

Duopolistic Competition with Choice-Overloaded Consumers^{*}

Georgios Gerasimou[†]
University of St Andrews

Mauro Papi[‡]
University of Aberdeen

October 2, 2017

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. An explanation for the occurrence of this effect is that consumers 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 and establish some (im)possibility results and characterizations under a variety of assumptions. We also propose a measure of market effectiveness that may be thought of as a proxy for consumer welfare in this environment, and use it alongside our model to provide a critical perspective on regulations that cap the number of products that firms could offer.

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

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

^{*}Previous title: “*Oligopolistic Competition with Choice-Overloaded Consumers*”. We thank Luke Froeb, Paola Manzini, Patrick Rey, Ariel Rubinstein, Yuval Salant, Rann Smorodinsky, Hugo Sonnenschein, David Ulph, Nikolaos Vettas, Rakesh Vohra, audiences at Bounded Rationality in Choice (2016, Northwestern), Bounded Rationality & Mechanism Design (2016, Glasgow), St Andrews, Technion, Athens University of Economics & Business and especially two referees of this journal for helpful 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 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 Iyengar and Lepper (2000), which 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 by now has reached the status of being discussed in leading undergraduate microeconomics textbooks.¹

Choice overload 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; Kamenica, 2012). At least partly in response to the large body of empirical work that has followed the original experimental evidence on this phenomenon (Chernev et al, 2015 is a thorough survey), 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, which intervened and forced energy suppliers to ban complex tariffs and restrict the number of products they could offer to *no more than four* (Ofgem, 2013, 2014).

In addition to such novel consumer welfare considerations and concerns, 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 they possibly can, and offering menus with sufficiently few products in order not to overload consumers and lose them to their competitors or drive them out of the market altogether. This strategic trade-off lies at the heart of the novel model of duopolistic competition in menu design that we introduce and analyse in this study.

Our model assumes that consumers are heterogeneous in their preferences as well as in their overload characteristics, and that firms have complete information about 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 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 proxy, the *market effectiveness* index, which associates each market state/strategy profile with the *actual* number of products that is available to consumers once the effects of overload have been accounted for. As such, it is increasing in the number of products that are offered by the two firms and decreasing in the number of consumers who are overloaded at these menus.

¹See Varian (2014, p. 589), for example.

We first consider the benchmark special case of the model where no consumer is ever overloaded. In this context, 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. Moreover, although this is not the unique equilibrium when preferences are uniformly distributed and markups differ, it does become one in strictly dominant strategies when products are equi-profitable.

Retaining the assumption of uniformly distributed preferences and equal markups, we then move on to the identification of necessary and sufficient conditions for minimum-variety/minimum-differentiation equilibria to exist in which firms offer the same one product. We also establish an equivalence between the existence of such equilibria and ones where firms offer single but distinct products. We then demonstrate that full product differentiation equilibria with equally sized menus are impossible in this environment for higher numbers of products whenever the consumers' overload thresholds are either uniformly or geometrically/exponentially distributed. We further show that full product differentiation is impossible under any overload distribution when markups are allowed to differ and the most profitable product is sufficiently more profitable (not necessarily more than twice as much) relative to the second most profitable one. We conclude this section by focusing on the conditions under which symmetric profiles with three or four products (corresponding to cases that appear to have particular real-world relevance) can be obtained as equilibria under the above overload distributions. The testable empirical predictions of our model here include, for example, that in order for these equilibria to be possible it must be that roughly half of all consumers become overloaded in menus with more than six products.

We then move to the general version of our model where preferences are not assumed to be distributed in a particular way. In this more general context we first provide a general characterization of symmetric equilibria and show that whenever such equilibria exist, they take an essentially unique form, in terms of both the number and structure of the products that are offered by both firms. Finally, we raise and answer questions on the relation between equilibrium and market effectiveness within the class of symmetric profiles in this more general environment. In particular, we note that in many recent market models with boundedly rational consumers, firms have a strategic incentive to obfuscate –e.g. to enhance product variety by offering many substitute products; increase the dimensionality of fees; frame the offered products in a certain way.² In these models, greater complexity induces consumer confusion and, as a result, increases the probability that the consumer makes the wrong decision by choosing an inferior product, enabling firms to profit out of it (Spiegler, 2016). As mentioned above, however, increasing the complexity of the offered menu in the market we study has a negative effect on firms as well, because a choice-overloaded consumer that finds a menu offered by a firm too complex either purchases from the rival firm or defers choice altogether. Therefore, it is not obvious *a priori* whether, in our model, the market by itself is able to achieve the optimal (from the market-effectiveness point of view) degree of “complexity” or if a regulatory intervention is needed to correct the market failure. We investigate

²For example, Spiegler (2006) shows that in a market model populated by consumers with limited ability to understand complicated products firms employ complex pricing structures. Chiovenau and Zhou (2013), on the other hand, study a price-frame competition model and find that an increase in competition leads firms to increase frame complexity. A third example is given by Gabaix and Laibson (2006) which shows that the presence of myopic consumers induces firms to shroud relevant product information in equilibrium.

this issue by showing that under both uniformly and geometrically distributed overload, the symmetric profiles that maximize our notion of market effectiveness can be obtained as the free-market equilibrium outcomes only in the extreme cases where optimality is attained when the commonly offered menu contains either a single product or all products that a firm could possibly offer. Moreover, it turns out that in both the uniform and geometric cases the equilibrium potential of the optimal symmetric profile is destroyed by firms' profitable deviations to *smaller* menus. Therefore, a novel policy implication of our analysis is that, in such environments, interventions that are motivated by considerations such as those captured by our market effectiveness welfare proxy might improve upon the market outcome by imposing a *lower* rather than an upper bound on the number of products that firms can offer.

The paper is organized as follows. Section 2 lays out our duopolistic model and introduces the proposed index of market effectiveness. Section 3 discusses the benchmark case of fully rational consumers under various assumptions on preferences and markups. Sections 4 and 5 analyse the model under the assumption of uniformly distributed preferences and under general preferences, respectively. Section 6 discusses the related literature and Section 7 concludes. Appendix A provides additional details on the decision process that our consumers are assumed to follow, much of which coincides with the model of overload-constrained utility maximization that is introduced in Gerasimou (2018). Appendix B contains the proofs of all results that appear in the main body of the paper, and also the statements and proofs of some auxiliary results that may be of some independent interest. Finally, Appendix C provides four real-world contemporary market examples where firms compete in small menus of sizes that contain between three and five products.

2 The Model

2.1 General Setup

We consider a market with two firms and let $X := \{x_1, x_2, \dots, x_k\}$ 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 the game with finitely many strategies that is induced by our model, the existence of pure-strategy equilibria 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. We acknowledge, however, that our assumption is 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, particularly after relevant empirical evidence becomes available that could 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 (hence choosable) 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 does not belong to the set X but, from the point of view of one-period decision making, can be considered objectively 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 such cases our consumers are assumed to avoid/indefinitely defer choice by opting for their outside option. In line with the relevant empirical evidence that was discussed in the introduction, we let the complexity of a menu coincide with its cardinality.

More specifically, we assume that at every strategy profile (A, B) , each consumer first pre-scans 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, denoted $x^*(D)$.³ 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. Finally, if he finds both menus to be complex, the consumer buys nothing and opts for the deferral/outside option.

Let us elaborate with an example that emphasizes the sequential nature of the above decision process, with the latter laid out formally in Appendix A. 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 from one of the two providers that are available in his region (both of which were previously unknown to him) which phone & broadband package to buy. 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 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

³We assume that consumers' preferences are strict, which implies that $x^*(D)$ is always a singleton.

faced with all products that are available in the market. Therefore, the decision process that we have imposed is internally consistent.⁴

Consumers in our model are assumed to be heterogeneous in terms of their preferences as well as their overload characteristics. To capture preference heterogeneity we assume that the probability $p_A(x_i)$ of some product x_i being the consumer's most preferred alternative in menu A satisfies

$$0 < p_A(x_i) < 1. \quad (1)$$

This assumption imposes a strong form of preference heterogeneity by requiring that there be no menu where all consumers agree on which alternative is the best. In addition, we assume that the choice probability of product x_i when menus A and B are available is given by

$$p_{A \cup B}(x_i) = \frac{p_X(x_i)}{\sum_{x \in A \cup B} p_X(x)}. \quad (2)$$

Thus, 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 model due to Luce (1959).

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 $\{1, \dots, k+n\}$, where n is an integer that may be weakly positive or negative, and is restricted so that $k+n \geq 1$. 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.

We note that the consumers' preferences and cognitive characteristics/overload thresholds will be treated as independent random variables in our model. This assumption is not without loss of generality. However, one reason why it may be justified in our setting is the general and unstructured nature of the products in the set X . Specifically, if the dimensionality of the products' attributes had been explicitly modelled, then correlations between the two components of the model would also need to be allowed for. For example, of particular relevance there would be the cases where lower overload thresholds are associated with higher

⁴We 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 whenever there are just 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.

degrees of preference for simple products that can be described in few attributes. However, one of the primary objectives of our paper is to formalize the idea that choice overload is an interesting and relevant phenomenon in the context of firm competition. Hence, as a starting point of this analysis we have chosen to focus on the simpler case where a specific structure on the products' attributes is not imposed, and, as a result, no clear direction exists that a correlation between preferences and cognitive characteristics should be assumed to take.

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 A and its opponent offers B we let its baseline payoff that corresponds to the case where no consumer is overloaded at either of these menus be defined by

$$R_1(A, B) = \sum_{i \in I_{A \setminus B}} p_{A \cup B}(x_i) w_i + \frac{1}{2} \sum_{j \in I_{A \cap B}} p_{A \cup B}(x_j) w_j. \quad (3)$$

In words, each product x_i that is offered only by the first firm is associated with an expected payoff that is given by its markup, 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.⁵

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

$$\pi_1(A, B) = \begin{cases} q(A) \cdot R_1(A, B), & \text{if } |A| \geq |B| \\ q(B) \cdot R_1(A, B) + [q(A) - q(B)] \cdot \sum_{i \in I_A} p_A(x_i) w_i, & \text{if } |B| > |A| \end{cases} \quad (4)$$

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,

⁵Our formulation assumes that markups are independent of how many firms are offering the relevant products. A more general approach would have been to assume that markups decrease by some parameter $\alpha \in (0, 1]$ whenever both firms are offering those products. This would change the definition of the baseline payoff function to $R_1(A, B) = \sum_{i \in I_{A \setminus B}} p_{A \cup B}(x_i) w_i + \alpha \frac{1}{2} \sum_{j \in I_{A \cap B}} p_{A \cup B}(x_j) w_j$, and would convey the

basic intuition of the effects of increasing competition whereby a given product's profitability is decreasing in the number of firms offering it. Our current formulation here would be captured by the special case where $\alpha = 1$. We have chosen not to pursue this more general version of the model in this paper mainly because this parameter α would also need to be allowed to depend on the strategy profile (A, B) , and such context-dependence would in turn lead to a significant loss of tractability and discipline. We acknowledge, however, that such an extension of our model would be desirable, especially if it is informed by relevant empirical work that motivates specific assumptions on the values in the collection $\{\alpha_{A,B} : A, B \in \mathcal{M}\}$ of context-dependent competition-effect parameters.

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 product x_i in A is $p_A(x_i)$ and *not* $p_{A \cup B}(x_i)$. On the other hand, consumers who are not overloaded at either A or B will consider both menus and hence will 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.

This payoff function captures the key tradeoff between variety and complexity that firms are facing when consumers are potentially overloaded, which is the main motivating force for the present study. 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 Market Effectiveness and Consumer Welfare

Since our model abstracts from firms' pricing decisions, traditional welfare measures such as consumer surplus are not applicable in our setting. A key aspect of our model, however, is that some overloaded consumers may ultimately choose none of the products offered in the market even though, by our product-desirability assumption, choosing something would have been objectively better than choosing nothing. Moreover, even when overload does not lead to choice deferral because only one of the two available menus is complex for a given consumer, it still has an adverse effect on that consumer's welfare by essentially narrowing his market options. These facts motivate the need for some welfare measure that captures the complexity-variety tradeoff that arises from the consumers' point of view. One such measure that may be viewed as a proxy for consumer welfare in this setting could be defined by what we will refer to as the *market effectiveness index* $W : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}$, which associates each strategy profile (A, B) with the value

$$W(A, B) := \begin{cases} q(A) \cdot |A \cup B| + [q(B) - q(A)] \cdot |B|, & \text{if } |A| > |B| \\ q(A) \cdot |A \cup B| = q(B) \cdot |A \cup B|, & \text{if } |A| = |B| \\ q(B) \cdot |A \cup B| + [q(A) - q(B)] \cdot |A|, & \text{if } |A| < |B| \end{cases} \quad (5)$$

This welfare proxy measures how effective a particular market state (A, B) is from the consumers' point of view by mapping this state into the *actual* number of products that is

available to consumers once the effects of overload have been accounted for. Specifically, 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 market effectiveness index defined by (5) accounts for these forces by adding the total number of products and those in the least complex menu after these have been scaled down by the fractions of consumers who consider the former and the latter, respectively.⁶

We note that our index of market effectiveness 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 minimum when $q(A) = q(B) = 0$, i.e. when all consumers are overloaded at both available menus.

If all products are *equi-profitable* in the sense that their markups coincide and equal some $w > 0$, it follows from (4) 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, in this class of symmetric profiles market effectiveness is generally not maximized when firms offer the same singleton menu. We will return to this point later.

3 The Rational-Consumers Benchmark: Maximum-Variety Equilibrium

We first consider the benchmark case where no consumer is overloaded. This revolves around the equilibrium capacity of the *maximum-variety* symmetric profile (X, X) where each firm offers all products. In particular, this is true both under general preferences and equi-profitable products, i.e.

$$w_1 = \dots = w_k,$$

as well as under general markups and *uniformly distributed* preferences. The latter assumption amounts to

$$p_A(x) = \frac{1}{|A|}$$

for every menu A , and hence implies that all products in the market are always equally likely to be chosen.

⁶Given that, as we argued in footnote 5, our model abstracts from the possibility of markups decreasing whenever the relevant products are offered by both firms, it is also the case that the potentially beneficial effects to consumers from this kind of competition do not enter the market effectiveness index. We note, however, that a modification in the direction of allowing for such an effect can easily be made even without introducing markups as the function's arguments. Specifically, if the firms' payoffs are modified as described above by means of a collection $\{\alpha_{A,B} : A, B \in \mathcal{M}\}$, then (5) could be augmented by adding the term $(1 - \alpha_{A,B}) \cdot |A \cap B|$ to each of its three components, multiplied by $q(A)$, $q(A) = q(B)$ and $q(B)$, respectively. This formulation encompasses (5) whenever $\alpha_{A,B} = 1$ for all A, B , whereas if $\alpha_{A,B} < 1$, the index is increasing in the number of products offered by both firms and decreasing in $\alpha_{A,B}$, consistent with intuition.

Proposition 1.

Suppose no consumer is overloaded at X . Then:

1. Under general preferences and equal markups, (X, X) is the unique equilibrium.
2. Under uniformly distributed preferences and general markups, (X, X) is a strict but generally not unique equilibrium.
3. Under uniformly distributed preferences and equal markups, (X, X) is an equilibrium in strictly dominant strategies.

Indeed, consider first the case where preferences are general and all markups are equal and normalized to one. Conditional on a firm offering menu X , its opponent's payoff when itself offers X is $\frac{1}{2}$, while its payoff from offering any $A \subset X$ is strictly lower. Clearly, this is so because the firm would be foregoing a payoff of $\frac{1}{2} \cdot p_X(x_i) > 0$ for each $x_i \in X \setminus A$ and, by assumption, it would be unable to attract any overloaded consumers by offering the less complex menu A . Hence, (X, X) is a strict equilibrium. Uniqueness follows as a corollary to Proposition 5 that is stated and discussed in Section 5.

The argument showing that (X, X) is a strict equilibrium in the case of uniformly distributed preferences and general markups is analogous to the case above.⁷ The more interesting aspect of this statement, however, 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 (4) 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 harmed due to the associated decrease in all products' choice probabilities outweighing the gains from being the only firm that offers these low-markup products.

Finally, coming to the case where preferences are uniformly distributed and all products are equi-profitable, offering X becomes a strictly dominant strategy because, 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. Notably, 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 our market effectiveness index.

⁷We 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 .

4 Uniformly Distributed Preferences

In this section we analyse further the special case of our model where the consumers' preferences are uniformly distributed. This assumption allows us to keep the analysis simple and also makes it easier to identify the effects of choice overload on the menus offered by firms in equilibrium. One way of motivating it 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. We relax this assumption in the next section.

4.1 Minimum-Variety Equilibria


We now turn to the analysis of the opposite extreme case where the market equilibrates in a *minimum-variety* state in which the same one product is offered by both firms. It turns out that our characterization of this special case is also informative in relation to whether other types of equilibria may also be possible at the same time. Notation-wise, in what follows it is understood that A_i refers to a menu with i products.

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 either coincides with or first-order stochastically dominates the cdf \hat{q} defined by $\hat{q}(l) := \frac{l}{2l-1}$.
3. (A_h, A_h) is an equilibrium if and only if $h = 1$.
4. (A_1, B_1) with $A_1 \neq B_1$ is an equilibrium.

Figure 1: Overload cdf for minimum-variety equilibria



The equivalence between the first two statements suggests that in order 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 $l > 1$ be bounded above by the threshold $\hat{q}(l) = \frac{l}{2l-1}$. As also shown in Fig. 1, this threshold cdf \hat{q} features a sharp drop in the fraction of non-overloaded consumers when the menu size increases from one to two, and, being strictly convex, is also such that further drops are decreasing in subsequent menu-size increments. An example of an overload distribution that satisfies this condition is the geometric/exponential cdf defined by $q(h) := p^{h-1}$, provided that $p < \frac{2}{3}$ (Fig. 1). Notably, the value of the target cdf \hat{q} as the menu size increases converges to $\frac{1}{2}$ and not to 0. Therefore, as far as the existence of minimum-variety equilibria is concerned, whether more than half of the consumers eventually become overloaded as the menu complexity increases is irrelevant.

Proposition 2 further clarifies that for minimum-variety equilibria to occur it is necessary and sufficient that symmetric equilibria of higher menu complexity do not exist. Intuitively, the above condition on q shows that for minimum-variety profiles to be equilibria it is necessary that overload increase sharply as menu size goes above one. This, in turn, would make it profitable for a firm to deviate to a menu of smaller size whenever the starting point is some profile (A_h, A_h) with $h > 1$, as this would enable it to absorb a proportion of overloaded consumers that is sufficiently high to offset the revenue that would have been lost by offering fewer products. Conversely, for a higher-complexity symmetric profile (A_h, A_h) to be an equilibrium such deviations to smaller menus must be unprofitable, which necessitates more gradual overload increases in the consumer population than those necessary for (A_1, A_1) to be an equilibrium.

Proposition 2 finally demonstrates that minimum-variety equilibria are possible if and only if singleton-menu *full product differentiation* profiles (A_1, B_1) with $A_1 \neq B_1$ are also equilibria. Intuitively, the firms' profits are the same across these two equilibrium classes and, even though this is due to different reasons⁸, it drives this dual equilibrium property. Moreover, due to equilibrium profits between (A_1, A_1) and (A_1, B_1) being equal, and the fact that the value of the market effectiveness index (i.e. our proxy for consumer welfare) is twice as high when two distinct products are offered in the market (indeed, both are considered by *all* consumers), equilibria in the latter class *Pareto-dominate* the minimum-variety ones.

4.2 Full Product Differentiation

We now turn to the analysis of more general full product differentiation profiles (A, B) in which A and B are disjoint menus. Allowing markups to differ, our next result identifies a sufficient condition for such equilibria 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

⁸For profile (A_1, A_1) the $\frac{w}{2}$ equilibrium 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 factor, together with the existence of two distinct products in the market, each of which is offered uniquely by each firm and is therefore associated with a choice probability of $\frac{1}{2}$.

there is no overload cdf q under which a full product differentiation equilibrium exists.

The noteworthy aspect of this result is not the fact that a full product differentiation equilibrium is impossible when the highest markup is sufficiently higher than the rest. Rather, what is interesting here is that the factor by which this markup needs to be higher than the second-largest markup for such a general impossibility to come about is bounded above by 2. Whenever the stated condition is satisfied, both firms will choose to offer a menu that necessarily includes the most profitable alternative because all potential gains from full product differentiation would be eliminated, making the shape of the overload distribution irrelevant. Moreover, the condition is such that the distance between the highest- and second-highest markups depends positively on the total number of products. Intuitively, when there are few products that firms can choose from, all 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 profitable product to stay relatively close, in profitability terms, 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 firms in *every* equilibrium its markup must be sufficiently high to offset the lower payoff that is associated with both firms offering x_1 when many other products could have been offered instead.

Turning now to the case where all markups are equal, our next result establishes that non-singleton full product differentiation equilibria where both firms offer the same number of products are impossible under a large class of overload distributions.

Proposition 4.


If all products are equi-profitable and overload is uniformly or geometrically distributed, then a full product differentiation equilibrium (A_m, B_m) does not exist for any $m \geq 2$.

The assumption of uniformly distributed overload is a natural theoretical benchmark. It implies that $q(h) = \frac{k+n+1-h}{k+h}$, where, as previously noted, $k + n$ stands for the menu size of the least overloaded consumer. Since the parameter n is allowed to be positive or negative, Proposition 4 makes the somewhat surprising prediction that full product differentiation equilibria with equally sized menus are impossible in such linear environments, regardless of how fast or slow consumers become overloaded (with the speed determined by $k + n$) as complexity increases. The assumption of geometrically/exponentially distributed overload is also a natural theoretical benchmark and, as already mentioned, it implies that the overload cdf q is defined by $q(h) = p^{h-1}$, with $p \in (0, 1)$. In our framework, this family of distributions encompasses a wide range of situations that range from such extreme ones where p is so close to 1 that consumers are practically never overloaded, to those where p is so close to 0 that essentially all consumers are overloaded in menus with just two products.

Together, Propositions 2 and 4 imply that the only situation where a full product differentiation equilibrium with equally sized menus is possible in this environment is the degenerate

one in which these menus are singletons.⁹ Given that all markups and choice probabilities are the same in this environment, our model makes the strong prediction that the pure effect of choice overload in this market rules out full product differentiation equilibria under a wide range of interesting cases. Intuitively, for such profiles (A_m, B_m) to be equilibria the overload cdf q must be such that the ratio $\frac{q(m)}{q(m-h)}$ is strictly decreasing at a certain increasing rate (for h both positive and negative), and, due to the fact that the payoff function changes as the deviating firm's menu becomes the relatively more complex one of the two, this ratio (a function of h) features a kink at $h = 0$.

Figure 2: Impossibility of a full product differentiation equilibrium
(example of $\frac{q(m)}{q(m-h)}$ when $m = 4$)



While $\frac{q(m)}{q(m-h)}$ is indeed strictly decreasing under both the uniform and geometric distributions, the rate at which consumers become more or less overloaded as the menu size of the firm that deviates from profile (A_m, B_m) goes above or below m , respectively, is not sufficiently high for all possible deviations. Fig. 2 illustrates this in the case of $m = 4$, where deviations in lower-complexity menus ($h > 0$) are not profitable but those in higher-complexity menus ($h < 0$) are.

4.3 Three- and Four-Product Symmetric Equilibria

In Appendix C we present four real-world contemporary examples from the subscription TV, broadband, retail banking and laptop computer industries where, in each case, two competing firms offer menus that contain a minimum of three and a maximum of five products. Notably, in each of these examples at least one firm is offering a menu with just three products, while in two cases (subscription TV contracts and high-end laptops) all four firms have done so. Several additional real-world market examples that attest to the focality of this number can also be reported. Although the examples that we provide are such that the sets of three products offered by each firm are not identical, their degree of substitutability can neverthe-

⁹As already noted, if overload is geometrically distributed this happens whenever $p < \frac{2}{3}$, while with uniform overload this is possible whenever $k + n \leq 4$.

less be considered rather high. To the extent that this is so, and given the impossibility of equal-menu-size full product differentiation equilibria that we established earlier, such three-product market outcomes can be modelled in our framework as symmetric profiles (A_3, A_3) .

Be that as it may, a question that one could raise and answer using our model is: Which specific overload distributions can sustain this symmetric three-product-menu equilibrium? We answer this question below for the case of uniformly and geometrically distributed overload.¹⁰

Observation 1.

Suppose all products are equi-profitable. Then:

1. (A_3, A_3) is an equilibrium under uniformly distributed overload if and only if $6 \leq k + n \leq 10$.
2. (A_3, A_3) is an equilibrium under geometrically distributed overload if and only if $p \in [\frac{4}{5}, \frac{7}{8}]$.

We note next that another case worth considering along these lines is the one where both firms offer the same menu of *four* products. This number is relevant, for example, because, in addition to being featured in a number of relevant real-world examples, it was, as already mentioned, also the number that was chosen by the UK energy market regulator, Ofgem, as the upper bound on the number of gas and electricity tariffs that firms operating in this industry could offer to consumers.


Observation 2.

Suppose all products are equi-profitable. Then:

1. (A_4, A_4) is an equilibrium under uniformly distributed overload if and only if $9 \leq k + n \leq 13$.
2. (A_4, A_4) is an equilibrium under geometrically distributed overload if and only if $p \in [\frac{6}{7}, \frac{9}{10}]$.

Figure 3: Overload distributions for three- and four-product symmetric equilibria

¹⁰The simple computational verification is left to the reader.



The class of such observations that are generated by our model lead to testable empirical hypotheses.¹¹ In particular, under the maintained assumption that firms in actual markets are indeed at an equilibrium state when offering three or four products, and that this number is the same for all firms, it follows from the above two observations that our model makes rather sharp predictions on how overload evolves. In particular, it suggests that the threshold menu size of the least overloaded consumer is either 9 or 10 if overload is uniformly distributed, and that the exponential parameter p is between $\frac{6}{7}$ and $\frac{7}{8}$ if overload is geometrically distributed.

As Fig. 3 illustrates in the upper bound case of this “intersection” where $k + n = 10$ and $p = \frac{7}{8}$, these distributions almost coincide in their overload predictions as the menu size increases to contain up to six products, while these predictions diverge significantly beyond that point. Notably, however, both distributions predict that approximately half of the consumers are overloaded in menus with more than six products. The latter therefore emerges as a novel insight and a potentially key property for such symmetric profiles to be equilibria in general. Empirical work that would shed light on the validity of these predictions in real markets would be particularly welcome and would allow for further modelling refinements.

5 General Preferences

We now relax the assumption that the consumers’ preferences are uniformly distributed. In particular, throughout this section we assume that products are equi-profitable and focus on symmetric profiles without imposing any restriction on consumers’ preferences other than the underlying requirement that all products’ choice probabilities be strictly positive. While the assumption that all markups are equal is somewhat limiting, it could be justified on the grounds that it allows for a fairly general analysis to be made in a relatively simple way, and also on the basis that preferences and markups play a similar (though, as explained in detail

¹¹Observations 1 and 2 could also be extended to the case of general preferences and equal markups that we consider in Section 4.

later, not identical/substitutable) role in our model.

5.1 Existence and Uniqueness of Symmetric Equilibria

We begin our analysis of the general-preference version of our model with the following characterization.

Proposition 5.

A symmetric profile (A, A) is an equilibrium if and only if the following hold:

1. *A contains $|A|$ most preferred products in X ;*
2. *$q(B)$ is sufficiently small relative to $q(A)$ for any menu B such that $|B| > |A|$;*
3. *$q(C)$ is not too large relative to $q(A)$ for any menu C such that $|C| < |A|$.¹²*

Moreover, if (A, A) is an equilibrium, then (B, B) is also an equilibrium if and only if $|A| = |B|$ and B contains $|B|$ most preferred products in X .

This characterization establishes that, in addition to containing $|A|$ most preferred products, in order for a menu A to be a symmetric equilibrium strategy it must be the case that sufficiently many consumers are overloaded at menus of sizes different from A relative to the fraction of consumers that are overloaded at A . The reason is that if the fraction of non-overloaded consumers is high enough at menus of size different from A , then firms have an incentive to deviate regardless of what the consumers' preferences are. It turns out, interestingly, that whenever offering a menu of a certain size is a symmetric equilibrium strategy, then offering a menu of that size (and with as many most preferred products) is the unique equilibrium strategy. The intuition is simple and, by now, familiar: assuming that a menu A is a symmetric equilibrium strategy is equivalent to requiring that sufficiently many consumers are overloaded at some simpler or more complex menu C relative to A . However, if this is the case, then C can never be an equilibrium strategy because the fraction of non-overloaded consumers at A relative to the fraction of non-overloaded consumers at C is too high for it to be the case. We note, finally, that if offering a menu A is a symmetric equilibrium strategy and there are unique $|A|$ most preferred products in X , then (A, A) is the unique symmetric equilibrium in general.

As was anticipated in the discussion of Proposition 1, a corollary to Proposition 5 is that if no consumer is overloaded at X , then (X, X) is the unique equilibrium. Indeed, it holds in this case that $q(C) = q(X)$ for all $C \subseteq X$, hence the third condition of Proposition 5 is always satisfied, while the first two conditions are trivially satisfied too.

5.2 Symmetric Equilibria and Market Effectiveness

We now focus on the uniform and geometric families of overload distributions that we also considered in the previous section, and raise the following questions: What is the "consumer-

¹²See Lemma 3 in Appendix B for a precise formulation of conditions 2 and 3.

optimal” menu size in each of these cases? Does it differ from the equilibrium menu size? If so, what kind of quantity regulation should a policy maker impose?

Throughout this subsection we denote by A^* a menu in the class of symmetric profiles (A, A) such that (A^*, A^*) maximizes the market effectiveness index. We will call such a menu *consumer-optimal* or simply *optimal*. Unless stated otherwise, A^* is assumed to contain $|A^*|$ most preferred products in X . Also, given some real number r , we denote by $\lceil r \rceil$ the value of the nearest integer function, where half-integers are rounded up to the closest higher integer number.

We first note that it is a simple implication of the definition of the market effectiveness index W and the uniform overload cdf $q(h) = \frac{k+n+1-h}{k+n}$ that

$$|A^*| = \begin{cases} \lceil \frac{k+n+1}{2} \rceil, & \text{if } \lceil \frac{k+n+1}{2} \rceil \leq k \\ k, & \text{otherwise} \end{cases} \quad (6)$$

We can now state the following result.

Proposition 6.

Assume that overload is uniformly distributed. Then:

1. (A^*, A^*) is an equilibrium if and only if $|A^*| = 1$ or $|A^*| = k$ and $k + n$ is sufficiently large;
2. (A^*, A^*) is not an equilibrium if and only if firms can profitably deviate to simpler menus.

Proposition 6 contains two economically relevant messages. First, a symmetric profile that maximises market effectiveness is an equilibrium in this environment only in the polar cases in which minimum- or maximum-variety profiles maximise market effectiveness. In the latter case, the menu-size threshold at which the least overloaded consumers actively choose (hence the parameter n too) has to be sufficiently high to ensure that offering the most complex menu of all is an equilibrium strategy. The reason is that if a maximum-variety profile is consumer-optimal and the fraction of non-overloaded consumers is not high enough, then firms can profitably deviate to simpler menus.¹³

Second, except in the extreme cases discussed above, a symmetric profile that maximises market effectiveness is *never* attained as an equilibrium, as firms can *always* profitably deviate to simpler menus. Moreover, Proposition 6 also suggests that deviating to more complex menus is never profitable. Hence, in equilibrium, the degree of product variety is smaller compared to what would be optimal for consumers. The intuition behind this is two-fold. On the one hand, under uniform overload the fraction of non-overloaded consumers at optimal or simpler menus is large enough (i.e., at least $\frac{1}{2}$) to make deviations to more complex menus unprofitable. On the other hand, the fraction of non-overloaded consumers at simpler menus relative to the fraction of non-overloaded consumers at the optimal menu is too large

¹³For example, assume that $X = \{x_1, x_2, x_3\}$ and $n = 2$. Then, the optimal menu size is 3. Assume that choice probabilities are $p_X(x_1) = 0.4$, $p_X(x_2) = 0.35$, and $p_X(x_3) = 0.25$. Notice that $\pi_1(X, X) = 0.3 < \pi_1(\{x_1\}, X) = 0.52$.

for the latter to be offered in a symmetric equilibrium. This is so because, under uniform overload, the rate at which the fraction of non-overloaded consumer increases as complexity decreases is too high compared to what would be the optimal amount. Therefore, assuming that overload is uniformly distributed and the market environment is such that only symmetric equilibria are possible, a policy implication of Proposition 6 is that a regulatory intervention should impose a *lower bound* on product variety.

We now turn to the case where overload is geometrically distributed and hence $q(h) = p^{h-1}$, $p \in (0, 1)$. We first note that a joint implication of this assumption and the definition of the market effectiveness index W is that

$$|A^*| = \begin{cases} 1, & \text{if } p \leq \frac{1}{e} \\ \left[-\frac{1}{\ln(p)}\right], & \text{if } p \in \left[\frac{1}{e}, \frac{1}{e^{1/k}}\right] \\ k, & \text{otherwise.}^{14} \end{cases} \quad (7)$$

The following result can now be stated.

Proposition 7.

Assume that overload is geometrically distributed. Then:

1. *If $|A^*| = 1$, (A^*, A^*) is an equilibrium;*
2. *If (A^*, A^*) is an equilibrium, either $|A^*| = 1$ or $|A^*| = k$;*
3. *(A^*, A^*) is not an equilibrium if and only if firms can profitably deviate to simpler menus.*

In the uniform overload case, increasing menu complexity results in the fraction of overloaded consumers to increase linearly. On the contrary, under geometric overload, the fraction of overloaded consumers grows exponentially as menu complexity increases. Despite this difference, and similar to the impossibility result that we obtained in Proposition 4 under these two families of distributions, the market-effectiveness results under geometric overload are analogous from a qualitative point of view to those under uniform overload. That is, offering an optimal menu is an equilibrium strategy only in the polar cases in which the optimal menu size is either 1 or k .

However, unlike in the uniform case, whenever the optimal menu size is k , increasing the menu-size threshold $k + n$ of the least overloaded consumers does not make unprofitable any deviations to menus simpler than the maximum-variety menu. The reason is that, under geometric overload, a change in $k + n$ does not affect the shape of the overload distribution. In addition, it can be shown that when the parameter p is sufficiently small, offering a maximum-variety menu is an equilibrium strategy whenever it is consumer-optimal to do so.¹⁵ However, if p is large enough, then whether offering a maximum-variety menu is an equilibrium depends on consumers' preferences.¹⁶

¹⁵Details are available from the authors upon request.

¹⁶An example is provided at the end of Appendix B.

Proposition 7 also suggests that, like in the uniform case, firms can profitably deviate from a consumer-optimal menu only by offering a simpler menu. The reasons are the same as those under the uniform case: (i) the fraction of consumers that are not overloaded at more complex menus is too small to make it profitable for firms to increase product variety; (ii) too many consumers are not overloaded at menus simpler than the optimal menu. We draw two economic implications. First, assuming that only symmetric equilibria are possible, the market by itself is able to achieve consumer optimality only in extreme cases. Second, in all those other cases where the market fails to do so, a policy intervention can improve upon the market outcome by imposing a *lower bound* on product variety.

This novel lower-cap policy implication that our model delivers in these environments may seem counter-intuitive. However, we note that this is partly driven by our building-block assumption (which, as was discussed, is backed by a large body of relevant empirical work) that consumers who are overloaded by both firms are lost to the market and thus yield zero profit to firms. This exercises a downward pressure on the size of the menus that firms choose to offer. In the above two cases it turns out that it is only in extreme cases that this force is sufficiently strong for the symmetric optima to be attainable as equilibria. Yet, we note that there is no reason to expect that this would continue to be the policy message if the payoff and welfare proxy/market effectiveness functions had been specified differently so as to capture, for example, the intuition that consumers choose at random instead of deferring choice altogether as menu complexity increases.

5.3 On the Role of Preferences and Markups in the Model

We conclude this section with a qualitative comparison of our analysis under the uniform preferences/general markups and general preferences/equal markups frameworks that we have considered in this paper. We will argue that this dichotomy is non-redundant and allows us to gain insights from different perspectives on the market we study.

In particular, assuming a given and fixed overload distribution, it follows from our analysis that even though a unique level of menu/market complexity may arise in symmetric equilibrium when preferences are unrestricted (cf Proposition 6 for the case of uniformly distributed overload), multiple symmetric equilibria –with varying menu sizes– may exist when markups are unrestricted instead (cf Proposition 1 in the case where overload is distributed in such a way that $q(X) = 1$). Indeed, consider the example where $X = \{x_1, x_2, x_3, x_4\}$ and suppose that no consumer is overloaded at X . As in the non-unique equilibrium example that was discussed after the statement of Proposition 1, assume first that markups are such that $w_1 = w_2 = w_3 = 1$ and $w_4 = 1/10$, and that choice probabilities $p_x(x_i)$ are equal for all $i \leq 4$. In this case, both (X, X) and $(\{x_1, x_2, x_3\}, \{x_1, x_2, x_3\})$ are equilibria. Next, suppose that markups are such that $w_1 = w_2 = w_3 = w_4 = 1$. Notice that offering X is a best response to the opponent offering any menu, regardless of how choice probabilities are structured. Therefore, (X, X) is the unique symmetric equilibrium.

The above remarks suggest that if we were to define a (possibly multi-valued) mapping from the set of symmetric-equilibrium market complexity levels into the set of overload distributions, then we would say that such a mapping is injective only when preferences are

allowed to remain general/unrestricted. Indeed, unlike in this case, the same overload distribution may correspond to different equilibrium complexity levels when markups are heterogeneous. Hence, although similar, markups and preferences do indeed play distinct roles in our model, with markup heterogeneity often encompassing more degrees of freedom.

Further to the above, the example below illustrates the strong sense in which preferences and markups in our model are not substitutable. Suppose $X = \{x_1, x_2, x_3, x_4\}$ and let the overload distribution q be such that $q(1) = 1, q(2) = 0.9, q(3) = 0.8$ and $q(4) = 0.7$. Consider first the case where markups are equal and such that $w_1 = w_2 = w_3 = w_4 = 0.25$ while preferences are heterogeneous and such that $p_x(x_1) = p_x(x_2) = p_x(x_3) = 0.3, p_x(x_4) = 0.1$. It holds that $\pi_1(X, X) = 0.0875, \pi_1(\{x_1, x_2\}, X) = 0.08375, \pi_1(\{x_1, x_2, x_3\}, X) = 0.07$, etc. Hence, (X, X) is an equilibrium. On the other hand, $\pi_1(\{x_1, x_2\}, \{x_1, x_2\}) = 0.1125$ and $\pi_1(X, \{x_1, x_2\}) = 0.12$. Hence, $(\{x_1, x_2\}, \{x_1, x_2\})$ is not an equilibrium. Consider now the symmetric case of the example in which the above non-uniform preference structure is replaced by an identical heterogeneous-markup structure, i.e. $w_1 = w_2 = w_3 = 0.3, w_4 = 0.1$, and assume also that choice probabilities are equal: $p_x(x_1) = p_x(x_2) = p_x(x_3) = p_x(x_4) = 0.25$. It now holds that $\pi_1(X, X) = 0.0875, \pi_1(\{x_1, x_2, x_3\}, X) = 0.1125$ and, in addition, $\pi_1(\{x_1, x_2\}, \{x_1, x_2\}) = 0.27, \pi_1(\{x_1, x_2, x_3\}, \{x_1, x_2\}) = 0.24, \pi_1(X, \{x_1, x_2\}) = 0.18$, etc. Therefore, in sharp contrast to the predictions made previously, $(\{x_1, x_2\}, \{x_1, x_2\})$ is an equilibrium here but not (X, X) .

6 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, 2017), *bounded-rational expectations* (Gabaix and Laibson, 2006; Spiegler, 2006), *comparison difficulty* (Piccione and Spiegler, 2012; Bachi and Spiegler, 2014; Papi, 2014) and *satisficing* (Papi, 2017).¹⁷ 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. Bachi and Spiegler (2014) propose a class of duopolistic models where consumers are presented with two-attribute products in Euclidean

¹⁷A textbook treatment of this literature is provided in Spiegler (2011), while Spiegler (2016) surveys and synthesizes some more recent developments and trends.

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. Although choice overload as analyzed here is not a source of deferral for consumers in the Bachi-Spiegler model, that paper studies the effects that indecisiveness-driven deferral¹⁸ 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.¹⁹ By contrast, consumer demand in our setting is not determined by some random utility model and the outside option 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 (2018) 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), Tyson (2008), Kuksov and Villas-Boas (2010), Ortoleva (2013), Buturak and Evren (2018), Frick (2016) and Dean et al (2017). The reasons for such behavior vary across these models and include regret as well as cognitive costs/attention constraints. Finally, the reader is referred 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.

7 Concluding Remarks

When oligopolistically competitive firms sell their products to consumers who are potentially choice-overloaded in the sense that they avoid/indefinitely defer making an active 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

¹⁸For a revealed-preference analysis of indecisiveness-driven choice deferral that is caused by the absence, respectively, of a partially or totally dominant option due to preference incompleteness, the reader is referred to Gerasimou (2016, 2018).

¹⁹Gerasimou (2018) also provides a revealed-preference analysis of undesirability/unattractiveness of the alternatives as a source of choice deferral in a deterministic setting.

consumers' tastes, but at the same time potentially overloads these consumers and hence either loses them to its rival or drives them out of the market altogether. This paper proposes the first model in the literature that aims to provide a framework for thinking about this strategic trade-off in a simple and general way.

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, due to cognitive constraints, a potentially significant fraction of the consumer population may not actually benefit from the products that are made available through the equilibrium market outcome. Our paper proposes a simple notion of market effectiveness that may be viewed as a proxy for consumer welfare in such an environment, and which may be relevant in this regard. Using this measure alongside our market model enabled us to analyse the equilibrium potential in the class of consumer-optimal symmetric strategy profiles and, perhaps surprisingly, arrive at the conclusion that, under uniformly or geometrically/exponentially distributed overload, a policy intervention can improve upon the market outcome by imposing a lower rather than an upper bound on the number of products that firms can offer.

Although our model assumes that consumers are not already endowed with a default market alternative, a potentially interesting extension would feature a formulation that would allow for such defaults to exist and influence the consumers' (hence the firms') decisions accordingly. 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 (2018): "Choice Overload and Asymmetric Regret," *Theoretical Economics*, forthcoming.
- 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.
- CHIOVENAU, I., AND J. ZHOU (2013): "Price Competition with Consumer Confusion," *Management Science*, 59(11), 2450–2469.
- 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 (2017): "Limited Attention and Status Quo Bias," *Journal of Economic Theory*, 169, 93–127.
- ELIAZ, K., AND R. SPIEGLER (2011): "Consideration Sets and Competitive Marketing," *Review of Economic Studies*, 78, 235–262.
- 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. (2016): "Asymmetric Dominance, Deferral and Status Quo Bias in a Behavioral Model of Choice," *Theory and Decision*, 80, 295–312.
- (2018): "Indecisiveness, Undesirability and Overload Revealed Through Rational Choice Deferral," *Economic Journal*, forthcoming.
- 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 (2017): "Competing for Attention: Is the Showiest also the Best?," *Economic Journal*, forthcoming.
- 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.
- (2017): "Price Competition with Satisficing Consumers," *International Journal of Industrial Organization (Special Issue on 'Economics of Consumer Search')*, forthcoming.
- 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 Over Agents with Boundedly Rational Expectations," *Theoretical Economics*, 1, 207–231.
- (2011): *Bounded Rationality and Industrial Organization*. New York: Oxford University Press.
- (2016): "Choice Complexity and Market Competition," *Annual Review of Economics*, 8, 1–25.
- 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.

Appendices

A The Consumer's Decision Process

As discussed in the main text, our consumers are assumed to face cognitive costs during their decision-making process which, when high enough, induce overload-driven choice deferral. In particular, we assume that their behaviour is described by a special case of the overload-constrained utility maximization model that is developed and axiomatically characterized in Gerasimou (2018). In this special case of the model the consumer's within-menu behavior is guided by a utility function $u : X \rightarrow \mathbb{R}$ 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 ,

$$C(A) = \begin{cases} \arg \max_{x \in A} u(x), & \text{iff } |A| \leq n \\ \emptyset, & \text{otherwise} \end{cases} \quad (8)$$

In our model we assumed that preferences were strict, which translates into the utility function u being injective.

Turning to the consumer's between-menu decision in the duopolistic market that we consider, let us define $\mathcal{D} := \mathcal{M} \times \mathcal{M} = \{(A, B) : A, B \in \mathcal{M}\}$ as the collection of all pairs of menus in \mathcal{M} , and let $\mathcal{M}^* := \mathcal{M} \cup \emptyset$. Consistent with the way in which the model in the main text was set up and motivated, the consumer's active choice or deferral in this class of decision problems is captured by the mapping $S : \mathcal{D} \rightarrow \mathcal{M}^* \cup \mathcal{D}$ that is defined by

$$S(A, B) = \begin{cases} A, & \text{if } |A| \leq n < |B|, \text{ or} \\ & |A|, |B| \leq n \ \& \ \max_{x \in A} u(x) > \max_{x \in B} u(x) \\ B, & \text{if } |B| \leq n < |A|, \text{ or} \\ & |A|, |B| \leq n \ \& \ \max_{x \in B} u(x) > \max_{x \in A} u(x) \\ \{A, B\}, & \text{if } |A|, |B| \leq n \ \& \ \max_{x \in A} u(x) = \max_{x \in B} u(x) \\ \emptyset, & \text{if } |A|, |B| > n \end{cases} \quad (9)$$

In words, the consumer chooses menu A (offered by firm 1) over B (offered by firm 2) either if B is considered complex and A is not, or if both A and B are non-complex but the agent's most preferred option is in A . 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. Finally, the consumer defers when both A and B are complex.

This specification effectively 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, as we also highlighted in the main text, 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 at some point in time, first according to the cardinality criterion and, if it is considered to be non-complex, then and only then as employing his preferences. He is also assumed to do the same with menu B at a later point. If both menus are non-complex, he compares the two best options in each of them and chooses his menu accordingly. Therefore, he is never confronted with $A \cup B$ as a single menu.

To capture the implications of the above deterministic consumer decision process for aggregate market shares, and thus to motivate the payoff function that is laid out in (4), let $P(x|A, B)$ stand for the choice probability of x when A and B are offered by the two firms. In light of the above, it holds that

$$P(x|A, B) = \begin{cases} \frac{1}{2} \cdot q(A) \cdot p_{A \cup B}(x) + [q(B) - q(A)]p_B(x), & \text{if } x \in A \cap B \ \& \ |A| > |B| \\ \frac{1}{2} \cdot q(B) \cdot p_{A \cup B}(x) + [q(A) - q(B)]p_A(x), & \text{if } x \in A \cap B \ \& \ |B| > |A| \\ \frac{1}{2} \cdot q(A) \cdot p_{A \cup B}(x) = q(B) \cdot \frac{1}{2} \cdot p_{A \cup B}(x), & \text{if } x \in A \cap B \ \& \ |A| = |B| \\ q(A) \cdot p_{A \cup B}(x), & \text{if } x \in A \setminus B \\ q(B) \cdot p_{A \cup B}(x), & \text{if } x \in B \setminus A \end{cases}$$

Therefore, as the above equations suggest, due to the fact that the consumers' outside option is not modelled as an element of X , the choice probabilities of the elements in $A \cup B$ do not sum to 1 whenever $q(A) < 1$ or $q(B) < 1$. This is a reflection of the effect that choice overload has in our market.

B Proofs

Lemma 1.

Assume that preferences are uniformly distributed and all markups are equal. 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 of the proposition is straightforward and omitted (see also the discussion in the main text). We will prove the second part. In particular, we will show that if preferences are uniformly distributed and all markups are equal, then, for any $A, B \in \mathcal{M}$, $\pi_1(X, B) > \pi_1(A, B)$. Without loss of generality, normalize the common 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

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

and

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

Moreover, by assumption, $p_{E \cup F}(x_i) = p_{E \cup F}(x_j) := \frac{1}{|E \cup F|}$ for all $E, F \in \mathcal{M}$ and all $x_i, x_j \in E \cup F$. Therefore,

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

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. \blacksquare

Proof of Proposition 2.

1 \Leftrightarrow 2. The common markup is normalized to 1 throughout the proof. Assume that both firms offer the same menu A_1 , where $|A_1| = 1$. Since $q(1) = 1$ by assumption, (A_1, A_1) is an equilibrium if and only if $\pi_1(A, A_1) \leq \frac{1}{2}$ for every $A \in \mathcal{M}$. Note first that $\pi_1(A, A_1) = \frac{1}{2} = \pi_1(A_1, A_1)$ for all $A \in \mathcal{M}$ such that $|A| = 1$. Suppose a firm deviates to $A \in \mathcal{M}$ such that $|A| \in \{2, \dots, k\}$. By Lemma 1, profits from this deviation are maximized whenever $|A \cap A_1|$ is maximized. Hence, without loss of generality assume that $A \supset A_1$. 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 \hat{q} by $\hat{q}(A) := \frac{|A|}{2|A|-1}$. It follows from the above that (A_1, A_1) is an equilibrium if and only if q coincides with or first-order stochastically dominates \hat{q} .

2 \Rightarrow 3. In view of the above equivalence, (A_1, A_1) is an equilibrium by assumption. Suppose to the contrary that (A_h, A_h) is also an equilibrium for some $h > 1$. It holds that $\pi_1(A_h, A_h) = \frac{q(h)}{2} \geq \pi_1(A, A_h)$ for all $A \in \mathcal{M}$. From Lemma (1), it is not profitable for firm 1 to deviate to a menu A with $A \cap A_h \neq \emptyset$. Consider, therefore, a deviation to some menu B_m such that $B_m \cap A_h \neq \emptyset$. It holds that $\pi_1(A_h, A_h) \geq \pi_1(B_m, A_h)$ for all possible m . Suppose $m = 1$. It follows from the above inequality that $\frac{q(h)}{2} \geq \frac{1}{h+1} + 1 - q(h)$ which is true if and only if $q(h) \geq \frac{2h+4}{3h+3}$. However, it holds by assumption that $q(h) \leq \frac{h}{2h-1}$. Moreover, $\frac{2h+4}{3h+3} > \frac{h}{2h-1}$ is true for all $h > 1$. Therefore, we arrive at the contradiction that $\pi_1(B_1, A_h) > \pi_1(A_h, A_h)$.

3 \Rightarrow 4. Suppose (A_h, A_h) is an equilibrium if and only if $h = 1$. It holds that $\pi_1(A_1, A_1) = \frac{1}{2} = \pi_1(A_1, B_1)$ for $A_1 \neq B_1$. Consider a deviation by firm 1 to a profile A where $|A| > 1$. Assume first that $A \cap B_1 = \emptyset$. In view of the 1 \Leftrightarrow 2 equivalence that was established above, and given

the definition of \hat{q} , it follows that $\pi_1(A, B_1) \leq \frac{|A|}{2|A|-1} \cdot \frac{|A|}{|A|+1} = \frac{|A|^2}{2|A|^2+|A|-1} < \frac{1}{2}$. Now suppose $A \cap B_1 \neq \emptyset$. It follows from Lemma 1 that $\pi_1(A, B_1)$ is maximized when $A = B_1$. Therefore, (A_1, B_1) with $A_1 \neq B_1$ is also an equilibrium.

4 \Rightarrow 1. Suppose (A_1, B_1) is an equilibrium for $A_1 \neq B_1$. It holds that $\pi_1(A_1, B_1) = \frac{1}{2} \geq \pi_1(A, B_1)$ for all $A \in \mathcal{M}$. In particular, $\pi_1(B, B_1) \leq \frac{1}{2}$ for all $B \supseteq B_1$ and $\pi_1(A, B_1) \leq \frac{1}{2}$ for all $A \supseteq A_1$. It follows from the first set of equilibrium inequalities that $\frac{1}{2} = \pi_1(B_1, B_1) \geq \pi_1(B, B_1)$ for all $B \supseteq B_1$. It also follows from the second set of equilibrium inequalities that $\frac{1}{2} = \pi_1(B_1, B_1) \geq \pi_1(A, B_1)$ for all $A \in \mathcal{M}$ such that $A \cap B_1 = \emptyset$. Therefore, (B_1, B_1) is an equilibrium. \blacksquare

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 $w_1^A \geq w_2^A \geq \dots \geq w_{|A|}^A$ and those in B such that $w_1^B \geq w_2^B \geq \dots \geq w_{|B|}^B$. We have

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

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

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

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

$$q(A) \cdot \left(\frac{K}{|A| + |B|} + \frac{w_{|A|}^A}{|A| + |B|} \right) \geq q(A) \cdot \left(\frac{K}{|A| + |B| - 1} + \frac{1}{2} \frac{w_1^B}{|A| + |B| - 1} \right) \quad (12)$$

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 (12) and rearranging, we get

$$w_{|A|}^A \left(|A| + |B| - 1 - \frac{\alpha}{2}|A| - \frac{\alpha}{2}|B| \right) \geq K \quad (13)$$

Observe that the assumption on how markups in A are distributed implies that $K \geq (|A| - 1) \cdot w_{|A|}^A$. Therefore, it follows from (13) 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

$$\pi(B, A) = q(A) \cdot \left(\frac{\sum_{i \in I_B} w_i^B}{|A| + |B|} \right) + q(B) - q(A) \quad (14)$$

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

$$\pi(B', A) = q(A) \left(\frac{\sum_{i \in I_B \setminus \{|B|\}} w_i^B}{|A| + |B| - 1} + \frac{1}{2} \frac{w_1^A}{|A| + |B| - 1} \right) + q(B) - q(A) \quad (15)$$

Let $L := \sum_{i \in I_B \setminus \{|B|\}} w_i^B$. Suppose that (A, B) is an equilibrium. It follows from (14) and (15) that

$$\begin{aligned} q(A) \left(\frac{L}{|A| + |B|} + \frac{w_{|B|}^B}{|A| + |B|} \right) + q(B) - q(A) \\ \geq \\ q(A) \left(\frac{L}{|A| + |B| - 1} + \frac{1}{2} \frac{w_1^A}{|A| + |B| - 1} \right) + q(B) - q(A) \end{aligned} \quad (16)$$

By rearranging (16) and solving for w_1^A we get

$$w_1^A \leq \frac{2|A|}{|A| + |B|} w'_{|B|} + G \quad (17)$$

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, (17) reduces to

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

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 (17) reduces to

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

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. ■

Lemma 2.

Suppose all products are equi-profitable. A full product differentiation equilibrium (A_m, B_m) with $m > 1$ exists if and only if the overload cdf q satisfies

$$\frac{q(m)}{q(m-h)} \geq \begin{cases} \frac{6m-4h}{6m-3h}, & h = 1, \dots, m-1 \\ \frac{2m-2h}{2m-h}, & h = -(k-m), \dots, -1 \end{cases} \quad (18)$$

Proof of Lemma 2.

By definition, a profile (A_m, B_m) with $A_m \cap B_m = \emptyset$ is an equilibrium if and only if $\pi_1(A_m, B_m) \geq \pi_1(A, B_m)$ for every menu A . Since $A_m \cap B_m = \emptyset$ and $|A_m| = |B_m| = m$, it holds that

$$\pi_1(A_m, B_m) = \frac{q(m)}{2}.$$

We consider three possible cases for the deviating menu A :

Case 1: $A = A_{m-h} \subset A_m$ for $h = 1, \dots, m-1$. It holds that $\pi_1(A_{m-h}, B_m) = q(m-h) \frac{m-h}{2m-h} + q(m-h) - q(m)$. Therefore, $\pi_1(A_m, B_m) \geq \pi_1(A, B_m)$ if and only if $\frac{q(m)}{2} \geq q(m-h) \frac{m-h}{2m-h} + q(m-h) - q(m)$ which is true if and only if

$$q(m) \geq \frac{6m-4h}{6m-3h} \cdot q(m-h) \quad (19)$$

Case 2: $A \cap B_m \neq \emptyset$. It follows from Lemma 1 that

$$\max_{A \cap B_m \neq \emptyset} \pi_1(A, B_m) = \pi_1(B_m, B_m) = \frac{q(m)}{2} = \pi_1(A_m, B_m).$$

Therefore, deviating to such a menu A is not profitable.

Case 3: $A_{m+h} \cap B_m = \emptyset$ for $h = 1, \dots, k-m$. It holds that $\pi_1(A_{m+h}, B_m) = q(m+h) \frac{m+h}{2m+h}$. Thus, $\pi_1(A_m, B_m) \geq \pi_1(A_{m+h}, B_m)$ if and only if $\frac{q(m)}{2} \geq q(m+h) \frac{m+h}{2m+h}$ which is true if and only if

$$q(m) \geq \frac{2m+2h}{2m+h} q(m+h) \quad (20)$$

Rearranging (19), (20) and rewriting the latter inequality in an equivalent way where h takes negative instead of positive values results in condition (18). \blacksquare

Proof of Proposition 4.

Consider the case of uniformly distributed overload. The cdf q in this case is defined by $q(h) = \frac{k+n+1-h}{k+n}$. Let m be such that $1 < m \leq k+n$ (the case where $k+n < m \leq k$ can be ignored as both firms earn zero payoffs at such (A_m, B_m) , which implies that this profile

cannot be an equilibrium). We have

$$\frac{q(m)}{q(m-h)} = \frac{k+n+1-m}{k+n+1-m+h}, \quad h = -\min\{(k+n-m), (k-m)\}, \dots, m-1. \quad (21)$$

Given Lemma 2 and condition (18), the following are implied by (21):

$$m \leq \frac{4h+k+n+1}{7}, \quad h = 1, \dots, m-1 \quad (22)$$

$$m \geq \frac{k+n+1-2h}{3}, \quad h = 1, \dots, \min\{(k+n-m), (k-m)\} \quad (23)$$

Since the terms on the right hand side of (22) and (23) are increasing and decreasing in h , respectively, they attain their lowest and highest values, respectively, when $h = 1$. In particular, when $h = 1$ they imply that $7m \leq 6+k+n$ and $3m \geq k+n-1$. Subtracting the latter inequality from the former yields $4m \leq 7$, which is impossible given the requirement that $m \geq 2$.

Now suppose that overload is geometrically distributed according to $q(h) = p^{h-1}$, $p \in (0, 1)$. For any $m \in \{2, \dots, \lfloor \frac{k}{2} \rfloor\}$, (18) becomes

$$\frac{p^{m-1}}{p^{m-1-h}} \equiv p^h \geq \begin{cases} \frac{6m-4h}{6m-3h}, & h = 1, \dots, m-1 \\ \frac{2m-2h}{2m-h}, & h = -(k-m), \dots, -1 \end{cases} \quad (24)$$

We will first show that, for any $m \geq 2$, $g(h) := p^h \geq \frac{6m-4h}{6m-3h} := f_m(h)$ for all $h > 1$ if it is true for $h = 1$. Observe that $g'(h) = \ln(p) \cdot p^h < 0$ and $g''(h) = \ln(p)^2 \cdot p^h > 0$ for all $p \in (0, 1)$. Moreover, $f'_m(h) = -\frac{6m}{(6m-3h)^2} < 0$ and $f''_m(h) = -\frac{36m}{(6m-3h)^3} < 0$ for all $h \geq 1$. Therefore, for all $m \geq 2$ it holds that $g''(h) > f''_m(h)$ for all $h \geq 1$ (i.e. f is decreasing at a higher rate than g). Thus, given an arbitrary $m \geq 2$, in order to show that $g(h) \geq f_m(h)$ for all $h = 1, \dots, m-1$ it suffices to show that there exists $p \in (0, 1)$ such that $g(1) \equiv p \geq f_m(1)$.

Similarly, let $s_m(h) := \frac{2m-2h}{2m-h}$. As above, we will show that in order for $g(h) \geq s_m(h)$ to be true for all $h < -1$ given some arbitrary $m \geq 2$ it suffices to show that it is true for $h = -1$ under some $p \in (0, 1)$. Indeed, observe that $g''(h) > 0$ and $s''_m(h) = -\frac{4m}{(2m-h)^3} < 0$ for all $h \leq -1$. Therefore, for any $m \geq 2$ it holds that $g''(h) > s''_m(h)$ for all $h \leq -1$.

It follows from the above that, given some $m \geq 2$, (24) is satisfied only if

$$p \geq \frac{6m-4}{6m-3} \quad (25)$$

$$\frac{1}{p} \geq \frac{2m+2}{2m+1} \quad (26)$$

However, since $\frac{6m-4}{6m-3} > \frac{2m+1}{2m+2}$ for all $m \geq 2$, there is no $p \in (0, 1)$ that makes both (25) and (26) true. Therefore, for each $m \geq 2$ there is no $p \in (0, 1)$ such that (A_m, B_m) is a full product differentiation equilibrium if $q(h) = p^{h-1}$. \blacksquare

Notational remark: Given a strategy profile (A, B) , we let $U_{AB}^A := \sum_{i \in I_{A \setminus B}} p_{A \cup B}(x_i)$ and $S_{AB} := \sum_{i \in I_{A \cap B}} p_{A \cup B}(x_i)$. Notice that, by definition, $U_{AB}^A + S_{AB} + U_{AB}^B = 1$.

Normalization: The following results assume products to be equi-profitable. For this reason, the common markup is normalized to 1 throughout the proofs.

Lemma 3.

A profile (A, A) is an equilibrium if and only if A contains $|A|$ most preferred products in X and, for any $B, C \in \mathcal{M}$ such that $|B| > |A|$ and $|C| < |A|$, the following two conditions hold:

$$\frac{q(B)}{q(A)} \leq \frac{1}{2U_{AB}^B + S_{AB}} \quad (27)$$

$$\frac{q(C)}{q(A)} \leq \frac{3 - 2U_{AB}^B - S_{AB}}{2} \quad (28)$$

Proof of Lemma 3:

It is immediate that if $A \in \mathcal{M}$ does not contain $|A|$ most preferred products in X , then each firm can profitably deviate by replacing a less preferred product with a more preferred one. Let $B, C \in \mathcal{M}$ be such that $|B| > |A|$ and $|C| < |A|$. Notice that

$$\pi_1(A, A) = \frac{q(A)}{2} \quad (29)$$

$$\pi_1(B, A) = q(B) \left(U_{AB}^B + \frac{S_{AB}}{2} \right) \quad (30)$$

$$\pi_1(C, A) = q(A) \left(U_{AC}^C + \frac{S_{AC}}{2} \right) + q(C) - q(A) \quad (31)$$

In order for A to be a symmetric equilibrium strategy it must be that the right hand side of (29) is weakly greater than those in (30) and (31). Solving for $q(B)$ and $q(C)$, respectively, and rearranging establishes (27) and (28). ■

Proof of Proposition 5.

For the existence part of the proof, see Lemma 3. For the uniqueness part, assume by contradiction that (A, A) and (C, C) are both equilibria, and, without loss of generality, $|A| > |C|$. Suppose first that $q(A) = q(C) = \hat{q}$. Therefore, $q(D) = \hat{q}$ for any $D \in \mathcal{M}$ such that $|C| \leq |D| \leq |A|$. Let $\tilde{x} \in X \setminus C$. Such a product \tilde{x} necessarily exists, as $|A| > |C|$ by assumption. Notice that $\pi(C, C) = \frac{\hat{q}}{2}$ and $\pi(C \cup \{\tilde{x}\}, C) = q(C \cup \{\tilde{x}\}) \left(\frac{p_X(\tilde{x})}{S_{CC} + p_X(\tilde{x})} + \frac{1}{2} \frac{S_{CC}}{S_{CC} + p_X(\tilde{x})} \right) = \hat{q} \left(\frac{2p_X(\tilde{x}) + S_{CC}}{2(S_{CC} + p_X(\tilde{x}))} \right)$. Notice that as long as $p_X(\tilde{x}) > 0$, which holds by assumption, the latter is strictly greater than the former. Therefore, each firm offering C has an incentive to increase the size of its menu by including a product in $X \setminus C$, a contradiction. Next, assume that $q(A) < q(C)$. Since (A, A) is an equilibrium, then, by condition (28), $q(C) \leq q(A) \frac{3 - 2U_{AC}^C - S_{AC}}{2}$.

Moreover, since $q(C) > q(A)$, then it must be that $2U_{AC}^C + S_{AC} < 1$. Similarly, since (C, C) is an equilibrium as well, then, by condition (27), $q(C) \geq q(A)(2U_{AC}^A + S_{AC})$. By putting these two together, $3 \geq 4U_{AC}^A + 3S_{AC} + 2U_{AC}^A$. Since, by definition, $U_{AC}^A + S_{AC} + U_{AC}^C = 1$, then $3 \geq 2 + 2U_{AC}^A + S_{AC}$. This inequality is always false, as $2U_{AC}^C + S_{AC} < 1$ implies that $2U_{AC}^A + S_{AC} > 1$, a contradiction. ■

Proof of Proposition 6:

Assume first that $|A^*| = k$, so that $A^* = X$. We now investigate the conditions under which X is a symmetric-equilibrium strategy. Given that X is the largest menu, then profitable deviations can consist of only simpler menus C . By applying condition (28) of Lemma 3, which ensures that deviations to simpler menus are not profitable, we obtain

$$\frac{k + n + 1 - |C|}{k + n + 1 - k} \leq \frac{3}{2} - \frac{S_{XC}}{2}$$

The right-hand side of this inequality is always greater than 1, because $\frac{S_{XC}}{2}$ is strictly smaller than $\frac{1}{2}$. On the other hand, as $n \rightarrow \infty$, then the left-hand side converges towards 1. This implies that as the menu size threshold of the least overloaded consumer grows to infinity, then X is a symmetric equilibrium strategy.

However, if n is sufficiently small, but large enough to ensure that X is the welfare-maximising menu, then X is not necessarily a symmetric equilibrium strategy as the example in the main body of the paper illustrates.

Assume now that $|A^*| = \left\lceil \frac{k+n+1}{2} \right\rceil$. Firstly, we show that deviating to more complex menus B is not profitable. By applying condition 27 of Lemma (3), which ensures that deviations to more complex menus are not profitable, we obtain

$$(k + n + 1 - |B|) \left(\left\lceil \frac{k + n + 1}{2} \right\rceil \right)^{-1} \leq \frac{1}{2U_{AB}^B + S_{AB}} \quad (32)$$

Notice that in the most unfavourable case preferences are uniformly distributed. The reason is that, by proposition 5, A^* must contain the $|A^*|$ most preferred products. Therefore, any product in A^* must have a choice probability at least as high as any product outside $X \setminus A^*$ implying that in the most unfavourable case the products outside A have the same choice probabilities of the products in A . Hence, we assume that preferences are uniform. By Lemma (3), expected profits to deviations to more complex menus B are maximised when the size of the intersection $|A^* \cap B|$ is maximised. Hence, the right-hand side of the above inequality reduces to $\frac{1}{2 \frac{|B|-|A^*|}{|B|} + \frac{|A^*|}{|B|}} = \frac{|B|}{2|B| - \lceil \frac{k+n+1}{2} \rceil}$. In the supplementary material (available from the authors upon request) we show that the resulting inequality is always true.

Secondly, we prove that deviating to simpler menus C is always profitable. By applying condition 28 of Lemma (3), we obtain

$$(k+n+1-|C|) \left(\left[\frac{k+n+1}{2} \right] \right)^{-1} \leq \frac{3 - S_{AC} - 2U_{AC}^C}{2} \quad (33)$$

If $\left[\frac{k+n+1}{2} \right] = 1$, then the statement is vacuously true. Since deviating to more complex menus is never profitable, then, whenever $\left[\frac{k+n+1}{2} \right] = 1$, offering a menu of size 1 is an equilibrium strategy. Next, assume that $\left[\frac{k+n+1}{2} \right] \geq 2$. Let $|C| = 1$. Then, the left-hand side reduces to $1 \times \left(\frac{k+n+1 - \left(\left[\frac{k+n+1}{2} \right] \right)}{k+n} \right)^{-1}$. In the most favourable case the right-hand side is equal to $\frac{3}{2} - \epsilon$ for some arbitrarily small $\epsilon > 0$. In the supplementary material we show that the resulting inequality is false. ■

Proof of Proposition 7.

Assume first that $|A^*| = 1$. Then, profitable deviations can be only be to more complex menus B . By applying condition (27) of Lemma (3),

$$p^{|B|-1} \leq \frac{1}{2U_{AB}^B + S_{AB}} \quad (34)$$

In the most unfavourable case, the left-hand side is equal to $p = e^{-2/3} \simeq 0.513417$ (i.e., $e^{-2/3}$ solves $1.5 = -\frac{1}{\ln(p)}$ and is the highest value of p that makes the welfare-maximising menu size equal to 1 by taking the rounding rule into account) and, by the arguments in the proof of proposition (6), preferences are uniformly distributed. Therefore, the right-hand side reduces to $\frac{|B|}{2|B|-1}$. We show in the supplementary material that the resulting inequality is always true. Hence, whenever the welfare-maximizing menu size is 1 offering a singleton menu is an equilibrium strategy.

Next, assume that $|A^*| = \left\lceil -\frac{1}{\ln(p)} \right\rceil$. By taking the inverse, the welfare-maximising probability parameter p is equal to $e^{-\frac{1}{|A^*|}}$. We first show that deviating to simpler menus C constitutes a profitable deviations. By applying condition 28 of Lemma (3)

$$\left(e^{-\frac{1}{|A^*|}} \right)^{|C|-|A^*|} \leq \frac{3 - 2U_{AC}^C - S_{AC}}{2} \quad (35)$$

In the most favourable case the right-hand side of equation (35) is equal to $3/2$. We show in the supplementary material that if $|C| = 1$, then the above inequality is false for any value of $|A^*|$. Hence, deviating to a sub-menu is always profitable.

Next, consider deviations to more complex menus B . We now prove that deviating to B is never profitable. As discussed previously, in the most unfavourable case preferences are uniformly distributed. Then, by applying condition (27) of Lemma (3),

$$\left(e^{-\frac{1}{|A^*|}}\right)^{|B|-|A^*|} \leq \frac{|B|}{2|B|-|A^*|} \quad (36)$$

We show in the supplementary material (available upon request) that the above inequality is always true.

Finally, consider the case where $|A^*| = k$. In the supplementary material we show that when the parameter p is sufficiently small, then offering a maximum-variety menu is an equilibrium strategy. Moreover, in the main body of the paper we provide an example showing that when p is large, whether offering a menu with k alternatives is an equilibrium strategy depends upon consumers' preferences. ■

Finally, as referred to in footnote 15, we conclude with an example showing that if overload is geometrically distributed and p is large enough, then whether offering a maximum-variety menu is an equilibrium or not depends on consumers' preferences. Indeed, assume that $X = \{x_1, x_2, x_3\}$. Assume also that $p = 0.98$. Notice that $0.98 > \frac{1}{e^{1/3}} \simeq 0.72$. Therefore, a menu of size 3 maximises welfare. Assume next that choice probabilities are $p_X(x_1) = p_X(x_2) = p_X(x_3) = 1/3$. Then, condition 28 of Lemma 3 reduces to $1.04 \simeq 0.98^{1-3} \leq \frac{3-\frac{1}{3}}{2} \simeq 1.33$ when $|C| = 1$ and $1.02 \simeq 0.98^{2-3} \leq \frac{3-\frac{2}{3}}{2} \simeq 1.17$ when $|C| = 2$. In both cases the inequality is true. Therefore, X is a symmetric equilibrium strategy. Now assume, instead, that choice probabilities are $p_X(x_1) = 0.99$, $p_X(x_2) = p_X(x_3) = 0.005$. Then, condition (28) of Lemma 3 reduces to $1.02 \simeq 0.98^{2-3} \leq \frac{3-(0.99+0.005)}{2} \simeq 1.0025$ when $|C| = 2$, which leads to a contradiction. Hence, in this second case X is not a symmetric equilibrium strategy.

C Current Market Examples of Firms Competing in “Small” Menus

C.1 Subscription TV Contracts

Select your TV Bundle

Help me choose



Original & Sky Q

Great entertainment with award-winning shows.

£22 a month

+ £25 choice of reward
+ Sky Q, our next generation box

Select & personalise

TV set-up fees apply.



Variety & Sky Q

We've got it covered from drama queens, fact geeks, music fans and kids

£32 a month

+ £25 choice of reward
+ Sky Q, our next generation box

Select & personalise

TV set-up fees apply.



Box Sets & Sky Q

Discover over 350 Sky Box Sets plus over 50 channels in HD.

£38 a month

+ £25 choice of reward
+ Sky Q, our next generation box

Select & personalise

TV set-up fees apply.

Source: www.sky.com/shop/sky-bundles (accessed on 13th June 2017)

Mix TV	Fun TV	Full House TV
£20 a month for 12 months	£31.50 a month for 12 months	£41.50 a month for 12 months
£20.00 activation fee 12 month contract No Virgin Phone line	£20.00 activation fee 12 month contract No Virgin Phone line	12 month contract £20.00 activation fee No Virgin Phone line
Start your order	Start your order	Start your order

Source: www.virginmedia.com/shop/tv/tv-only.html (accessed on 28th June 2017; snapshot compilation)

C.2 Broadband Contracts

Unlimited Broadband	Unlimited Infinity 1	Unlimited Infinity 2
<p>£24.49/mth</p> <p>40% WELCOME DISCOUNT INCLUDED</p> <p>Includes BT Line rental, 12-month contract Upfront £9.99 Price after 12 months: £40.99</p> <p>OFFER ENDS 13d 09h 47m 25s</p> <p>Up to 17Mb speeds Unlimited monthly usage</p> <p>Choose this package ></p>	<p>£29.99/mth</p> <p>40% WELCOME DISCOUNT INCLUDED</p> <p>Includes BT Line rental, 12-month contract Upfront £59.99 Price after 12 months: £47.49</p> <p>OFFER ENDS 13d 09h 47m 25s</p> <p>Up to 52Mb speeds Unlimited monthly usage</p> <p>Choose this package ></p>	<p>£49.99/mth</p> <p>Includes BT Line rental, 12-month contract Upfront £59.99 Price after 12 months: £53.99</p> <p>Up to 76Mb speeds Unlimited monthly usage</p> <p>Choose this package ></p>




Source: www.productsandservices.bt.com/products/broadband-packages (accessed on 30th June 2017)

SKY BROADBAND UNLIMITED	SKY FIBRE	SKY FIBRE MAX	SKY FIBRE MAX
<p>£28.99 £18.99 a month</p> <p>for 12 months (^ then £28.99 a month) + £75 Prepaid MasterCard® Includes line rental & Sky Talk + one off £19.95 set-up fee</p> <p>Add to basket</p> <p>Find out more</p>	<p>£28.99 a month</p> <p>+ £75 Prepaid MasterCard® Includes line rental & Sky Talk + £59.95 one off set-up fee, 18 month contract</p> <p>Add to basket</p> <p>Find out more</p>	<p>£38.99 a month</p> <p>+ £75 Prepaid MasterCard® Includes line rental & Sky Talk + £59.95 one off set-up fee, 18 month contract</p> <p>Add to basket</p> <p>Find out more</p>	<p>£43.99 a month</p> <p>+ £75 Prepaid MasterCard® Includes line rental & Sky Talk + £59.95 one off set-up fee + free Sky Q Hub, 18 month contract</p> <p>Add to basket</p> <p>Find out more</p>

Source: www.sky.com/shop/broadband-talk (accessed on 30th June 2017)

C.3 Current Accounts

Featured current accounts
View our range of current accounts below or [check your eligibility](#) for a current account and overdraft.

 <p>Classic Account The everyday account for everyday banking</p> <p>Find out more ></p>	 <p>Club Lloyds An account with extras. Pay at least £1,500 into your account every month, otherwise it costs £3 per month</p> <p>Find out more ></p>	 <p>Platinum Packed with insurance benefits for only £17 per month</p> <p>Find out more ></p>
---	--	--







Source: <https://www.lloydsbank.com/current-accounts.asp> (accessed on 16th June 2017)

Compare our current accounts

<p>SELECT</p> <p>No fee Our standard bank account</p> <p>Apply online</p>	<p>REWARD</p> <p>£2 per month A bank account that Rewards you</p> <p>Find out more</p>	<p>REWARD SILVER</p> <p>£12 per month Rewards and lifestyle benefits</p> <p>Find out more</p>	<p>REWARD PLATINUM</p> <p>£19 per month Rewards and even more benefits</p> <p>Find out more</p>	<p>REWARD BLACK</p> <p>£28 per month An account for our Private Customers</p> <p>Find out more</p>
---	--	---	---	--




Source: <http://personal.rbs.co.uk/personal/current-accounts/compare-current-accounts.html> (accessed on 30th June 2017)

C.4 Laptop Computers

  Space Grey	  Space Grey	  Silver
<p>New</p> <p>Touch Bar and Touch ID 2.8GHz Processor 256GB Storage</p> <p>2.8GHz quad-core 7th-generation Intel Core i7 processor</p> <p>Turbo Boost up to 3.8GHz</p> <p>16GB 2133MHz LPDDR3 memory</p> <p>256GB SSD storage¹</p> <p>Radeon Pro 555 with 2GB memory</p> <p>Four Thunderbolt 3 ports</p> <p>Touch Bar and Touch ID</p> <p>£2,349.00</p>	<p>New</p> <p>Touch Bar and Touch ID 2.9GHz Processor 512GB Storage</p> <p>2.9GHz quad-core 7th-generation Intel Core i7 processor</p> <p>Turbo Boost up to 3.9GHz</p> <p>16GB 2133MHz LPDDR3 memory</p> <p>512GB SSD storage¹</p> <p>Radeon Pro 560 with 4GB memory</p> <p>Four Thunderbolt 3 ports</p> <p>Touch Bar and Touch ID</p> <p>£2,699.00</p>	<p>2.2GHz Processor 256GB Storage</p> <p>2.2GHz quad-core Intel Core i7 processor</p> <p>Turbo Boost up to 3.4GHz</p> <p>16GB 1600MHz memory</p> <p>256GB SSD storage¹</p> <p>Intel Iris Pro Graphics</p> <p>Two Thunderbolt 2 ports</p> <p>£1,899.00</p>

Source: <https://www.apple.com/uk/shop/buy-mac/macbook-pro/15-inch> (accessed on 28th June 2017; snapshot compilation)

Viewing 1-3 of 3

<input type="checkbox"/> Add to Compare	<input type="checkbox"/> Add to Compare	<input type="checkbox"/> Add to Compare
 <div style="background-color: #0070C0; color: white; padding: 2px; font-size: 8px;">24h Onsite Support included</div>	 <div style="background-color: #0070C0; color: white; padding: 2px; font-size: 8px;">24h Onsite Support included</div>	 <div style="background-color: #0070C0; color: white; padding: 2px; font-size: 8px;">24h Onsite Support included</div>
<p>XPS 15 £1,349.00</p> <hr/> <p>7th Generation Intel® Core™ i5 processor</p> <hr/> <p>Windows 10 Home</p> <hr/> <p>8GB Memory</p> <hr/> <p>1TB 5400RPM Hard Drive + 32GB Solid State Drive</p>	<p>XPS 15 £1,429.00</p> <hr/> <p>7th Generation Intel® Core™ i7 processor</p> <hr/> <p>Windows 10 Home</p> <hr/> <p>8GB Memory</p> <hr/> <p>256GB Solid State Drive</p>	<p>XPS 15 £1,599.00</p> <hr/> <p>7th Generation Intel® Core™ i7 processor</p> <hr/> <p>Windows 10 Home</p> <hr/> <p>16GB Memory</p> <hr/> <p>512GB Solid State Drive</p>

▶ Longest battery life. Up to 19.5 hours!

Source: <http://www.dell.com/uk/p/xps-15-9560-laptop/pd> (accessed on 30th June 2017)