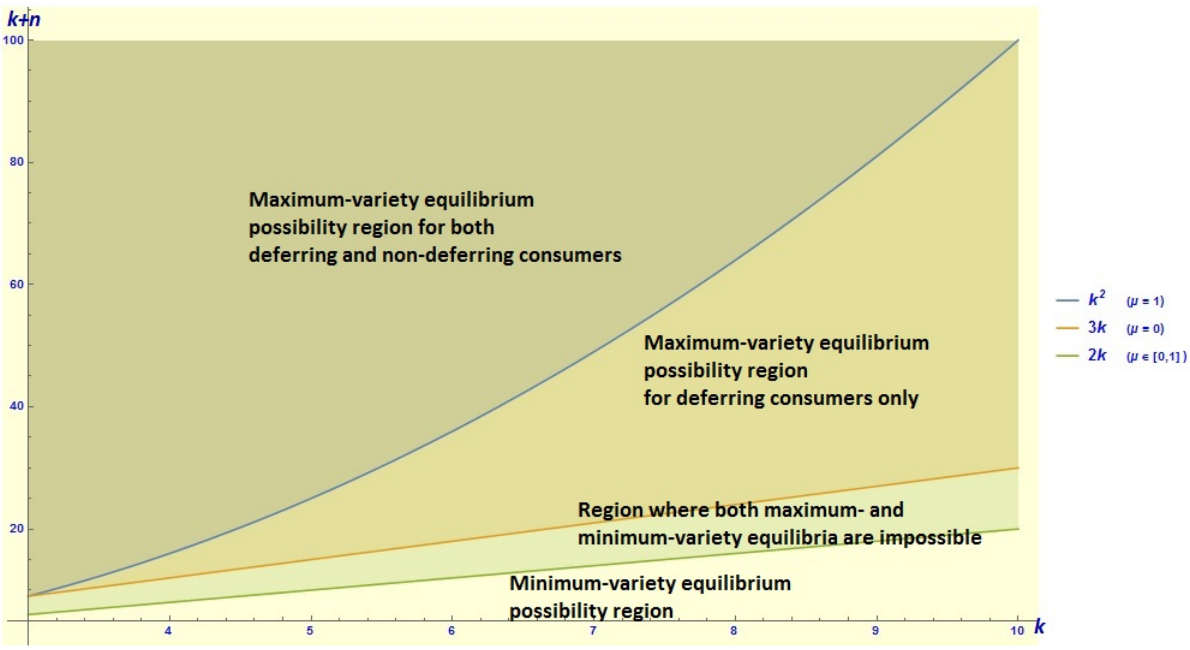

As conditions (I) and (II) make clear, under no circumstances is it possible for both the maximum- and minimum-variety profiles to be equilibria when overload is uniformly distributed. Not surprisingly, the condition for the former kind of equilibria to be possible sets a lower bound on the threshold of the least overloaded consumer relative to the total number of available products, while the opposite is true for minimum-variety equilibria where a lower bound is set instead.

For maximum-variety equilibria, in particular, the highest overload threshold $k+n$ increases quadratically in the total number of products k when consumers do not defer, and linearly so when they do defer once overloaded. The intuition for this disparity is similar to what was discussed in the context of Proposition 1. As the number k of products increases there is a need for a sufficiently high fraction of consumers to not be overloaded at menus of size k, irrespective of whether overloaded consumers defer or not. Since overload is uniformly distributed, however, the fraction of non-overloaded consumers at some fixed menu size $h<k$ relative to the fraction of non-overloaded consumers at k is strictly increasing in k if the upper bound $k+n$ in the support of q stays constant, making such h-sized menus more and more profitable. The rate at which this relative fraction increases with k is strictly higher when consumers do not defer than when they do, due to the reason that was discussed above. This results in a need for a proportionately higher increase in threshold $k+n$ for non-deferring relative to deferring consumers.

5 Unequal Markups and Homogeneous Overload

For simplicity, up until now we've been assuming that all products are equi-profitable in that they have the same markup. Our next result relaxes this assumption and examines the effects that this has in conjunction with another special case of the overload distribution whereby all consumers have the same overload threshold.

Proposition 5. *Suppose there is an integer l such that*

$$q(B)=\begin{cases} 1, & \textit{if } |B|<l \\ 0, & \textit{if } |B|\geq l \end{cases}$$

Then, for all $\mu\in[0,1]$, a profile (A,A) is an equilibrium if A consists of $l-1$ most profitable products. Moreover, if $\mu=0$, then this (A,A) is the unique symmetric equilibrium.

In words, when all consumers are known to ignore menus that exceed a certain size, both firms offering the most profitable and largest menu that consumers will be able to consider is an equilibrium. Intuitively, if one firm unilaterally deviates and offers a menu larger than that, it gets zero payoff because no consumer will look at that menu. If it offers a submenu with fewer options, then this deviation is not profitable either because, by assumption, no consumer is overloaded at the original menu, and therefore the firm cannot hope to attract any consumers this way; on the contrary, it only has to lose.

Less obvious is the fact that the firm has no profitable deviation either when it tries to exploit the trade-off between how profitable a product is and how many firms are offering that product. In particular, at the symmetric profile (A,A) where A consists of the $l-1$ most profitable products, each of them is equally likely to be chosen by either firm. One may think that if the difference in profitability between the $(l-1)$-th and l-th product is very small, then the firm may find it profitable to replace the former with the latter so as not to have this product's share split with the other firm. However, given that each product's choice probability was originally $\frac{1}{l-1}$ while after the firm's unilateral deviation this would decrease to $\frac{1}{l}$, any benefits to the firm that would come from the fact that it is the unique seller of the product would be offset by the fact that the expected payoff of all products in the market would decrease.

6 Related Literature

Our paper seems to be the first in the growing literature of behavioural industrial organization where firms are assumed to compete in the presence of possibly choice-overloaded consumers. However, we note some links between our paper and Kamenica (2008), Goldreich and Halaburda (2013) and Bachi and Spiegler (2015) which precede our work.

Kamenica (2008) studied a monopolist who sells to consumers that may be of two types depending on whether they are informed about their preferences or not. In that model, consumer demand for the firm's equilibrium "product line" may be decreasing in the number of products contained in it because uninformed consumers make the "contextual inference" that a smaller product line includes the most popular alternatives, and hence choose one of these due to the higher probability that it will match their preferences. By contrast, when faced with larger product lines, uninformed consumers may defer due to their lack of information about their preferences and the inability to make such a contextual inference. Our model differs from Kamenica (2008) in two important respects. First, small menus may be offered in the market as the equilibrium outcome of duopolistic competition. Second, such competition takes place in the presence of cognitively-constrained consumers who are fully aware of their preferences, as opposed to fully rational ones who are uncertain about theirs.

Goldreich and Halaburda (2013) documented empirically that, in the context of investment decisions over 401(k) pension plans, larger menus of options in their sample were of lower quality than smaller menus according to the Sharpe ratio of portfolio quality. Such smaller menus were not necessarily submenus of their larger counterparts. They attributed these observations to different levels of menu-setting ability by those in charge with designing the menus to be offered. This motivated a model where menu-setters were either experts or lower-ability ones, and conditions were identified under which the latter type of menu-setters offer larger menus than the former type.

Bachi and Spiegler (2015) study a duopolistic model where consumers are presented with two-attribute products and experience difficulties in making trade-offs across these attributes. Part of that paper deals with the case where consumers actually defer choice due to such comparison difficulties, for example when none of the feasible alternatives is clearly dominant in a menu. Although choice overload as analyzed here is not a source of deferral for consumers in that model, it is the first in this literature to study the effects of indecisiveness-driven choice deferral on market outcomes. For related work on models of individual choice where 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).

The broader theoretical literature where the 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); inattention (Eliaz and Spiegler, 2011; de Clippel, Eliaz, and Rozen, 2014; Bordalo, Gennaioli, and Shleifer, 2016); bounded-rational expectations (Spiegler, 2006); comparison / trade-off difficulty (Piccione and Spiegler, 2012; Papi, 2014, 2015; Fisher and Plan, 2015). A textbook treatment of this literature is provided in Spiegler (2011), while Spiegler (2015) surveys and synthesizes some more recent developments and trends. For recent empirical work on the effects that the presence of loss-averse, inattentive and inattentive/cognitively constrained consumers has on the telecommunications, retail and energy industries the reader is referred to Genakos et al (2015), Clerides and Courty (2016) and Hortacsu, Madanizadeh, and Puller (2015), respectively.

References

ANDERSON, C. J. (2003): "The Psychology of Doing Nothing: Forms of Decision Avoidance Result from Reason and Emotion," *Psychological Bulletin*, 129, 139–167.

BACHI, B., AND R. SPIEGLER (2015): "Buridanic Competition," *mimeo*.

BERGER, J., M. DRAGANSKA, AND I. SIMONSON (2007): "The Influence of Product Variety on Brand Perception and Choice," *Marketing Science*, 26, 460–472.

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

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

CHERNEV, A. (2006): "Decision Focus and Consumer Choice Among Assortments," *Journal of Consumer Research*, 33, 50–59.

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

CHERNEV, A., AND R. HAMILTON (2009): "Assortment Size and Option Attractiveness in Consumer Choice Among Retailers," *Journal of Marketing Research*, 46, 410–420.

CLERIDES, S., AND P. COURTY (2016): "Sales, Quantity Surcharge and Consumer Inattention," *Review of Economics and Statistics*, forthcoming.

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

DEAN, M. (2008): "Status Quo Bias in Large and Small Choice Sets," *mimeo*.

DEAN, M., O. KIBRIS, AND Y. MASATLIOGLU (2014): "Limited Attention and Status Quo Bias," *mimeo*.

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. (2015): "Monotone Threshold Representations," *Theoretical Economics*, forthcoming.

GENAKOS, C., C. ROUMANIAS, AND T. VALLETTI (2015): "Loss Aversion on the Phone," *mimeo*.

GERASIMOU, G. (2015): "Indecisiveness, Undesirability and Overload Revealed Through Rational Choice Deferral," *MPRA Working Paper 67290*.

——— (2016): "Asymmetric Dominance, Deferral and Status Quo Bias in a Behavioral Model of Choice," *Theory and Decision*, forthcoming.

GOLDREICH, D., AND H. HALABURDA (2013): "When Smaller Menus Are Better: Variability in Menu-Setting Ability," *Management Science*, 59, 2518–2535.

GOODMAN, J., AND S. MALKOC (2012): "Choosing Here and Now Versus There and Later: The Moderating Role of Psychological Distance on Assortment Size Preferences," *Journal of Consumer Research*, 39, 751–768.

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

HORTACSU, A., S. A. MADANIZADEH, AND S. L. PULLER (2015): "Power to Choose? An Analysis of Consumer Inertia in the Residential Electricity Market," *mimeo*.

HUBERMAN, G., S. S. IYENGAR, AND W. JIANG (2007): "Defined Contribution Pension Plans: Determinants of Participation and Contribution Rates," *Journal of Financial Services Research*, 31, 1–32.

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.

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

KREPS, D. M. (1979): "A Representation Theorem for 'Preference for Flexibility'," *Econometrica*, 47, 565–577.

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

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

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.

SPIEGLER, R. (2006): "Competition cver 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*.

Appendix: Proofs

We start with the following Lemma.

Lemma 1. *Assume that $\mu = 1$ and all markups are equal. Let $|A| = a$, $|B| = b$, and $|A \cap B| = c$. Then:*

1. *$\pi(A, B)$ is constant in c whenever $a = b$.*
2. *$\pi(A, B)$ is strictly increasing in c whenever $a > b$, and*
3. *$\pi(A, B)$ is strictly decreasing in c whenever $b > a$.*

Proof of Lemma 1:

Assume that $a = b$. Then, the profit of a firm offering A when the opponent offers B is $\frac{a-c}{a+b-c} + \frac{1}{2}\frac{c}{a+b-c} = \frac{2a-c}{2(a+b-c)}$. Denote this profit by $\bar{\pi}$. Note that $\frac{\partial \bar{\pi}}{\partial c} = \frac{a-b}{2(a+b-c)^2}$. Since $a = b$, it follows that $\frac{\partial \bar{\pi}}{\partial c} = 0$.

Next, suppose that $a > b$. Then, $\pi(A, B) = q(A)\bar{\pi}$. Note that $\frac{\partial \pi(A,B)}{\partial c} = \frac{\partial q(a)\bar{\pi}}{\partial c} = q(a)\frac{a-b}{2(a+b-c)^2}$. Since $a > b$, it follows that $\frac{\partial \pi(A,B)}{\partial c} > 0$.

Finally, assume that $b > a$. Then, $\pi(A, B) = (1 - q(B)) + q(B)\bar{\pi}$. Note that $\frac{\partial \pi(A,B)}{\partial c} = \frac{\partial q(B)\bar{\pi}}{\partial c} = q(B)\frac{a-b}{2(a+b-c)^2}$. Since $b > a$, we have $\frac{\partial \pi(A,B)}{\partial c} < 0$. ■

Proof of Proposition 1:

By assumption, $q(X) = 1$. Suppose to the contrary that (X, X) is not an equilibrium. In view of the symmetric payoffs it suffices to consider the profitable deviation for firm 1. Since (X, X) is not an equilibrium, there exists

$A \subset X$ such that $\pi_1(A,X) > \pi_1(X,X)$. We have

$$\begin{aligned}\pi_1(X,X) &= \frac{1}{2}\sum_{i\in I} p_X(x_i)w_i &= \frac{1}{2}\frac{1}{|X|}\sum_{i\in I} w_i \\ \pi_1(A,X) &= \frac{1}{2}\sum_{i\in I_A} p_X(x_i)w_i &= \frac{1}{2}\frac{1}{|X|}\sum_{i\in I_A} w_i\end{aligned}$$

Clearly, $A \subset X$ implies $\sum_{i\in I} w_i > \sum_{i\in I_A} w_i$ and therefore $\pi_1(X,X) > \pi_1(A,X)$, a contradiction.

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

$$\pi_1(A,B) = \sum_{i\in I_{A-B}} p_{AB}(x_i) + \frac{1}{2}\sum_{j\in I_{A\cap B}} p_{AB}(x_j) \tag{6}$$

$$\pi_1(X,B) = \sum_{i\in I_{X-B}} p_X(x_i) + \frac{1}{2}\sum_{j\in B} p_X(x_j) \tag{7}$$

By assumption, it holds that

$$\sum_{i\in A\setminus B} p_{AB}(x_i) + \frac{1}{2}\sum_{j\in A\cap B} p_{AB}(x_j) \geq \sum_{i\in I_{X\setminus B}} p_X(x_i) + \frac{1}{2}\sum_{j\in B} p_X(x_j) \tag{8}$$

From Assumption 1 we have $p_{EF}(x_i) = p_{EF}(x_j) = p_{EF}$ for all $E, F \in \mathcal{M}$ and all $x_i, x_j \in E \cup F$. Thus, we get

$$\sum_{i\in A\setminus B} p_{AB}(x_i) + \frac{1}{2}\sum_{j\in A\cap B} p_{AB}(x_j) = |A \setminus B| \cdot p_{AB} + \frac{1}{2}|A\cap B| \cdot p_{AB} \tag{9}$$

$$\sum_{i\in I_{X\setminus B}} p_X(x_i) + \frac{1}{2}\sum_{j\in B} p_X(x_j) = |X - B| \cdot p_X + \frac{1}{2}|B| \cdot p_X \tag{10}$$

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 this inequality is also wrong 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

$$\pi_1(A,B) = \frac{n-l}{m+n-l} + \frac{1}{2}\frac{l}{m+n-l} = \frac{2n-l}{2(m+n-l)} \tag{11}$$

$$\pi_1(X,B) = \frac{k-m}{k} + \frac{1}{2}\frac{m}{k} = \frac{2k-m}{2k} \tag{12}$$

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.

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

Proof of Proposition 2:

By assumption, all markups are equal. Without loss of generality, normalize them to equal 1. Consider first the case where $\mu = 1$. Given the symmetry of the payoff functions, it suffices to analyse the case of firm 1 only. Note that $\pi_1(X, X) = \frac{1}{2}$. Thus, (X, X) is an equilibrium if and only if $\pi_1(X, X) \geq \pi_1(A, X)$ for all $A \in \mathcal{M} \setminus \{X\}$. Since $\pi_1(A, X) = [1 - q(X)] + q(X)\left[\frac{1}{2}\frac{|A|}{k}\right]$, we have $\pi_1(X, X) \geq \pi_1(A, X)$ whenever $q(X) \geq \frac{k}{2k-|A|}$. Since $\frac{k}{2k-|A|}$ is increasing in $|A|$, it attains its maximum value of $\frac{k}{k+1}$ at $|A| = k - 1$. Thus, (X, X) is an equilibrium if and only if $q(X) \geq \frac{k}{k+1}$. Clearly, condition (4) need not hold for this to be true.

Now suppose $\mu = 0$. Note that, in this case, $\pi_1(X, X) = q(X)\frac{1}{2}$. Again, (X, X) is an equilibrium if and only if $\pi_1(X, X) \geq \pi_1(A, X)$ for all $A \in \mathcal{M} \setminus \{X\}$. This is equivalent to $\frac{q(X)}{2} \geq q(X)[\frac{1}{2}\frac{|A|}{k}] + q(A) - q(X)$ for all such A. Rearranging this we arrive at the condition

$$q(X) \geq \frac{2kq(A)}{3k - |A|} \qquad \text{for all } A \in \mathcal{M} \setminus X. \tag{13}$$

Now notice that $\frac{2kq(h)}{3k-h} \leq \frac{2kq(h-1)}{3k-h+1} \Leftrightarrow (3k - h + 1)q(h) \leq (3k - h)q(h - 1) \Leftrightarrow \frac{q(h)}{q(h-1)} \leq \frac{3k-h}{3k-h+1}$. Thus, the term on the right hand side of (13) is weakly decreasing in $|A|$ if and only if condition (4) is satisfied. Since this has been assumed to be the case, it follows that this term attains its maximum value of $\frac{2k}{3k-1}$ when $|A| = 1$. Therefore, (X, X) is an equilibrium if and only if $q(X) \geq \frac{2k}{3k-1}$. ■

Proof of Proposition 3:

As above, we normalize the weighted markups to 1. Consider first the case where no consumer defers ($\mu = 1$). Assume that both firms offer B such that $|B| = 1$. Since $q(1) = 1$ from Assumption 1, (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 $A \in \mathcal{M}$. The preceding argument shows that (B, B) such that $|B| = 1$ is an equilibrium if and only if q first-order stochastically dominates $\widehat{q}$. The proof for the case where $\mu \in [0, 1)$ is identical. ■

Proof of Corollary 4:

Once again, we normalize all markups to 1.

For the first part, recall that, by Proposition 3, a symmetric profile (B, B) with $|B| = 1$ is an equilibrium if and only if q first-order stochastically dominates the cdf $\widehat{q}$ defined by $\widehat{q}(A) = \frac{|A|}{2|A|-1}$. Since $q(A) = \frac{k+n-|A|}{k+n-1}$ by assumption, we must have $\frac{k+n-|A|}{k+n-1} \leq \frac{|A|}{2|A|-1}$ for all $A \in \mathcal{M}$. This holds if and only if $2|A|^2 - (k + n + 2)|A| + k + n \geq 0$ for all $A \in \mathcal{M}$. The latter inequality is satisfied if and only if $|A| = 1$ or $|A| \geq \frac{k+n}{2}$. Thus, (B, B) with $|B| = 1$ is an equilibrium if and only if $k + n \leq 2|A|$. Since $2|A|$ is strictly increasing in $|A|$ and attains its maximum value when $|A| = k$, it follows that (B, B) with $|B| = 1$ is an equilibrium if and only if $k + n \leq 2k$.

For the second part, recall that, by Proposition 2, the profile (X, X) is an equilibrium when $\mu = 1$ if and only if $q(X) \geq \frac{k}{k+1}$. Since overload is uniformly distributed over the interval $\{1, \ldots, k+n\}$ by assumption, $q(X) = \frac{n}{k+n-1}$. Now, it holds that $\frac{n}{k+n-1} \geq \frac{k}{k+1}$ if and only if $k^2 \leq k + n$. Finally, suppose $\mu = 0$. By Proposition 1 again, (X, X) is an equilibrium when $\mu = 0$ if and only if $\frac{n}{k+n-1} \geq \frac{2k}{3k-1} \Leftrightarrow 2k^2 - (n+2)k + n \leq 0 \Leftrightarrow k \leq \frac{n}{2} \Leftrightarrow k + n \geq 3k$, provided that (4) is satisfied. ∎

Proof of Proposition 5:

The proof applies to all $\mu \in [0, 1]$. Without loss of generality, suppose $A = \{x_1, \ldots, x_{l-1}\}$ is a set of $l - 1$ most profitable alternatives, and $w_1 \geq w_2 \geq \ldots \geq w_{l-1}$. We will show that (A, A) is an equilibrium. It suffices to show that there is no profitable deviation for firm 1. Suppose, per contra, that there is such a profitable deviation at menu $B \neq A$. If $|B| > l - 1$, then $\pi_1(B, A) = 0 < \pi_1(A, A)$. Moreover, it is clear that $B \subset A$ is not a profitable deviation either. Next, consider some $B \subseteq X - A$ with $|B| \leq l - 1$. We have

$$\begin{aligned} \pi_1(A, A) &= \frac{1}{|A|}\left(\frac{1}{2}\sum_{j=1}^{l-1} w_j\right) \\ \pi_1(B, A) &= \frac{1}{|A \cup B|}\left(\sum_{i \in I_B} w_i\right) \end{aligned}$$

Let $\overline{w}_j := \min_{j \in I_A} w_j$. It holds that $\pi_1(A, A) \geq \frac{1}{2} \cdot \overline{w}_j$.

Let $\widehat{w}_j := \max_{j \in I_B} w_j$. It holds that $\pi_1(B, A) \leq \frac{|B|}{|A \cup B|}\widehat{w}_j$.

It now suffices to show that

$$\frac{1}{2}\overline{w}_j \geq \frac{|B|}{|A \cup B|}\widehat{w}_j$$

Suppose not. Then $\overline{w}_j < \frac{2|B|}{|A \cup B|}\widehat{w}_j$. Since $\overline{w}_j \geq \widehat{w}_j$ by assumption, this is true only if $\frac{2|B|}{|A \cup B|} > 1$. But since $|A| = l - 1$, $|B| \leq l - 1$ and $|A \cap B| = 0$, this is impossible. A contradiction therefore obtains.

Finally, suppose $B \subseteq X$ is such that $A \cap B \neq \emptyset$, $A \neq B$ and $|B| \leq l - 1$. We have

$$\begin{aligned} \pi_1(A, A) &= \frac{1}{|A|}\left(\frac{1}{2}\sum_{j=1}^{l-1} w_j\right) \\ \pi_1(B, A) &= \frac{1}{|A \cup B|}\left(\sum_{i \in I_{B-A}} w_i + \frac{1}{2}\sum_{i \in I_{A \cap B}} w_i\right) \end{aligned}$$

By assumption, $|A| = l - 1$. Let $|A \cup B| = n$. By assumption, $n \in \{l, \ldots, 2l - 3\}$. Assume, per contra, that $\pi_1(B, A) > \pi_1(A, A)$. We have

$$\frac{1}{n}\left(\sum_{i \in I_{B-A}} w_i + \frac{1}{2}\sum_{i \in I_{A \cap B}} w_i\right) > \frac{1}{2(l-1)}\sum_{j=1}^{l-1} w_j \tag{14}$$

Let $\overline{w}_B := \max_{i \in I_{B-A}} w_i$ and $\underline{w}_A := \min_{j \in I_A} w_j$. Suppose $|B| = m$. There are $n - l + 1$ elements in the sum $\sum_{i \in I_{B-A}} w_i$ and $l + m - n - 1$ elements in the sum $\sum_{i \in I_{A \cap B}} w_i$. It follows then from (14) that

$$\frac{(n-l+1)}{n}\overline{w}_B + \frac{l+m-n-1}{2n}\underline{w}_A > \frac{1}{2(l-1)}(l-1)\underline{w}_A = \frac{1}{2}\underline{w}_A$$

This reduces to

$$\overline{w}_B > \frac{2n - l - m + 1}{2n - 2l + 2}\underline{w}_A. \tag{15}$$

Since $\underline{w}_A \geq \overline{w}_B$ by assumption, it follows from (15) that $2n - l - m + 1 < 2n - 2l + 2$, which implies $m = |B| > l - 1$. This is a contradiction. Therefore, no deviation to a menu $B \neq A$ is profitable when the opponent plays A.

Finally, suppose $\mu = 0$ and consider a symmetric profile (B, B) where $A \neq B$. If $|B| \leq |A|$, if follows from the arguments above that each firm profitably deviates to A. If $|B| > |A|$ instead, then, by assumption, all consumers choose nothing at (B, B), so that both firms receive a payoff of zero. Again, a unilateral deviation to A is profitable. Thus, (A, A) is the unique symmetric equilibrium in this case. ■