Buridanic Competition*

Benjamin Bachi[†] and Ran Spiegler[‡]

January 14, 2014

Abstract

We analyze a model in which two profit-maximizing firms compete in two-attribute products over agents who follow a non-compensatory choice procedure that responds purely to ordinal quality rankings: sticking to a default option when no market alternative dominates another, and focusing on a random attribute when choosing by default is impossible. Our equilibrium analysis highlights the effect of such trade-off avoidance on various aspects of the market outcome: total quality of the offered products, amount of obfuscation, prevalence of "hard choices", as well as market participation and consumer switching rates. We discuss the potential implications of this analysis for "default architecture".

*Spiegler acknowledges financial support from ERC Grant no. 230251. Bachi benefitted from ERC grant no. 269143. We are grateful to Kfir Eliaz and Ariel Rubinstein for useful comments.

[†]Tel Aviv University. E-mail: beniamin@post.tau.ac.il.

[‡]Tel Aviv University, University College London and CFM. URL: http://www.tau.ac.il/~rani. E-mail: r.spiegler@ucl.ac.uk.

1 Introduction

One of the biggest distinctions between economists' and psychologists' view of the decision process is the way they regard trade-offs. The standard economic approach assumes that the decision maker has well-defined preferences, and in the vast majority of applications these preferences are continuous and locally non-satiable, which implies that decision makers are "trade-off machines" who effortlessly weigh multiple considerations.

A viewpoint more typical of psychologists (e.g., Tversky (1972), Payne et al. (1993), Luce et al. (1999), Anderson (2003)) is that decision makers generally try to *avoid trade-offs*, for a variety of reasons. First, trade-offs may be hard to calculate. Second, in many cases, making a decision that lacks a solid justification and relies on intangible, subjective "decision weights", is emotionally difficult. Finally, when the decision maker is required to *justify* his choice to other people - a key motive in organizational behavior - he is expected to provide reasons for his choice, and decision weights are hard to justify in this manner.

Such "trade-off avoidance" implies that decision makers will tend to use so-called "non-compensatory choice procedures" that rely purely on the *ordinal* rankings over alternatives along each dimension. In particular, when they have a default option that enables them to "decide not to decide", they may exercise this option. What are the implications of this view of the decision process for market behavior?

This paper studies a model in which two profit-maximizing firms compete in two-attribute products for a continuum of homogenous agents. Firm i's product is characterized by a pair $(q_i^1, q_i^2) \in \mathbb{R}_+^2$, where q_i^k represents the quality of the firm product's along dimension k. Firm i's profit from selling one unit is $1 - c(q_i^1, q_i^2)$, where c is a continuous and strictly increasing cost function. The agents' choice set consists of the firms' products as well as an outside option represented by the quality pair $(0, 0)$. A conventionally rational agent would be endowed with a continuous, strictly increasing function $u(q^1, q^2)$, and would choose an alternative that maximizes u. In this case, our model would collapse into standard Bertrand competition: in Nash equilibrium, firms would offer quality pairs (q^1, q^2) that maximize u subject to $c(q^1, q^2) = 0$.

In contrast, we assume that the agents follow a non-compensatory choice procedure, which is based solely on *ordinal quality rankings.* When one market alternative dominates another (i.e., $q_i^k \geq q_j^k$ for some $i \neq j$ and for both $k = 1, 2$, with at least one strict inequality), the agent chooses the dominant product; we refer to this situation as an "easy choice", because it does not require the agent to perform any trade-offs. When neither market alternative dominates the other (a situation we refer to as a "hard choice"), the agent responds by "deciding not to decide", namely sticking to a *default* option. The default specification is a feature of the market's design. We examine three default policies: (*i*) "*opt in*": the default option is the outside option; (*ii*) "*opt out*": the default option is one of the firms (each with equal probability, for simplicity); (*iii*) "*no default*": the agent is forced to make an active choice and cannot choose by default. In this case, we assume that the agent resolves his tension by focusing entirely on a single, randomly selected attribute (again, avoiding trade-offs), where the probability with which each attribute is selected represents its relative *salience.* We also consider mixed specifications, in which different agents obey different default rules.

Note that by assumption, the outside option is inferior to any of the market alternatives. Thus, under "opt in", the agent is much like the proverbial Buridan's Ass: unable to resolve the trade-off between two superior market alternatives, he procrastinates choice and ends up settling for the inferior outside option. "Buridanic" situations of a similar nature have received considerable attention by researchers who studied empirically the intertwined phenomena of choice complexity and choice procrastination, in both experimental and "field" settings, notably retirement savings - see Iyengar and Lepper (2000), Iyengar et al. (2004), Madrian and Shea (2001) and Beshears et al. (2012), for a few important examples. Our contribution is to formulate a simple procedural model that exhibits trade-off avoidance and induces choice procrastination, in the context of an otherwise competitive market, and allows us to discuss the theoretical equilibrium implications of various default rules.

Our model fits market environments in which consumers find it hard to trade-off various product attributes. For example, think about weighing a car's safety against its fuel efficiency. Even if the consumer has access to hard data about each attribute, it may be hard for him to find the right scale for comparison. The difficulty is not only cognitive; as one of the dimensions is safety, the consumer ultimately has to trade off the risk of injury or death with lower fuel costs. Similarly, when workers choose between retirement savings plans that involve life insurance and disability benefits, they not only have to perform complex actuarial calculations, but also face the unnerving task of thinking about all sorts of terrible things that might happen to them in old age. Indeed, much of the above-cited research focused on the role of default architecture in addressing choice procrastination in this context. In other market settings, as in the case of buying a car, the natural default rule is either "opt in" or "no default" (the latter fits a situation in which the agent must buy a new car, such that delaying the purchase out of indecision is simply not an option). In the case of magazine subscription, "opt out" corresponds to an auto-renewal policy, while "opt in" corresponds to a market regulation that forbids auto-renewals.

Another interpretation of the model fits non-market, organizational settings. For instance, think of a company official who considers several candidates for a construction job, which involves several dimensions (total cost, speed of delivery, quality of materials, etc.). The decision maker is an official in an organization, who may have to justify his selection either ex-ante or ex-post. When the selection is not obvious and requires judgment, the official is more vulnerable to criticisms by his supervisors; it may be easier for him to justify his selection if he does not acknowledge some relevant dimensions, thus presenting the problem to his superiors as being simpler than it actually is, or if he goes for a default provider (in case one exists).

We analyze symmetric (mixed-strategy) Nash equilibria in the simultaneous-move game played between the two firms, under a variety of default specifications. We first make the simple observation that as far as the payoff structure of the firms' game is concerned, the "opt out" policy is equivalent to "no default" with symmetric attribute salience. Thus, default policies are logically linked to disclosure measures that affect the attributes' relative salience. Next, we focus on the "no default" case with arbitrary attribute salience. We show that if the cost function c is weakly super-modular, no market alternative ever dominates another in equilibrium: the realization of q^2 is a deterministic decreasing function of the the realization of q^1. Thus, the equilibrium market response to "no default" removes easy choices from the agents' landscape.

For our subsequent results, we strengthen the structure of c and assume it is additively separable, such that quality along any dimension can be measured by its cost to the firm, w.l.o.g. First, we provide a complete characterization of symmetric equilibrium under a pure "no default" policy, for arbitrary attribute salience. As the two attributes become more equally salient, the expected total cost of the products offered in equilibrium goes down, as does the cost variation across the products' two attributes. This means that there is a sense in which greater product quality is positively correlated with a greater amount of obfuscation (defined in terms of the gap between "true" and "perceived" quality, measured in cost units). We then proceed to incorporate opting in. The case of a pure "opt in" policy is simple: the firms' equilibrium behavior is "competitive", in the sense that firms mix somehow over products with $c(q^1, q^2) = 1$. Yet, since neither market alternative dominates another in equilibrium, agents adhere to their default and thus there is no market participation.

Finally, we consider a mixed default policy that assigns "opt in" and "opt out" to different groups of agents. This case is considerably more complex to analyze, and gives rise to multiple equilibria. We show that domination occurs in any equilibrium, such that agents face easy choices with positive probability, which induces positive switching rates. We isolate a class of symmetric equilibria, in which firms mix over total cost according to some continuous density, and independently obfuscate (i.e., shift the cost across attributes) according to a discrete uniform distribution, such that agents face an "easy choice" if and only if the realization of the obfuscation strategy is the same for both firms. We provide an upper bound on market participation and switching rates for this class of equilibria, and show that this class is fully characterized by two properties of interest. In the concluding section, we discuss the possible implications of these results for contemporary discussions of "default architecture".

Related literature
This paper belongs to the growing literature on "behavioral industrial organization" (see Spiegler (2011) for a textbook treatment). Within this literature, two papers are most closely related. Gabaix and Laibson (2006) analyze a model in which two firms compete in price pairs, where a fraction of consumers choose are unaware of dimension 2, and thus choose purely on the basis of price rankings along dimension 1, while the remaining consumers are conventionally rational and choose the firm with the lowest true price. Consumers have an outside option, the value of which is correlated with their type (the interpretation is that more sophisticated consumers are more likely to find a good outside option). Indeed, the "no default" version of our model can be interpreted in terms of unawareness: the agent focuses on one dimension because he is unaware of the other.

Spiegler (2006) analyzes a model in which n firms choose price *cdf*s over $(-\infty, 1]$. A firm's profit conditional on being chosen is the expected price according to its own *cdf*. The consumer chooses by taking a sample point from each of the *cdf*s and selecting the cheapest firm in his sample. As Spiegler (2006) notes, this can be viewed as a reduced form of a model in which firms choose infinite-dimensional price vectors and the consumer chooses according to the price ranking in a randomly selected dimension. This interpretation forms a clear link with the present model, and suggests an interesting generalization of our model to the case of n firms and K dimensions, in which consumers choose according to some probabilistic aggregation of the ordinal rankings along each dimension. From this perspective, Spiegler (2006) assumes a specific aggregation rule - random dictatorship - and takes the limit $K \to \infty$.

The present paper proposes an approach to modeling market competition when consumers have limited ability to make comparisons. Piccione and Spiegler (2012) suggest a different approach. A market alternative consists of a "real price" and a "price format", and consumers are able to make a price comparison (and thus choose the cheapest firm) if and only if the two firms' price formats are comparable, according to some primitive comparability structure. Carlin (2009) and Chioveanu and Zhou (2013) study special cases of this limited comparability formalism and extend them to the many-firms case. All these papers can be viewed as extensions of Varian (1980), who studied price competition when an exogenous fraction of the consumer population does not make comparisons. The new models essentially endogenize this parameter as a consequence of the firms equilibrium obfuscation tactics.

A few choice-theoretic works studied boundedly-rational, non-compensatory choice procedures. Rubinstein (1988) analyzes a procedure related to ours, where the decision maker regards one two-attribute alternative as dominating another if it is "approximately the same" along one dimension and significantly better along the other. Mariotti and Manzini (2007) axiomatize a "sequentially rationalizable" choice procedure that employs a succession of binary relations to eliminate alternatives from the choice set. Papi (2013) axiomatizes a variant that mixes compensatory and non-compensatory elements, and applies it to a Stackelberg model. In this model, the decision maker uses such a procedure only to shrink the size of his choice set, and then maximizes a well-defined utility function to the constructed consideration set. And of course, an important strand in decision theory, running through Bewley (1986) and Ok (2002), has axiomatized multi-utility representations of incomplete preferences, which in a model with a default option imply a default bias (see Masatlioglu and Ok (2005)). Dean (2008) conducts an experimental test of axioms that characterize various families of models of decision avoidance.

Finally, the role of attribute salience in consumer choice has been recently studied by Koszegi and Szeidl (2013) and Bordalo et al. (2013a). These papers model salience as a systematic distortion of decision weights, while the present paper captures the salience of an attribute by the probability it is considered by a trade-off avoiding decision maker. Bordalo et al. (2013b) and Spiegler (2013) analyze market models in which decision weights are endogenously determined by firms' pricing and marketing equilibrium strategies.

2 The Model

Two firms play the following symmetric simultaneous-move game. Each firm $i = 1, 2$ offers a product characterized by a pair $(q_i^1, q_i^2) \in \mathbb{R}_+^2$, where q_i^k represents the quality of product attribute k. The firms face a measure one of agents, whose choice set consists of the firms' two products, as well as an outside option represented by the quality pair $(0, 0)$. When an agent chooses firm i, the firm incurs a cost of $c(q_i^1, q_i^2)$. Firm i's payoff conditional on being selected is $1 - c(q_i^1, q_i^2)$.

Agents choose according to the following procedure. When $q_i^k \geq q_j^k$ for both $k = 1, 2$, with at least one strict inequality - i.e. when one market alternative dominates another - every agent chooses the firm offering the dominant product. We refer to such a situation as an "easy choice". When $q_i^1 > q_j^1$ and $q_i^2 < q_j^2$ (a case we refer to as a "difficult choice"), we distinguish between two cases: (i) when one of the three feasible alternatives is a default option, the agent sticks to this default; (ii) when the agent has no default option - i.e., he has to make an active choice - the agent chooses firm i with probability α^1 and firm j with probability α^2. We will often denote $\alpha^1 = \alpha$. The parameter α captures the relative salience of the first quality dimension. We will consider various default rules, including heterogeneity among agents in this regard, in the sequel.

To illustrate the firms' payoff function, consider a strategy profile in which $q_1^1 > q_2^1$ and $q_1^2 < q_2^2$. If all agents are initially assigned to the outside option as a default, both firms earn zero profits. Now consider an alternative default rule, according to which each firm serves as the default option for 50% of the agent population. Then, each firm i earns $\frac{1}{2}(1 - c(q_i^1, q_i^2))$. Finally, if agents have no default option, firm 1 earns $\alpha(1 - c(q_1^1, q_1^2))$ while firm 2 earns $(1 - \alpha)(1 - c(q_2^1, q_2^2))$. By comparison, if the strategy profile satisfied $q_1^k > q_2^k$ for both $k = 1, 2$, firm 1's payoff would be $1 - c(q_1^1, q_1^2)$, while firm 2 would earn zero profits, independently of the default rule.

The agent's choice behavior departs from rationality in several dimensions. First, it is sensitive to the default specification; this is a framing effect that conventionally rational decision makers do not exhibit. Second, even when we hold the default rule fixed, we can observe violations of rationality, even without imposing the assumption that agents' utility function is increasing. For instance, consider the "opt in" rule. When the available market alternatives are $(2, 2)$ and $(1, 1)$, the agents choose the former. However, if we replace the latter alternative with $(1, 3)$, the agents shift to the outside option $(0, 0)$.

An interpretational difficulty arises when both firms offer the same quality pair, i.e. $q_1 = q_2$. In this case, our procedure above assumes the agent chooses by default whenever this is possible. However, in our opinion, the case for such "Buridanic" behavior when the agent faces two literally identical market alternative seems much weaker than in the case in which he faces distinct, difficult-to-compare alternatives. For instance, a plausible alternative assumption is that the agent chooses each market option with probability $\frac{1}{2}$. Fortunately, except for one special case, this criticism will not affect our equilibrium analysis: $q^1 = q^2$ will typically be a zero-probability event in symmetric Nash equilibrium (under any assumption about how the agent decides in this event).

The cost function

Throughout the paper, we assume that c is continuous and strictly increasing, with $c(0,0) = 0$. Also, there exist finite q^1 and q^2 such that $c(0,q^2) = c(q^1,0) = 1$. When imposing additional structure on c, one should bear in mind that since the agents' choice procedure is based entirely on ordinal rankings, the cardinal meaning of the quality variables q^1 and q^2 is questionable, and therefore we should be cautious when making assumptions that are based on cardinal quality measurements. Thus, for some of the results, we will require c to be *weakly supermodular* - i.e., for every two quality pairs q and r, $c(q \vee r) + c(q \wedge r) \geq c(q) + c(r)$, where $q \vee r = (\max\{q^1, r^1\}, \max\{q^2, r^2\})$, $q \wedge r = (\min\{q^1, r^1\}, \min\{q^2, r^2\})$. For other results, we will impose the stronger assumption of *additive separability*: $c(q^1, q^2) = c^1(q^1) + c^2(q^2)$. Since the consumer's choice procedure is invariant to monotone transformations of c^k, additive separability means that we can assume w.l.o.g that $c^k(q^k) = \frac{1}{2}q^k$, i.e. quality along each dimension is measured by (twice) its cost to the firm, such that the cost of a firm's product is equal to its average quality.

Comment: Quality vs. prices

Although we refer to the two product attributes in terms of quality, it is of course possible to interpret them as prices. For instance, dimension 1 can be the product's price, such that $q^1 = -p$, while dimension 2 is a "proper" quality dimension. One difficulty with this interpretation is that the lower bound on q^1 now implies an upper bound on the product's price, and such a bound is harder to justify. It may represent an ex-post participation constraint: the agent's choice of a firm between firms means signing a contract that gives him the right to buy the firm's product, and the agent can later choose not to exercise this right, if he finds the price too high (at that stage, he does not have to make a comparison with other alternatives, and it is easier for him to figure out his willingness to pay for the product). It is also possible to interpret both dimensions as prices. For instance, the firms' product can be viewed as a contract that specifies state-contingent payouts. If firms maximize expected profits, their payoff will be linear in q. One source for the agent's trade-off avoidance is lack of a firm prior belief over states. This is in line with Bewley's (1986) notion of Knightian uncertainty as incomplete preferences.

Comment: Heterogeneous agents
Our model assumes that agents make active choices (by considering a single dimension at random) only when the choice is easy, or when they lack a default option; otherwise, they stick to the default in response to difficult choices. However, a formally equivalent assumption would be that there is an additional, "decisive" agent type, who chooses "actively" according to a randomly selected attribute even when he could choose by default. More precisely, for any mixed default rule in our model, we can find a distribution between these two agent types and some other mixed default rule that would give rise to the same payoff function for the firms. Thus, although our model seemingly makes the assumption that *all* agents exhibit an extreme default bias, there is an alternative interpretation that allows for heterogeneity among agents in terms of their response to difficult choices.

There is an element of heterogeneity that is entirely absent from our model, namely the coexistence of trade-off avoiding and conventionally rational agents (or, more generally, agents that follow choice procedures that are sensitive to *cardinal* quality measurements). Extending the model in this direction turns out to be non-trivial, and therefore it is left for future work. The present model should be viewed as an extreme case, which is diametrically opposed to the conventional model that regards consumers as impeccable "trade-off machines".

3 Symmetric Equilibrium Analysis

We now turn to analysis of symmetric Nash equilibria in the two-firm game described in Section 2, under various default rules.

3.1 No Default

We begin with the case in which all agents are under the "no default" regime, such that they choose according to dimension 1 (2) with probability α $(1-\alpha)$. Without loss of generality, let $\alpha \geq \frac{1}{2}$. The case of $\alpha = 1$ is simple (it is formally a special case of Gabaix and Laibson (2006)). In Nash equilibrium, firms will offer $q^2 = 0$ and q^1 will be determined by the equation $c(q^1, 0) = 1$. The reasoning is simple: since the agent never considers dimension 2, firms have no incentive complete on this dimension. In contrast, competitive pressures along dimension 1 drive its quality up in Bertrand fashion, such that in equilibrium firms must make zero profits.

The case of $\alpha \in [\frac{1}{2}, 1)$ is more interesting.

Proposition 1 (No easy choices) *Consider the "no default rule" and let $\alpha \in [\frac{1}{2}, 1)$. If c is weakly supermodular, then for any symmetric Nash equilibrium, there exist $\bar{q}^1, \bar{q}^2 > 0$ such that the support of the equilibrium strategy is a continuous and strictly decreasing curve that connects the points $(0, \bar{q}^2)$ and $(\bar{q}^1, 0)$.*

This result means that when agents are forced to make active choices and cannot choose by default, symmetric Nash equilibrium has the feature that no market alternative ever dominates another. That is, the agents always face difficult choices. What is the significance of this result? Note that one interpretation of our model is that making difficult choices involves a mental cost, which agents successfully avoid if they can choose by default. When they are forced to make an active choice, the mental cost is incurred whenever they face a difficult choice. Our result means that in that case, spontaneous competitive forces "conspire" to maximize this mental cost, as long as c is weakly supermodular. The key argument in the proof is that if there were two quality pairs q and r in the support of the equilibrium strategy that dominate one another, then deviating to either (q_1, r_2) or (r_1, q_2) would be profitable.

For our next results, we strengthen the structure of c and assume it is additively separable. Recall that w.l.o.g, $c(q^1, q^2) = \frac{1}{2}(q^1 + q^2)$.

Proposition 2

(i) Let $\alpha = \frac{1}{2}$. Under the "no default" rule, the game has a unique symmetric Nash equilibrium, in which firms play $q^1 \sim U[0, 1]$, and $q^2 = 1 - q^1$ with probability one, such that total cost is $\frac{1}{2}$ with probability one.

(ii) Let $\alpha \in (\frac{1}{2}, 1)$. Under the "no default" rule, the game has a unique symmetric Nash equilibrium, in which firms mix over total cost $\frac{1}{2}(q^1 + q^2) = c$ according to the *cdf*

$$G(c) = \frac{1-\alpha}{2\alpha - 1}\left[\frac{\alpha}{1-c} - 1\right]$$

defined over the interval $[1-\alpha, \alpha]$. The quality along each dimension is a deterministic function of c:

$$\begin{aligned} q^1 &= \frac{2\alpha}{2\alpha - 1}\left[c - (1-\alpha)\right] \\ q^2 &= \frac{2(1-\alpha)}{2\alpha - 1}\left[\alpha - c\right] \end{aligned}$$

When $\alpha \in (\frac{1}{2}, 1)$, firms mix over average quality c in equilibrium. The greater the asymmetry in the attributes' salience, the greater the range of values that c gets in equilibrium. Note that the $\alpha \to \frac{1}{2}$ and $\alpha \to 1$ limits of this equilibrium characterization coincide with our analysis for these extreme cases.

The expectation of c is

$$E_G(c) = 1 - \frac{\alpha(1-\alpha)}{2\alpha - 1} \ln\left(\frac{\alpha}{1-\alpha}\right)$$

which is strictly increasing in α in the range $(\frac{1}{2}, 1)$. Note that in equilibrium q^1 takes values in $[0, 2\alpha]$ while q^2 takes values in $[0, 2(1-\alpha)]$, and the two quality components are linked to each other deterministically by the linear equation

$$q^2 = 2(1-\alpha) - \frac{1-\alpha}{\alpha} q^1$$

Let us calculate equilibrium industry profits. Consider the quality pair $(q^1, q^2) = (2\alpha, 0)$, which is an extreme point in the support of the equilibrium strategy. When a firm plays this vector, it wins the agent if and only if he focuses on attribute 1. Therefore, the firm's payoff is $\alpha \cdot [1 - \frac{1}{2}(0 + 2\alpha)] = \alpha(1-\alpha)$, hence we obtain the following corollary.

Corollary 1 *For any $\alpha \in [\frac{1}{2}, 1]$, industry profits in symmetric Nash equilibrium under the "no default" rule are $2\alpha(1-\alpha)$.*

Thus, under the "no default" rule, equilibrium industry profits become more competitive as attribute salience becomes more asymmetric. The intuition is that the agents' ex-post choice behavior becomes more homogenous, which means that competitive pressures are stronger. At the same time, more asymmetric attribute salience also implies a greater amount of obfuscation, in the sense that the range of values that $|q^2 - q^1|$ gets becomes wider as α gets closer to 1. We will discuss the meaning of this coupling of greater competitiveness with greater obfuscation in Section 4.

Comment: Rational-choice interpretation
When the "no default" rule is held fixed, the agents' choice behavior has a simple rational-choice interpretation: a fraction α^k of the agent population is genuinely interested only in attribute k, i.e. they have a well-defined utility function that is increasing in q^k and constant in q^{-k}. For this interpretation to be valid, it is important to assume that the bundling of the two quality attributes is intrinsic to the product, such that competitive pressures will not lead to their "unbundling". And of course it is implausible when the two attributes are, say, price and overall quality. From this perspective, it is not surprising that when discrimination is impossible, greater heterogeneity in consumer preferences (captured by shifting α toward $\frac{1}{2}$) results in a less competitive market outcome. However, the rational-choice interpretation does not fit our model on the whole, as explained in the previous section.

Spurious attributes
Throughout this paper, we view the two attributes as intrinsic features of the product. However, suppose that the framing of products as if they have two attributes is spurious, in the sense that as far as firms are concerned, product quality is fully characterized by the scalar $x = q^1 + q^2$. In particular, their cost is purely a function of x. For instance, q^k could be interpreted as provision of a certain quantity in a state indexed by k, where the only thing that relevant for the cost is the total quantity across the two states. Under this assumption, weak supermodularity of c implies that it is a weakly convex function of x, i.e. $c''(x) \geq 0$ for every $x \geq 0$.

Proposition 2 is can be easily extended to this case (as the proof makes transparent). When $\alpha > \frac{1}{2}$, the characterization is the same, except that some parameters are defined implicitly; and we omit it for brevity. When

$\alpha = \frac{1}{2}$, however, the extension is simple and conveys an insight (the proof of Proposition 2 covers this case).

Claim 1 *Let* $\alpha = \frac{1}{2}$ *and assume that* c *is a weakly convex function of* $x = q^1 + q^2$*. Under the "no default" rule, the game has a unique symmetric Nash equilibrium, in which firms play* $q^1 \sim U[0, x^*]$*, and* $q^2 = x^* - q^1$ *with probability one, where*

$$x^* = \frac{1 - c(x^*)}{c'(x^*)}$$

Thus, the total quality that firms offer in equilibrium is exactly what a monopolist maximizing $x(1 - c(x))$ would choose. If we interpret x as a quantity, this expression means quantity multiplied by the profit per unit sold, namely total profits. The firms' equilibrium strategy maximizes this function. In this sense, firms competing for trade-off avoiding agents under the "no default" rule with $\alpha = \frac{1}{2}$ behave as monopolists. Another way of interpreting x^* is that each firm i chooses x_i as if firm j randomly draws x_j from $U[0, 1]$ and the consumer rationally chooses the firm that offers the highest x. This is a "dual" interpretation: rather than positing that firm j plays x^* deterministically and the consumer chooses randomly (because of the random breakdown of x^* into the two quality components), here we assume that firm j chooses x_j randomly and the consumer chooses deterministically.

3.2 Opt Out

The "no default" rule with $\alpha = \frac{1}{2}$ is equivalent to "opt out", in terms of the firms' payoff function. Therefore, as far as the firms' behavior is concerned, our equilibrium analysis in this sub-section holds for the latter default rule as well. Moreover, a mixed default rule - by which a fraction λ of the agents obey "no default" while the remaining agents obey "opt out" - is payoff-equivalent to a pure "no default" rule with a modified salience parameter $\alpha' = \alpha\lambda + \frac{1}{2}(1-\lambda)$. Note, however, that the distinction between the two rules is significant for the description of the agents' equilibrium behavior. Under "opt out", an agent switches away from his default option only if the other market alternative dominates it. Since this never happens in equilibrium, we have the following corollary.

Corollary 2 *Under the "opt out" rule, no consumers ever switch away from their default option in symmetric Nash equilibrium.*

Thus, while "opt out" ensures full market participation, it also implies no switching in equilibrium. In contrast, the notion of switching is of course meaningless under "no default".

3.3 Opt In

Let us now assume that a fraction $\lambda > 0$ of the agents are initially assigned to the outside option as a default - i.e. they are under the "opt in" regime, while every other agent obeys "opt out" or "no default". Throughout this section, we assume that c is additively separable, i.e. $c(q^1, q^2) = \frac{1}{2}(q^1 + q^2)$. Let us first consider the extreme case of $\lambda = 1$.

Proposition 3 *When $\lambda = 1$, firms play $q^1 + q^2 = 2$ with probability one in symmetric Nash equilibrium, and agents choose the outside option.*

The intuition is that under a pure "opt in" rule, agents participate in the market only when they face an easy choice. Thus, firms have nothing to gain from making comparison between market alternatives hard. This in turn implies strong competitive pressures, which raise quality along both dimensions and push profits to zero. However, since in equilibrium neither market alternative dominates the other, all agents end up sticking to their default alternative. The market outcome is somewhat paradoxical: the offered products are highly attractive, yet no agent picks any of them. This is of course a consequence of the extreme assumption that *all* agents respond to difficult choices by sticking to their default, while none make active choices unless forced to.

Proposition 3 highlights the interpretational difficulty pointed out in Section 2: if the symmetric equilibrium strategy is pure, this means that both firms offer the same product with probability one, and in this case the assumption that agents act like Buridan's Ass seems less plausible. Of course, the result does not preclude the possibility that the symmetric equilibrium strategy is a smooth density over the line $q^1 + q^2 = 2$, in which case the critique does not apply.

The case of $\lambda \in (0, 1)$ turns out to be considerably more difficult to analyze. Let us begin with the following observation.

Proposition 4 *When $\lambda \in (0, 1)$, market alternatives dominate one another with positive probability in any symmetric Nash equilibrium.*

Thus, when some (but not all) agents obey the "opt-in" rule, easy choices - and thus switching away from defaults - occur with positive probability in any symmetric Nash equilibrium. This means that the structure of equilibrium strategies is more elaborate than under the pure default regimes, since q^2 and q^1 can no longer be linked deterministically.

Let us restrict attention to the case in which all "non-opt-in" agents are under the "opt out" regime, i.e. each firm plays the role of a default option for a fraction $(1-\lambda)/2$ of the agents. (If some of the non-opt-in agents are under the "no default" regime, firms' payoffs will not change.) We present a class of symmetric equilibria. For this purpose, we introduce some new notation. First, we represent a pure strategy (q^1, q^2) by the pair (p, e), where $p = 1 - \frac{1}{2}(q^1 + q^2)$ is the *profit* that the quality pair generates for the firm conditional on being chosen, and $e = \frac{1}{2}(q^1 - q^2)$ represents the amount of *obfuscation* the strategy exhibits. Second, for any positive integer n, denote

$$\sigma = \frac{1-\lambda}{2} \cdot \left(1 - \frac{1}{n}\right)$$

The interpretation of σ is simple: it is the mass of "captured agents" that each firm would enjoy - namely, the fraction of agents who would choose the firm by default - if the probability of easy choices were $\frac{1}{n}$.

Let $d > 0$ and let $n \geq 3$ be an integer. Define $s^*(d,n)$ to be a mixed strategy that consists of independent randomizations over p and e, where:

(i) p is distributed according to the *cdf*

$$G(p) = (1 + \sigma n)\left(1 - \frac{d\sigma n}{p}\right)$$

defined over the interval $[d\sigma n, d\sigma n + d]$.

(ii) e is uniformly distributed over the discrete set

$$\left\{d\left(k - \frac{1}{2}(n-1)\right)\right\}_{k=0,1,\ldots n-1}$$

The following diagram represents the support of $s^*(0.4, 3)$ in the original (q^1, q^2) space.

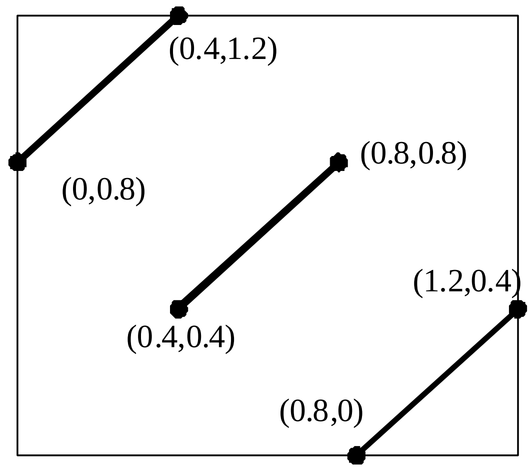

Figure 1

The support consists of three 45-degree line segments. Each segment corresponds to one of the three values that e can get, $-0.4, 0, 0.4$. Average quality gets values in $[0.4, 0.8]$, whereas quality along each dimension gets values in $[0, 1.2]$. If both firms play this strategy, domination occurs only within a segment, hence the probability of easy choices is $\frac{1}{3}$.

Claim 2 *Let $\lambda \in (0,1)$. If $d = \frac{2}{\lambda + n(2-\lambda)}$ and $n \in [1 + \frac{1}{\lambda}, 1 + \frac{3}{\lambda}]$, then $s^*(d,n)$ is a symmetric Nash equilibrium strategy.*

Let us elaborate on the properties of this class of equilibrium strategies.

Structure of the support and domination probability. The support of the equilibrium strategy is divided into n line segments, which are vertical in the (p,e) representation (they have a slope of $+1$ in (q^1,q^2) space, as in Figure 1). Each segment corresponds to a different value of e. The distance between adjacent segments is d, which is also the range of values that p can get. Therefore, domination occurs only within each segment - i.e., two realizations $(p_1,e_1),(p_2,e_2)$ constitute an easy choice if and only if $e_1 = e_2$. The probability agents face an easy choice is thus $\frac{1}{n}$.

Switching rates and market participation. The characterization of domination probabilities implies that the switching rate is $\frac{1}{n}$ $(\frac{1}{2n})$ for opt-in (opt-out) agents. The restriction on the values that n can get implies that the switching rate is at most $\lambda/(1+\lambda)$ for opt-in agents, and half that for opt-out agents. The overall equilibrium market participation rate for any $\lambda \in (0,1)$ is thus bounded from above by

$$(1-\lambda)+\lambda\cdot\frac{\lambda}{1+\lambda}=\frac{1}{1+\lambda}$$

Because n only gets integer values, these upper bounds on participation and switching are not tight. The maximal switching rate overall is $\frac{1}{3}$.

Quality. The marginal distributions over q^1 and q^2 are identical, with support $[0, dn]$. The upper bound of this interval is strictly above 1. The expected equilibrium average quality is

$$1-d\sigma n(1+\sigma n)\ln\left(\frac{1+\sigma n}{\sigma n}\right)$$

It can be verified that this is strictly greater than $\frac{1}{2}$. That is, expected average quality is higher than in pure "opt out". In other words, assigning some agents to the outside option makes the equilibrium market outcome more competitive, in the sense that expected quality is higher.

Limit equilibria and consumer welfare. As λ tends to 0 (approaching a pure "opt-out" rule), the permissible values of n diverge, and the collection of line segments becomes infinitely dense, approximating the line $q^1+q^2=1$. Equilibrium switching rates thus converge to zero. On the other hand, there are two limit equilibrium distributions when $\lambda \to 1$. In both of them, $q^1+q^2=2$ with probability one; in one of them, $n=3$, such that e is uniformly distributed over $\{-\frac{1}{2},0,\frac{1}{2}\}$; while in the other, $n=4$, such that e is uniformly distributed over $\{-\frac{3}{5},-\frac{1}{5},\frac{1}{5},\frac{3}{5}\}$. Thus, when the pure "opt-in" rule is slightly perturbed, the rate of market participation is at most $\frac{1}{3}$. If we define a consumer's welfare as the average quality he ends up getting, then the limit equilibrium for $\lambda \to 0$ induces a consumer surplus of $\frac{1}{2}$, whereas the limit equilibria for $\lambda \to 1$ induce a consumer surplus of at most $\frac{1}{3}$. Thus, given the set of equilibrium strategies we have focused on, "opt out" is superior to "opt in" in terms of consumer welfare: the increase in market participation outweighs the decrease in the quality of equilibrium products.

The class of equilibria under consideration turn out to be fully characterized by two properties that can be distilled from the above description. A mixed strategy s satisfies *independence* if it induces statistically independent distributions over p and e. We say that s satisfies *constant comparability* if $\Pr\{(p_1, e_1) \text{ dominates } (p_2, e_2) \mid (p_1, p_2)\}$ is the same for almost all (p_1, p_2), where (p_1, e_1) and (p_2, e_2) are two independent draws from s.

Proposition 5 *If a symmetric Nash equilibrium strategy satisfies independence and constant comparability, it must take the form $s^*(d, n)$, where $d = \frac{2}{\lambda+n(2-\lambda)}$.*

While the two properties lack an a priori justification, they are of interest because they provide a link to other models of price competition under limited comparability. In Varian (1980), the fraction of consumers who make price comparisons is assumed to be an exogenous constant. Therefore, equilibria in our model that exhibit constant comparison probability may be viewed as a "foundation" for this constant. In Piccione and Spiegler (2012), the two properties are logically linked by an underlying property of the comparability structure.

4 Conclusion

The robust empirical finding that people tend to stick to default options in the presence of difficult choices implies that "default architecture" can have dramatic implications for eventual choice patterns. This has led researchers (e.g. Thaler and Sunstein (2008), Beshears et al. (2012)) to advocate switching from opt-in default regimes to "no default" (i.e. active choice) or "opt-out" rules, in settings such as retirement savings by employees. It should be emphasized that in these settings, spontaneous and direct competition among firms over employees is not the norm; instead, contracts are mediated by employers. In other words, these markets are regulated de facto, except that the regulator is not the government but the employer. Nevertheless, it is interesting to speculate about the equilibrium effects of redesigning default options when spontaneous competition *is* the norm.

The analysis of our model implies that in such cases, a shift toward opting out or active choice (captured by a switch from high λ to low λ) does indeed raise the overall level of market participation, but at the same time it lowers expected quality and switching rates. The intuition is that the opt-out rule gives firms greater market power; because they benefit from consumer inertia, they have a stronger incentive to induce it by creating difficult choices for consumers.

Since the shift from "opt in" to "opt out" reduces the maximal probability of easy choices that is possible (in the class of equilibria we have focused on), there is a sense in which it increases the "mental cost" that consumers experience as a result of difficult choices. On the other hand, one could argue that since consumers choose by default whenever they face a difficult choice, they never actually incur this mental cost. The problem is that the mental cost does not have any revealed-preference manifestation in our model, because the default rule is *imposed* on agents. Arguably, if agents could express preferences over default rules, they would opt for a rule that saves the mental cost. It would be interesting to construct models of such high-order preferences.

A broad lesson from our exercise is that when we wish to analyze regulatory interventions that address consumer decision errors, it is important to have an explicit procedural model of consumer choice, which provides a concrete "story" behind the consumers' errors, and enables us to speculate about the market equilibrium's response to the intervention. For further exploration of this theme, see Spiegler (2014).

References

[1] Anderson, C. J. (2003): "The Psychology of Doing Nothing: Forms of Decision Avoidance Result from Reason and Emotion," Psychological Bulletin129, 139-167.

[2] Beshears, J., J. J.Choi, D. Laibson and B.C. Madrian (2013): "Simplification and Saving," Journal of Economic Behavior & Organization 95, 130-145.

[3] Bewley, T. (1986) Knightian Uncertainty Theory: Part I, Cowles Foundation Discussion Paper No. 807.

[4] Bordalo, P., N. Gennaioli and A. Shleifer (2013): "Salience and Asset Prices," American Economic Review Papers and Proceedings 103, 623-628.

[5] Bordalo, P., N. Gennaioli and A. Shleifer (2013): "Salience and Consumer Choice," Journal of Political Economy 121, 803-843.

[6] Carlin, B. I. (2009): "Strategic Price Complexity in Retail Financial Markets," Journal of Financial Economics 91, 278-287.

[7] Chioveanu, I. and J. Zhou, (2013): "Price Competition with Consumer Confusion," Management Science 59, 2450–2469.

[8] Dean. M. (2008): "Status Quo Bias in Large and Small Choice Sets", Mimeo, NYU.

[9] Gabaix, X. and D. Laibson (2006): “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets,” Quarterly Journal of Economics, 121, 505–540.

[10] Iyengar, S. S., G. Huberman, and W. Jiang. (2004): "How Much Choice is Too Much: Determinants of Individual Contributions in 401K Retirement Plans." In Pension Design and Structure: New Lessons from Behavioral Finance, ed. O. S. Mitchell and S. Utkus, 83-95. Oxford: Oxford University Press.

[11] 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.

[12] Kőszegi, B., and A. Szeidl (2013): "A Model of Focusing in Economic Choice," The Quarterly Journal of Economics 128, 53-104.

[13] Luce, M. F., J. W. Payne, and J. R. Bettman (1999): "Emotional Trade-off Difficulty and Choice," Journal of Marketing Research, 36, 143-159.

[14] Madrian, B. C., and D. F. Shea (2001): "The Power of Suggestion: Inertia in 401 (k) Participation and Savings Behavior," The Quarterly Journal of Economics 116, 1149-1187.

[15] Manzini, P., and M. Mariotti (2007): "Sequentially Rationalizable Choice," American Economic Review 97, 1824-1839.

[16] Masatlioglu, Y., and E. A. Ok (2005): "Rational choice with status quo bias", Journal of Economic Theory 121, 1-29.

[17] Ok, E. A.(2002): "Utility Representation of an Incomplete Preference Relation," Journal of Economic. Theory 104, 429–449.

[18] Payne, J. W., J. R. Bettman and E. J. Johnson (1993): The Adaptive Decision Maker. Cambridge University Press.

[19] Papi, M. (2014): "Noncompensatory Consideration and Compensatory Choice: An Application to Stackelberg Competition," Economic Theory Bulletin (forthcoming).

[20] Piccione, M. and R. Spiegler (2012): "Price Competition under Limited Comperability," *Quarterly Journal of Economics*, 127, 97-135.

[21] Rubinstein, A. (1988): "Similarity and Decision-Making under Risk (Is there a Utility Theory Resolution to the Allais Paradox?)," Journal of Economic Theory 46, 145-153.

[22] Spiegler, R. (2006): "Competition over Agents with Boundedly Rational Expectations," Theoretical Economics 1, 207-231.

[23] Spiegler, R. (2011): Bounded Rationality and Industrial Organization. Oxford University Press, New York.

[24] Spiegler, R. (2014): "On the Equilibrium Effects of Nudging," Mimeo, Tel Aviv University.

[25] Thaler, R. H. and C. R. Sunstein (2008): Nudge: Improving Decisions about Health, Wealth, and Happiness. Yale University Press.

[26] Tversky, A. (1972): "Elimination by Aspects: A Theory of Choice," Psychological Review 79, 281-299.

[27] Varian, H. R. (1980): "A Model of Sales," American Economic Review 70, 651-659.

5 Appendix: Proofs

Let us introduce some notation that will serve us in what follows. A symmetric mixed equilibrium strategy is a probability measure μ over the set

$\{(q^1, q^2) \mid c(q^1, q^2) \leq 1\}$. Let F^k denote the marginal *cdf* over q^k induced by μ. That is,

$$F^1(q^1) = \int_{q^2} \int_{r^1 \leq q^1} d\mu(r^1, q^2)$$
$$F^2(q^2) = \int_{q^1} \int_{r^2 \leq q^2} d\mu(q^1, r^2)$$

Let F^{k-} be the left limit of F^k, i.e. $F^{k-}(x) = \lim_{y \to x^-} F^k(y)$. Let $\bar{q}^k$ denote the supremum of the support of F^k.

When c is additively separable, namely $c(q^1, q^2) = \frac{1}{2}(q^1 + q^2)$, we will sometimes denote $p^k = 1 - q^k$ and interpret it as the profit, or "price", associated with dimension k, and use $p = \frac{1}{2}(p^1 + p^2)$ to denote the firm's profit conditional on being chosen. In the same vein, we sometimes denote $p(q) = 1 - c(q)$.

5.1 Proposition 1

Let us first establish that μ is continuous, such that $F^{k-} \equiv F^k$. Assume, in contradiction, that w.l.o.g F^1 contains an atom on some q^1. Consider the lowest q^2 such that the support of μ contains (q^1, q^2). If $c(q^1, q^2) = 1$ then firms make zero profits in equilibrium. However, this is impossible when $\alpha \in (\frac{1}{2}, 1)$, because firms can secure a strictly positive profits, by playing a mixed strategy with full support on $\{(q^1, q^2) \mid c(q^1, q^2) \leq 1\}$. Now consider the case in which $c(q^1, q^2) < 1$. Here a conventional "undercutting" argument applies: if a firm deviates to $(q^1 + \varepsilon, q^2)$, where $\varepsilon > 0$ is arbitrarily small, the increase in the firm's probability of being chosen overweights its loss in profit conditional on being chosen.

Next, we show that the supports of F^1 and F^2 contain no gaps - that is, for each $k = 1, 2$, the support of F^k is $[0, \bar{q}^k]$, where $\bar{q}^k > 0$. Assume the contrary. W.l.o.g, let $[a, b]$ be a maximal interval such that $F^1(q^1) = c$ for any $q^1 \in [a, b]$. Let (b, q^2) be some element in the support of μ. If a firm deviates to (q^1, q^2) where $q^1 \in (a, b)$, the firm's probability of being chosen does not change, but its profit conditional on being chosen increases, hence the deviation is profitable.

Having established these two properties of F^1 and F^2, we reach the key argument in the proof. Suppose that the support of μ contains two pairs $q = (q^1, q^2)$, $r = (r^1, r^2)$, such that $r > q$. By the previous argument, $F^k(r^k) - F^k(q^k) = \varepsilon^k > 0$. Both q and r are best-replies to μ. In order for deviations to either (q^1, r^2) or (q^2, r^1) to be unprofitable, we must have

$$\begin{aligned}
(1 - c(q^1, r^2))(\alpha^1 F^1(q^1) + \alpha^2 F^2(r^2)) &\leq (1 - c(q^1, q^2))(\alpha^1 F^1(q^1) + \alpha^2 F^2(q^2)) \\
(1 - c(r^1, q^2))(\alpha^1 F^1(r^1) + \alpha^2 F^2(q^2)) &\leq (1 - c(r^1, r^2))(\alpha^1 F^1(r^1) + \alpha^2 F^2(r^2))
\end{aligned}$$

where the R.H.S (L.H.S) of the first inequality is the payoff from q $((q^1, r^2))$, and the R.H.S (L.H.S) of the second inequality is the payoff from r $((r^1, q^2))$. Adding up the two inequalities and rearranging, we obtain

$$\begin{aligned}
& \left[\alpha^1 F^1(q^1) + \alpha^2 F^2(q^2)\right]\left[c(q^1, r^2) + c(q^2, r^1) - c(q^1, q^2) - c(r^1, r^2)\right] \\
\geq\ & \alpha^1 \varepsilon^1 \left[c(r^1, r^2) - c(r^1, q^2)\right] + \alpha^2 \varepsilon^2 \left[c(r^1, r^2) - c(q^1, r^2)\right]
\end{aligned}$$

While the R.H.S of this inequality is strictly positive, weak supermodularity of c implies that the L.H.S is non-positive (note that $q = (q^1, r^2) \wedge (r^1, q^2)$, $r = (q^1, r^2) \vee (r^1, q^2)$), a contradiction.

Since we have established that the support of F^k, $k = 1, 2$, is $\left[0, \bar{q}^k\right]$, where $\bar{q}^k > 0$, and that the support does not contain points that dominate one another, it follows that the support is a continuous curve that connects $(0, \bar{q}^2)$ and $(\bar{q}^1, 0)$.

5.2 Proposition 2

By Proposition 1, we can describe the support of μ by a continuous and strictly decreasing function $g : [0, \bar{q}^1] \to [0, \bar{q}^2]$, where for each q^1 in the support of F^1, $g(q^1)$ is the unique q^2 for which (q^1, q^2) is in the support of μ. Therefore,

$$F^2\left(g\left(q^1\right)\right) = 1 - F^1\left(q^1\right) \tag{1}$$

for every $q^1 \in [0, \bar{q}^1]$. Since F^1 and F^2 are strictly increasing, they are differentiable almost everywhere, such that the slope of g is

$$g' = -\frac{dF\left(q^1\right)/dq^1}{dF^2\left(q^2\right)/dq^2} \tag{2}$$

for almost every (q^1, q^2) along the graph of g.

Let us now write down an individual firm's payoff function when the opponent plays μ:

$$\pi\left(q^1, q^2\right) = \left[1 - c(q^1, q^2)\right]\left[\alpha^1 F^1(q^1) + \alpha^2 F^2(q^2)\right]$$

In equilibrium, first-order conditions must hold. Thus, for both $k = 1, 2$, the equation

$$[\alpha^1 F^1(q^1) + \alpha^2 F^2(q^2)] \cdot \frac{\partial c(q^1, q^2)}{\partial q^k} = \alpha^k \cdot \frac{dF^k(q^k)}{dq^k} \cdot [1 - c(q^1, q^2)] \tag{3}$$

must hold almost everywhere along the graph of g. Let us now invoke the property that c is purely a function of $q^1 + q^2$, and with a sight abuse of notation, write $c(q^1 + q^2)$. Then, the L.H.S of the equations for $k = 1$ and $k = 2$ are identical, and so we obtain

$$\frac{dF^1\left(q^1\right)/dq^1}{dF^2\left(q^2\right)/dq^2} = \frac{\alpha^2}{\alpha^1}$$

By (2), we conclude that

$$g'\left(q^1\right) = -\frac{\alpha^2}{\alpha^1}$$

almost everywhere along the graph of g. Therefore, we can write g as follows:

$$g(q^1) = \bar{q}^2 - \frac{\alpha^2}{\alpha^1} q^1 \tag{4}$$

Let us now distinguish between two cases.

(i) $\alpha^1 = \frac{1}{2}$. Then, $g' = -1$. This means that $\bar{q}^1 = \bar{q}^2 = q^1 + q^2$ and $c\left(q^1 + q^2\right)$ is constant for every (q^1, q^2) in the support of μ. Plug (1) into (3) and obtain the simplified equation

$$\frac{\partial c(q^1 + q^2)}{\partial q^k} = \frac{dF^k\left(q^k\right)}{dq^k} \cdot \left[1 - c(q^1 + q^2)\right]$$

Since $c\left(q^1 + q^2\right)$ is constant, $\frac{\partial c(q^1+q^2)}{\partial q^k}$ is constant as well. Thus, $\frac{dF^k\left(q^k\right)}{dq^k}$ is constant as well, which implies that q^k is distributed uniformly over $\left[0, \bar{q}^k\right]$ and

$$\frac{dF^k\left(q^k\right)}{dq^k} = \frac{1}{\bar{q}^k}$$

for almost every $q^k \in [0, \bar{q}^k]$ and (3) can be written as:

$$c'(x) \cdot x = 1 - c(x)$$

where $x = \bar{q}^k$. Finally, plug $c(q^1 + q^2) = \frac{1}{2}(q^1 + q^2)$, and obtain $x = 1$, which pins down the characterization.

(ii) $\alpha^1 \in (\frac{1}{2}, 1)$. The two extreme points in the support, $(0, \bar{q}^2)$ and $(\bar{q}^1, 0)$, must both generate the equilibrium payoff:

$$\alpha^1 \cdot (1 - c(\bar{q}^1 + 0)) = \alpha^2 \cdot (1 - c(0 + \bar{q}^2)) = \pi \tag{5}$$

The two points are also linked by (4), if we plug $g(\bar{q}^1) = 0$. Combining these two equations, we obtain a solution for $\bar{q}^1$, $\bar{q}^2$ and for the equilibrium payoff π. Moreover, according to (4), every realization of total cost $c\left(q^1 + q^2\right)$ in this interval is associated with a unique (q^1, q^2), as given in the statement of the proposition. Let us derive F^1. Since every (q^1, q^2) in the support of μ must be a best-reply, we must have that for every $q^1 \in [0, \bar{q}^1]$:

$$\left[1 - c(q^1 + g(q^1))\right] \cdot \left[\alpha^1 F^1\left(q^1\right) + \alpha^2 F^2\left(g(q^1)\right)\right] = \pi \tag{6}$$

Since every q^1 is associated with a unique $c = c(q^1 + g\left(q^1\right))$ which increases with q^1, $F^1(q^1) = G(c(q^1))$, where G is the induced *cdf* over c. Plugging (1) and (4) into (6), we obtain an explicit expression for F^1 over $[0, \bar{q}^1]$, and hence also for G:

$$G(c) = \frac{1}{2\alpha - 1}\left[\frac{\pi}{1-c} - (1-\alpha)\right]$$

Let us now plug $c(q^1 + q^2) = \frac{1}{2}(q^1 + q^2)$. By (4) and (5), $\bar{q}^1 = 2\alpha$ and $\bar{q}^2 = 2(1-\alpha)$, such that the equilibrium payoff is $\pi = \alpha(1-\alpha)$. This pins down G and g, hence the values that c can get, as well as the values of (q^1, q^2) as a function of c.

The last step is checking that there are no profitable deviations. It suffices to consider deviations to pure strategies $(q^1, q^2) \in [0, \bar{q}^1] \times [0, \bar{q}^2]$. It is easy to verify that given the explicit expressions for F^1 and F^2, the payoff function

$$\pi\left(q^1, q^2\right) = \left[1 - c(q^1 + q^2)\right]\left[\alpha^1 F^1(q^1) + \alpha^2 F^2(q^2)\right]$$

is decreasing (increasing) in both arguments when (q^1, q^2) is above (below) the graph of g, hence the maximal payoff is obtained at the points along g.

5.3 Proposition 3

Assume the contrary. Fix some symmetric Nash equilibrium, and let $x^* < 2$ denote the lowest value of $q^1 + q^2$ in the support of the equilibrium strategy μ. Consider some (q^1, q^2) in the support of μ for which $q^1 + q^2 = x^*$. By definition, (q^1, q^2) does not dominate any quality pair in the support of μ. Since $\lambda = 1$, (q^1, q^2) generates zero profits, hence this is the firms' equilibrium payoff. However, if firm 1, say, deviates by mixing with full support over $\{(r^1, r^2) \mid x^* < r^1 + r^2 < 2\}$, it will dominate firm 2's quality pair with positive probability, and thus make strictly positive profits, a contradiction.

5.4 Proposition 4

First, note that w.l.o.g, we can assume that the agents who are not initially assigned to the outside option as a default choose purely according to q^k with probability α^k, where $\alpha^1 + \alpha^2 = 1$ and $\alpha^1 \geq \frac{1}{2}$. In particular, when $q_1^1 > q_2^1$ and $q_1^2 < q_2^2$, firm k is chosen with probability $\alpha^k(1-\lambda)$. Second, the proof that in symmetric Nash equilibrium, F^1 and F^2 are atomless and their supports contain no "holes" (see the proof of Proposition 1) can be extended to the $\lambda > 0$ case. We omit the proof for brevity - the arguments essentially are the same.

Now suppose that contrary to the claim, there is a symmetric Nash equilibrium in which domination occurs with zero probability. Then, we can describe the support of the equilibrium strategy μ by a continuous and strictly decreasing function $g : [0, \bar{q}^1] \to [0, \bar{q}^2]$, where for each q^1 in the support of F^1, $g(q^1)$ is the unique q^2 for which (q^1, q^2) is in the support of μ. Thus, $F^1(q^1) = 1 - F^2(q^2)$ for (q^1, q^2) along the support.

Consider a pure strategy (q^1, q^2) in the interior of the support of μ. The payoff from this strategy is

$$\left[1 - \frac{1}{2}(q^1 + q^2)\right] \cdot (1-\lambda) \cdot \left[\alpha^1 F^1(q^1) + \alpha^2 F^2(q^2)\right] \tag{7}$$

hence this is the firms' equilibrium payoff. Let q, r be two points in the support of F. Assume w.l.o.g $q^1 < r^1$, $q^2 > r^2$. Consider deviations to $q \wedge r = (q^1, r^2)$ and $q \vee r = (r^1, q^2)$. If a firm deviates to $q \wedge r$, its payoff will be

$$\left[1-\frac{1}{2}(q^1+r^2)\right]\cdot(1-\lambda)\cdot\left[\alpha^1F^1(q^1)+\alpha^2F^2(r^2)\right]$$

On the other hand, if the firm deviates to $q\vee r$, its payoff will be

$$\left[1-\frac{1}{2}(r^1+q^2)\right]\cdot\left[(1-\lambda)\left(\alpha^1F^1(r^1)+\alpha^2F^2(q^2)\right)+\lambda\left(F^1(r^1)-F^1(q^1)\right)\right]$$

In order for μ to be an equilibrium, both expressions must be weakly below the payoff at q, which is the same as the payoff at r.

Denote $A=F^1\left(q^1\right)$, $B=F^1\left(r^1\right)-F^1\left(q^1\right)$ and $C=1-F^1\left(r^1\right)$. Then, the payoffs at the four points q, r, $q\wedge r$ and $q\vee r$ can be written as follows:

$$\begin{aligned}\pi\left(q\right) &= p\left(q\right)\left(1-\lambda\right)\left(\alpha^1A+\alpha^2B+\alpha^2C\right)\\ \pi\left(r\right) &= p\left(r\right)\left(1-\lambda\right)\left(\alpha^1A+\alpha^1B+\alpha^2C\right)\\ \pi\left(q\wedge r\right) &= p\left(q\wedge r\right)\left(1-\lambda\right)\left(\alpha^1A+\alpha^2C\right)\\ \pi\left(q\vee r\right) &= p\left(q\vee r\right)\left(1-\lambda\right)\left(\alpha^1A+\alpha^1B+\alpha^2B+\frac{\lambda}{1-\lambda}B+\alpha^2C\right)\end{aligned}$$

It follows that

$$\begin{aligned}&\pi(p)+\pi(q)-\pi\left(q\wedge r\right)-\pi\left(q\vee r\right)\\ =\;& B\left(1-\lambda\right)\left(\alpha^1\left(p\left(r\right)-p\left(q\vee r\right)\right)+\alpha^2\left(p\left(q\right)-p\left(q\vee r\right)\right)-\frac{\lambda}{1-\lambda}\right)\end{aligned}\tag{8}$$

Note that $B>0$. If q and r are sufficiently close, we have

$$\begin{aligned}p\left(r\right)-p\left(q\vee r\right) &< \frac{\lambda}{1-\lambda}\frac{1}{2\alpha^1}\\ p\left(q\right)-p\left(q\vee r\right) &< \frac{\lambda}{1-\lambda}\frac{1}{2\alpha^2}\end{aligned}$$

such that expression (8) is strictly negative, which means that the deviation to $q\vee r$ or $q\wedge r$ is profitable, a contradiction.

5.5 Claim 2

Let $d = \frac{2}{\lambda + n(2-\lambda)}$. The strategy $s^*(d,n)$ induces a marginal distribution over q^k, with support $[0, \bar{q}^k]$ where $\bar{q}^k = nd$, $k = 1,2$. Clearly, when we look for profitable deviations from $s^*(d,n)$, we need only look for pure strategies $(q^1, q^2) \in [0, \bar{q}^1] \times [0, \bar{q}^2]$. From now on, we adhere to the (p,e) representation of strategies. We index the n values that e by $k = 0, 1, ..., n-1$. Let $l = d\sigma n$ and $h = d(\sigma n + 1)$ denote the lowest and highest values in the support of the marginal distribution over p. Define $L^k = \{(p,e) \mid p \in [l,h] \text{ and } e = e^k\}$. That is, L^k is one of the n line segments that constitute the support of $s^*(d,n)$, which is associated with e^k. There are three cases to consider.

Case 1: Deviation to (p,e) where $p \geq h$.
For any $p \geq h$, it suffices to look for the most profitable deviation (p,e). The fact that e is uniformly distributed over evenly spaced values independently of p, and that $d = h - l$ and $d = e^k - e^{k-1}$, the total length of the $\{L^k\}$ segments that (p,e) is dominated by is independent of e. Moreover, the number of segments that partially dominate (p,e) is at most 2. Because of the concavity of G, it is more profitable to be partially dominated by one segment (the dominating prices on that segment being $[l, l+2x]$ for some x) than being partially dominated by 2 segments (in each the dominating prices are $[pl, l+x]$). This implies that for a given p the most profitable e maximizes the number of line segments L^k that entirely dominates (p,e). Therefore, in the sequel we restrict attention w.l.o.g to $e = 1-p$, i.e., to $(0,q)$, where $q < 1-h$, in the (q^1,q^2) representation.

Consider a deviation to $p = h + (m+x)\frac{d}{2}$, $m = 0, 1, .., n-2$, $x \in [0,1]$. The payoff is

$$\left(h + (m+x)\frac{d}{2}\right)\frac{1-\lambda}{2}\left(1 - \frac{1}{n}(1 + m + G(l + dx))\right)$$

Note that for $x = 0$, the payoff at $m = 0$ (which corresponds to no deviation) is higher than at $m=1$ if and only if $n \leq 1+\frac{3}{\lambda}$. Second, if this is the case, then the payoff continues to decrease for any $m>1$ ($n \leq 1+\frac{3}{\lambda}$ is a sufficient condition for the derivative of the payoff with respect to m is negative for $m>1$ and $x=0$).

Moreover, the derivative of the payoff function w.r.t x (for a given m) is increasing. Thus, for each m, the maximal payoff is achieved at $x \in \{0,1\}$.This, together with the previous result, imply that deviations to $p \geq h$ are unprofitable if and only if $n \leq 1+\frac{3}{\lambda}$

Case 2: Deviation to (p,e) where $p \leq l$.
By the same argument as in Case 1, the most profitable deviation for a given $p \leq l$ is to e that maximizes the number of entire segments L^k which are dominated by (p,e). Therefore, in the sequel we restrict attention w.l.o.g to $e=(l+e^0)-p$, i.e., to $(\bar{q}^1, q)$, where $q>1-l$, in the (q^1,q^2) representation.

Consider a deviation to $p=l-(m+x)\frac{d}{2}$, $m=1,..,n-2$, $x \in [0,1]$. The payoff is:

$$\left(l-(m+x)\frac{d}{2}\right)\left(\frac{1-\lambda}{2}+\frac{1+\lambda}{2}\frac{1}{n}\left(1+m+\left(1-G\left(h-dx\right)\right)\right)\right)$$

Note that the payoff at $m=0$ (corresponding to no deviation) is higher than at $m=1$ if and only if $n \leq 1+\frac{3-\lambda}{(1-\lambda)\lambda}$. Second, if this is the case, then the payoff continues to decrease for any $m>1$. Note that $n \leq 1+\frac{3}{\lambda}$ implies $n \leq 1+\frac{3-\lambda}{(1-\lambda)\lambda}$.

The derivative of this function w.r.t x implies the following: (i) it is increasing in x for $m \leq \sigma n-1$; (ii) it is negative for $m>\sigma n-1$. Thus, it is enough to check for deviation to $x=0$ and $m \leq \sigma n-1$, and by the previous result, these deviation are unprofitable for $n \leq 1+\frac{3}{\lambda}$.

Case 3: Deviation to (p,e) where $l \leq p \leq h$.
Fix $p \in [l,h]$. Because any (p,e) where p is in this interval is comparable to points in at most 2 segments, and because all segments have the same

probability distribution, it is enough to check for deviations from (p, e^0) to $(p, e^0 + x)$, where $x \in \left(0, \frac{d}{2}\right)$. Thus, $(p, e^0 + x)$ is comparable only to points on L^0 and L^1. Consider these three cases:

(i) $p + x \leq h$ and $p - x \geq l$. In this case $(p, e^0 + x)$ is not dominating, nor being dominated by, any point in L^1. As x increases, $(p, e^0 + x)$ is dominated by less points on L^0 but also dominates less. The firm's net gain of market share is

$$\frac{1}{n}\frac{1-\lambda}{2}\left[G(p) - G(p-x)\right] - \frac{1}{n}\frac{1+\lambda}{2}\left[G(p+x) - G(p)\right]$$

Substituting G, we obtain the following condition for the deviation's profitability:

$$\frac{p+x}{p-x} > \frac{1+\lambda}{1-\lambda}$$

It is easy to verify that the L.H.S is maximized at $p = l + \frac{d}{2}$ and $x = \frac{d}{2}$, and the inequality is satisfies iff $n < 1 + \frac{1}{\lambda}$.

(ii) $p + x > h$. In this case $(p, e^0 + x)$ is dominated by some prices in L^0 and in L^1,but not dominating any point. Because the total length of the segments of L^0 and L^1 that dominate $(p, e^0 + x)$ is constant for any such x, the concavity of G implies that it is more profitable to be dominated by L^0 alone than by both. That is, this deviation is strictly less profitable than the deviation to $(p, e^0 + h - p)$ which is covered in case (i).

(iii) $p - x < l$. In this case $(p, e^0 + x)$ is dominating some prices in L^0 and in L^1, but not being dominated by any point. Because the total length of segments of L^0 and L^1 that $(p, e^0 + x)$ dominates is constant for any such x, the concavity of G implies that it is more profitable to dominated L^0 alone. That is, this deviation is strictly less profitable than the deviation to $(p, e^0 + p - l)$, which is covered in case (i) as well.

5.6 Proposition 5

Consider a symmetric Nash equilibrium strategy μ that satisfies independence and constant comparability. The feature that the induced marginal distribution over q^k has no atoms and no holes carries over to the present setting. From now on, we adhere to the (p, e) representation of pure strategies. The proof proceeds stepwise.

Step 1: The marginal distribution over p is atomless.
Proof: Assume the contrary - i.e., that some price p is realized with positive probability. Then, with positive probability $p_1 = p_2 = p$. In this case, (p_1, e_1) and (p_2, e_2) necessarily do not dominate one another. Thus, conditional on $p_1 = p_2 = p$, the probability of domination is zero. By constant comparability, the probability of domination must be zero in equilibrium, in contradiction to Proposition 4.

The following two steps state properties that hold for *almost all* pairs of realizations of a symmetric equilibrium strategy. For expositional convenience, we state and prove the claims with slight imprecision, as if they hold for *all* realizations.

Step 2: If (p', e') is dominated by (p, e), then (p'', e') is dominated by (p, e) for every $p'' \in (p, p')$.
Proof: Let $p' > p'' > p$ be three prices in the support of the marginal distribution over p. By definition, if (p'', e') is dominated by (p, e), then (p', e') is dominated by (p, e) as well. Now, calculate the probability of domination conditional on $(p_1, p_2) = (p', p)$, by integrating over all possible values of e_1, e_2, and do the same for $(p_1, p_2) = (p'', p)$. By independence, e_1 and e_2 are *i.i.d.* Therefore, if (contrary to the claim) there is positive probability that (p', e') is dominated by (p, e) yet (p'', e') is not dominated by (p, e), we will get a violation of constant comparability, because the domination probability conditional on (p', p) will be strictly higher than the domination probability conditional on (p'', p).

Step 3: For every (p, e) and (p', e') in the support of μ with $e \neq e'$, $|e' - e| \geq |p' - p|$.

Proof: Assume the contrary, i.e., $|e - e'| < h - l$ for e, e' in the support of the marginal obfuscation distribution (where h and l are as defined in the proof of Claim 2). By Step 1, we can find a price $p \in (l, h)$ in the support of the marginal distribution over p, such that $p - l < |e - e'|$. This means that (h, e) will be dominated by (l, e') and yet (p, e) will not be dominated by $(l.e')$, contradicting Step 2.

Step 4: The marginal equilibrium distribution over e is uniform with support $\{e^0, ..., e^{n-1}\}$, where $e^{k+1} - e^k = h - l$ for every $k = 0, ..., n-2$, and $h + e^{n-1} = h - e^0 = 1$.

Proof: Step 3 immediately implies that the gap between two adjacent realizations $e < e'$ cannot be less than $h - l$. Assume the gap is strictly greater than $h - l$. Then, a firm can profitably deviate from (h, e) to $(h, e + \delta)$, where $\delta > 0$ is arbitrarily small. The reason is that since the distribution over p is atomless, it assigns positive probability to prices arbitrarily close to h. Thus, by switching to $(h, e + \delta)$, the firm reduces the probability of being dominated by strategies (p, e) for $p < h$, without affecting the probability of being dominated by strategies (p, e''), $e'' \neq e$. Since the marginal distributions over q^k have no holes, $h + e^{n-1} = h + e^0 = 1$. Finally, the reason that the distribution is uniform is as follows. In equilibrium, firms are indifferent among all (h, e^k). By construction, the payoff from (h, e^k) is $h \cdot \frac{1-\lambda}{2} \cdot (1 - \Pr(e^k))$, because (h, e^k) is dominated by (p, e) if and only if $p < h$ and $e = e^k$.

Step 5: $h = 1 - \frac{n-1}{2}(h - l)$, $l = 1 - \frac{n+1}{2}(h - l)$.

Proof: Recall that $h + e^{n-1} = h - e^0 = 1$. Therefore, $e^0 = -e^{n-1}$. Since values of e are evenly spaced by intervals of length $h - l$, it follows that the distribution of e is symmetric around zero, such that $e^{n-1} = \frac{n-1}{2}(h - l)$, and the result follows.

To complete the proof, we add the equation that the profits at l and h

coincide:

$$h \cdot \frac{1-\lambda}{2} \cdot (1-\frac{1}{n}) = l \cdot \left[\frac{1-\lambda}{2} + (1-\frac{1-\lambda}{2}) \cdot \frac{1}{n}\right]$$

This equation, coupled with Step 5, gives us the solutions to h and l, as well as the equilibrium profit, as a well-defined function of n. We can retrieve the marginal distribution G over p from the following equation:

$$\frac{\sigma(2\sigma n+2)}{2\sigma n+n+1} = p \cdot \left[\frac{1-\lambda}{2} \cdot \left(1-\frac{1}{n}+\frac{1}{n}(1-G(p))\right) + (1-\frac{1-\lambda}{2}) \cdot \frac{1}{n} \cdot (1-G(p)\right]$$

Since this equation holds for every p in the support of G, the support of G cannot have holes inside $[l,h]$, for otherwise there would be an atom, contradicting Step 1.