Competition over Context-Sensitive Consumers *

Arno Apffelstaedt† Lydia Mechtenberg‡

Abstract: When preferences are sensitive to context, firms may influence purchase decisions by designing the environment of consumption choices. Confirming anecdotal evidence on retailer marketing tricks, we show that competitive retailers exploit context-sensitivity by designing choice environments that drive a wedge between preferences before and after entering a store. This wedge induces any consumer who slightly under-estimates her sensitivity to context to switch preference from a competitive bait product to a more profitable product at the store. Depending on the quality preferences and budget of consumers, the market is either in an up-selling equilibrium or in a down-selling equilibrium. In the former, firms attract consumers with low-quality products, compete on prices, and design context to ultimately sell a product of higher price. In the latter, firms attract consumers with high-price products, compete on quality, and design context to ultimately sell a product of lower quality. When modeling context-sensitivity according to the theories of *Salience* (Bordalo, Gennaioli and Shleifer, 2013), *Focusing* (Kőszegi and Szeidl, 2013), or *Relative Thinking* (Bushong, Rabin and Schwartzstein, 2016), designing context comes down to the introduction of a single decoy product. This decoy draws consumer attention at the store to the favorable attributes of the product the firm aims to sell. The exploitation of context-sensitive naïves is robust to the presence of sophisticated or rational consumers.

Keywords: Choice Context, Retailer Competition, Up-Selling, Down-Selling, Decoys

JEL Codes: D03, D11, D41

This Version: April 2017

*We are grateful to Pedro Bordalo, Tom Cunningham, Markus Dertwinkel-Kalt, Andrew Ellis, Nicola Gennaioli, Michael D. Grubb, Paul Heidhues, Botond Kőszegi, Francesco Nava, Matthew Rabin, Joshua Schwartzstein, Andrei Shleifer, and Adam Szeidl for extremely useful discussions and indispensable comments.

†Corresponding Author. Address: Universität Hamburg, Department of Economics, Von-Melle-Park 5, 20146 Hamburg, Germany. Phone: +49 177 4498 604. Email: arno.apffelstaedt@uni-hamburg.de.

‡Universität Hamburg, Department of Economics. Email: lydia.mechtenberg@uni-hamburg.de.

1 Introduction

Evidence that consumer choice is context-sensitive is abundant. Most people perceive \$10 for a given bottle of wine to be expensive when accompanied by cheaper alternatives (say, at a discount store), but cheap at an exclusive liquor store where alternatives cost \$20 on average. A range of promising theories have recently emerged to model such behavior, reflecting the observation that consumers judge alternatives relative to the immediate environment in which they are presented, among them the theories of *Salience* (Bordalo, Gennaioli and Shleifer, 2013), *Focusing* (Kőszegi and Szeidl, 2013), and *Relative Thinking* (Bushong, Rabin and Schwartzstein, 2016).

We study the optimal response of competitive firms to this well-known behavioral anomaly of consumers. Our model reflects the typical retail market structure: Each firm owns a store where it sells a line of alternative products (differentiated in quality and price) and competition is on consumer entry: Consumers first observe the product lines of all stores and then enter one firm to buy a product. The local choice context at the store can lead consumers to overvalue the quality or price of products relative to their outside assessment, depending on how the firm designs the product line. Consider yourself planning the purchase of that bottle of wine at home. Are you aware that you are likely to be willing to spend more money for a similar bottle when you enter a nice liquor store than when you purchase the wine at a discount supermarket? We show that if (and only if) consumers under-estimate—even just marginally—the effect of context on their choice, firms will exploit this bias by designing choice environments that drive a wedge between the preferences inside and outside of the store. Firms then use this wedge to compete for the consumer with an unprofitable *attraction product*, knowing that context effects will induce her to buy a more profitable *target product* at the store. When in-store context is modeled along the theories of Salience, Focusing or Relative Thinking, firms generate the preference distortion by presenting the consumer with a third option—a *decoy*—that, while being unattractive as an option itself, makes the target stand out in relative value at the store.

To put this prediction in the context of our example: While you might have been attracted to the liquor store in the belief of buying a competitively priced, medium-quality bottle of wine, you end up leaving the store with a considerably more expensive high-quality wine instead. In the jargon of marketing experts you have been “up-sold”. Up-selling is touted among these experts as one of the most powerful, not-to-be-missed marketing tricks and most consumers come across such attempts on a regular basis, for example when purchasing airline tickets.[1] Ellison and Ellison (2009) present evidence of up-selling in the online retail

[1]See, for example, Max Nisen on “Super cheap airline fares lures in lots of fliers, but most shell out to

market for computer parts. Facebook, Shopify and SAP offer up-selling software to make it easier for smaller retail firms to use such strategies.[2] Our novel prediction is that the "up-sell"—the switch from a cheaper to a more expensive product *inside* the store—is part of a wider marketing strategy that also includes the design of an adequate "bait" product to deal with competition *outside* the store—and, importantly, that naïve context-sensitivity may be at the core of many such phenomena.

There is considerable suggestive evidence for this claim to be true. One marketing blog talks of "[d]rawing people in with a low offer and then presenting them with better, more expensive options [of being] the bread and butter of upselling", making clear that the bait is as important as the switch.[3] Others describe up-selling as "getting the consumer to make a higher cost purchase than he or she *orginally planned*", selling "a product that is more expensive than the one *they initially came to buy*" or something more profitable "than the original product they *intended to buy*", hinting at the naïveté of consumers when selecting a firm.[4] Finally, while one is inclined to equate up-selling with pushy salespeople, marketing experts are aware that letting the consumer decide for herself and inducing the switch with a smart presentation of options and relative comparisons is the more subtle and successful way for an up-sell. In fact, many firms seem to inflate their product line with additional options to make the target product stand out in comparison and thereby draw consumers away from the (unprofitable) attraction product; a strategy that resonates with the classical, experimental literature on context and decoy effects and is also predicted by our model when context is modelled according to Salience, Focusing, or Relative Thinking.[5] One of the two firms studied by Ellison and Ellison (2009, see Figure 2, p.434) could also be argued to do just that.

The model makes more subtle novel predictions. One of them is that attracting consumers with a cheap, low-quality product and then inducing them to switch to a more expensive, higher quality target (the classical result associated with up-selling) is only one possible equi-

upgrade" (Quartz, 16th July 2015, retrieved from https://qz.com/456017, accessed 02-23-2017)

[2]See https://www.facebook.com/business/help/1604184966521384, https://apps.shopify.com/ultimate-upsell, and http://help-legacy.sap.com/saphelp_crm60/helpdata/en/46/6d7f1de28c7183e10000000a114a6b/content.htm (all three accessed 02-22-2017).

[3]See https://econsultancy.com/blog/66879-10-powerful-examples-of-upselling-online/ (accessed 02-22-2017).

[4]See www.forbes.com/sites/neilpatel/2015/12/21/how-to-upsell-any-customer, http://www.brainsins.com/en/blog/upselling-increasing-profits/1488, and https://www.123-reg.co.uk/blog/ecommerce/how-to-increase-revenue-with-up-selling-and-cross-selling/ (all three have emphasis added and were accessed 02-23-2017).

[5]For a good range of examples of firms using such strategies, see https://econsultancy.com/blog/66879-10-powerful-examples-of-upselling-online/ (accessed 02-22-2017). Two seminal papers on the effect of adding unwanted products to the choice set in order to increase the choice-probability of "target" products are Huber, Payne and Puto (1982) and Simonson (1989).

librium outcome. Depending on parameter values, firms may in fact find it more profitable to do the opposite, that is, to use a *down-selling* strategy. In this case, consumers expect to purchase an expensive, high-quality product when entering a store, but purchase a cheaper product of lower quality instead.[6] Context also works in the opposite way, making the consumer *more* (instead of less) price-sensitive at the store. We predict down-selling schemes to become more profitable as the maximum amount of money consumers are willing to spend increases: This allows firms to attract consumers with more expensive products, leading to a stronger (and thus more profitable) "bargain effect" when the consumer switches to a product of lower price. This finding resonates well with the anecdotal evidence on down-selling, which mainly associates retailers of up-scale, luxury products with the phenomenon.[7]

Because the exploitation targets naïve context-sensitive consumers, one might expect that our results are sensitive to the presence of sophisticated or rational consumers. We show in two extensions that this is not the case. The market reacts to sophisticated consumers by providing additional, non-distortionary stores where the consumer can commit to a product of her outside preference. In reality, no-frills discount stores such as Aldi in the market for grocery goods might serve such a purpose. Rational consumers, on the other hand, will enter the exploitative firms and re-exploit them by purchasing the non-profitable attraction product. However, this does not stop firms from using this practice. Instead, firms increase the exploitation of naïves in order to substitute for the losses made on rational consumers.

Theoretical contributions dealing with the question of how firms react to context-sensitivity in market settings are rare. Kamenica (2008) shows that, given that there is also uncertainty about the production cost, a monopolist may be able to "manipulate" the quality perception of rational, uninformed consumers by adding decoy products to the product line. While this is an important result that sheds new light on the importance of consumer inference, it is definitely not the end of the story. Context-effects have been found in experimental settings with no explanatory room for inference, see, e.g., Herne (1999), Ariely, Loewenstein and Prelec (2003), Mazar, Kőszegi and Ariely (2014) and Jahedi (2011). Moreover, the conjecture that context-sensitive shopping behavior is largely irrational seems corroborated by the extensive online discussion of context-related marketing techniques that all seem to "manipulate" or "trick" consumers into purchase decisions.

Earlier literature in behavioral economics has made the point that "context matters", but

[6]The down-sell is relative to the product the consumer was attracted with. Relative to the rational benchmark, firms may still be providing overly high quality.

[7]Christina Binkley makes a convincing case for this marketing strategy to be wide-spread in the high-fashion industry in her aptly named article "The Psychology of the $14,000 Handbag: How Luxury Brands Alter Shoppers' Price Perceptions; Buying a Keychain Instead" (The Wall Street Journal, 9th August 2007, retrieved from https://www.wsj.com/articles/SB118662048221792463, accessed 02-23-17).

has not formally studied its strategic role in competitive markets.[8] Instead, it has offered theories that are able to explain and model context-dependent preferences. Our model is sufficiently general to encompass these theories, and we produce results for the three most recent ones (Salience, Focusing, and Relative Thinking) in this paper. We highlight a hitherto unstudied strategic use of context that only exists in competitive markets: Designing choice environments that drive a wedge between consumer preferences in the moment of competition with other firms and preferences in the moment of purchase. It is this particular exploitation of naïve context-sensitivity that generates product lines with three distinct products for just one type of consumer: a "false competitor" (a.k.a. the attraction product), a target, and a decoy. Such choice sets have inspired early experimental research on context effects (see, in particular, Huber, Payne and Puto, 1982), and have been used as rationale to offer theories of context-dependent consumer choice (most recently by Bordalo, Gennaioli and Shleifer 2013 and Bushong, Rabin and Schwartzstein 2016), but their existence in markets has so far never been questioned nor explained.

The strategic use of in-store context we describe is very different to the role of "salience effects" for product choice in models of "direct" competition as studied by Bordalo, Gennaioli and Shleifer (2016) where consumers do *not* make their purchase decision in two steps. Most obviously, the main results of our paper stem from the possibility that preferences may change *after entering* the store of a particular firm and can therefore not be reproduced in a direct market. We discuss more subtle differences between Bordalo, Gennaioli and Shleifer (2016) and our paper in the conclusion of this paper. There are other papers in the literature on competition over biased consumers that feature a two-phase choice procedure by which consumers first select a firm and then a product. However, they do not allow the choice environment to affect consumer preferences. Some of these papers relate to ours by the idea that "marketing devices" or "frames" play a strategic role when attracting consumers (Eliaz and Spiegler 2011*a*, Eliaz and Spiegler 2011*b*, Piccione and Spiegler 2012), others more technically by the fact that there exists an element of naïve time-inconsistency that firms may try to exploit (among others, Gabaix and Laibson 2006, Ellison 2005, DellaVigna and Malmendier 2004, Heidhues and Kőszegi 2008, and Heidhues, Kőszegi and Murooka 2017). Our results are in many regards novel with regard to both of these streams. A more detailed discussion of our contribution to this literature is relegated to the conclusion.

The remainder of the paper is organized as follows. We introduce a formal model in the next section. In section 3 we first derive a rational benchmark and then carve out the major impact of assuming context-sensitivity in retail markets, which is the possibility of

[8] A notable exception is Bordalo, Gennaioli and Shleifer (2016), whose contribution in relation to ours we discuss further below.

firms to "fool" (i.e., up- or down-sell) naïve consumers. We also show in this section how the profitability of such strategies depends on the (partial) naïveté of consumers and the type of choice environment that the firm selects. Section 4 addresses the questions of *what* types of environment firms will construct in equilibrium and *how* such context is constructed when it is a function of the choice set as suggested by the theories of Salience, Focusing and Relative Thinking. Section 5 concludes with a discussion of our results with regard to model assumptions and highlighting differences between our findings and earlier results in the literature on competition over biased consumers.

2 A Model

A unit mass of consumers has demand for a good that can be differentiated in quality $q \in \mathbb{R}$ and price $p \in \mathbb{R}$, where quality and price are both measured in dollars. There is a minimum quality $\underline{q} > 0$ and a maximum price $b > 0$ agents are willing to accept and pay, respectively. Each consumer demands one good. There is a large number K of firms in the market. Each firm k owns a store. To purchase from firm k, a consumer has to enter its store. At the store, the firm can offer any menu of products J^k. Each product $j \in J^k$ implements the good at some level of quality $q_j \in \mathbb{R}$ and price $p_j \in \mathbb{R}$. The set $M^k = ((q_j, p_j))_{j \in J^K}$ is called the *product line* of firm k. Each firm k also chooses *how* to present its product line to consumers who enter its store. This choice is represented by the variable Θ^k, which we call the *in-store context* of firm k. (We define Θ^k in detail further below). Instead of entering a store and purchasing a product, consumers can select the outside option of no purchase. The sequence of events is as follows.

- Firms simultaneously commit to a product line M^k and an in-store context Θ^k.
- Each consumer then moves in two stages:
  - Stage 1: The consumer observes the product lines M^k of all firms and then decides to enter one store to make a purchase *or* to exercise the outside option and leave the market without purchase.
  - Stage 2: If the consumer has entered store k, she selects a product $j \in J^k$ in context Θ^k.[9]

[9] In principle, the consumer could enter multiple stores before purchasing a product. However, in comparison to search models, a consumer in our model does not gain information from entering multiple stores as she has full information about the *entire* choice set (i.e., the products of all firms) already in stage 1. Because firms commit to a product line before consumers move, the assumption of a two-phase time structure therefore does not lead to qualitatively different results than a more "realistic" structure. We discuss this

Context-Sensitive Consumers. When evaluating products *outside* stores, consumers value a product of given quality and price at all firms equally. Without loss of generality (henceforth w.l.o.g.), let this (global) surplus function be given by

$$u_j = q_j - p_j. \tag{1}$$

We assume (w.l.o.g.) that the outside option of no purchase generates surplus $u_0 = 0$. *Inside* a firm-specific store, the local valuation of firm k's products may differ from Equation (1) due to the consumer now being exposed to the local context of the store: Let Θ^k be a vector that has as many entries as the firm has products in the product line (i.e., $|\Theta^k| = |J^k|$). Element $\theta_j^k \in \Theta^k$ identifies the effect of local context at store k—for instance, the color of price-tags or the relative position of product j in the product line—on the valuation of product j. We assume, in particular, that in-store context can either increase the perceived quality ($\theta_j^k = Q$) or the perceived price ($\theta_j^k = P$) of a product, thereby leading to an inflation or deflation of product value relative to Equation (1). If the local context at store k has no influence on the valuation of product j, we write $\theta_j^k = N$. With a given context Θ^k, consumers then value products $j \in J^k$ inside store k with the surplus function

$$\hat{u}_j^k = \begin{cases} q_j - p_j & \text{if } \theta_j^k = N, \\ \beta q_j - p_j & \text{if } \theta_j^k = Q, \\ q_j - \beta p_j & \text{if } \theta_j^k = P, \end{cases} \tag{2}$$

where $\beta \geq 1$ measures the size of contextual distortions; the possibility of $\beta = 1$ nests the rational model.[10]

When making their entry decision in stage 1, consumers observe the product lines of all stores and form an expectation about their purchase in stage 2. This expectation depends, of course, on the consumer's awareness of possible preference distortions at the store. We allow for different types. A perfectly *sophisticated* type knows Θ^k and β, and will therefore always predict her behavior correctly. On the other end there is a perfectly *naïve* type who is either (completely) unaware of context effects or (falsely) believes that her valuations are consistent across different contexts. We capture these two extremes as well as their convex

point in more detail in section 5 after we have presented the results to our (in that sense "reduced-form") set-up. Note also that our results remain qualitatively unchanged if we would introduce an outside option in stage 2.

[10]Note, importantly, that we do *not* claim that the outside assessment of products is free of distortions. The crucial element of our model is not the particular form of Equation (1), but that—once that consumers have entered the store—preferences may change *relative* to this outside assessment. Our results go through for any limitation of "local" context effects to small values, i.e., for any β arbitrarily close to 1.

hull by assuming that all consumers are aware of the environment at firm k (and thus of Θ^k), but are heterogeneous in their belief about the size of β. Specifically, each consumer has a point belief $E(\beta) = \tilde{\beta}$ and predicts herself to value products inside of store k with the surplus function

$$E_{\tilde{\beta}}\left[\hat{u}_j^k\right] = \hat{u}_j^k|_{\beta=\tilde{\beta}}. \tag{3}$$

The distribution of types in the population is $F(\tilde{\beta})$, with density $f(\tilde{\beta})$. We assume $f(\tilde{\beta}) = 0$ for any $\tilde{\beta} \leq 1$, which implies that agents may mispredict the size of contextual distortions, but not the direction. The lower bound $\tilde{\beta} = 1$ identifies the perfectly naïve type. Note that any type $\tilde{\beta} \neq \beta$ is *partially* naïve. We subdivide naïves into *under-estimators* ($\tilde{\beta} < \beta$) and *over-estimators* ($\tilde{\beta} > \beta$) of the effect of context on their choice. This categorization will play an important role for market supply in equilibrium.

Firms. Each firm maximizes its profit π^k by choosing a product line and a context for its store. For large parts of the paper it will be sufficient to consider a reduced form model in which the firm chooses the distortion Θ^k (i.e., whether a product is quality- or price-inflated inside the store) directly. This allows us to capture the effect of environmental variables on consumer choice in a very general manner. When solving the model, we first consider the direct choice of Θ^k under different technological restrictions and then consider an extended model where context Θ^k is a function of the product line M^k, nesting the models of Salience (Bordalo, Gennaioli and Shleifer, 2013), Focusing (Kőszegi and Szeidl, 2013) and Relative Thinking (Bushong, Rabin and Schwartzstein, 2016).

Throughout the paper, firms simultaneously commit to a (finite) menu of products M^k (with $(q_j, p_j) \in \mathbb{R}^2$) and a context Θ^k before consumers move. All variables of a firm (the number of products, product qualities, prices, and distortions) are set simultaneously. We assume that firms—e.g., through historical observations of consumer choice—have perfect knowledge of the context-sensitivity parameter β and of the distribution of beliefs $F(\tilde{\beta})$, but cannot observe the type of individual consumers. Firms have symmetric cost functions. When a consumer purchases a good of quality q from firm k, the firm incurs a cost $c(q)$ that we assume is strictly convex increasing in the quality delivered, $c'(q) > 0$, $c''(q) > 0$, and satisfies $c(0) = c'(0) = 0$. These standard Inada conditions imply that for a given, context-dependent surplus function (see Equation (2)) there exists a unique, strictly positive quality q^c that is cost-efficient. In particular,

$$q^c = \begin{cases} q^* \ := \arg\max_q [q - c(q)] \ \ \Leftrightarrow c'(q^c) = 1 & \text{if } \theta_j^k = N, \\ q^Q := \arg\max_q [\beta q - c(q)] \Leftrightarrow c'(q^c) = \beta & \text{if } \theta_j^k = Q, \\ q^P := \arg\max_q [q - \beta c(q)] \Leftrightarrow c'(q^c) = \dfrac{1}{\beta} & \text{if } \theta_j^k = P. \end{cases}$$

Note that $q^Q > q^* > q^P > 0$. We concentrate on interior results by demanding that minimum quality $\underline{q}$ is sufficiently low and maximum willingness to pay b sufficiently high that consumers do not *per-se* reject buying cost-efficient quality q^c at cost. This is true with any context effect $\theta_j^k \in \{N, Q, P\}$ if and only if $\underline{q} \leq q^P$ and $b \geq c(q^Q)$, which we assume henceforth. To help us define equilibria that pin down in-store context and the size of product lines exaxtly, we assume that any distortion of context (choosing a store context other than $\theta_j^k = N \ \forall j \in J^k$) entails a positive but infinitely small cost as does the inclusion of one additional product in the product line.[11] This assumption sustains the results of an analysis without set-up costs but requires that firms distort preferences or add products only if doing so has a strictly positive effect on profits.

Solution Concept. We analyze market supply in the competitive equilibrium, where the latter is defined as as a tuple $(\boldsymbol{M}, \boldsymbol{\Theta})$, $\boldsymbol{M} := (M^k)_{k=1,...,K}$, $\boldsymbol{\Theta} := (\Theta^k)_{k=1,...,K}$, with the following properties:

1. (Nash Equilibrium) Firms play mutual best responses. For every $k \in \{1, ..., K\}$, $\pi^k((M^k, \Theta^k), (\boldsymbol{M}^{-k}, \boldsymbol{\Theta}^{-k})) \geq \pi^k((M^{k'}, \Theta^{k'}), (\boldsymbol{M}^{-k}, \boldsymbol{\Theta}^{-k}) \ \forall (M^{k'}, \Theta^{k'}) \neq (M^k, \Theta^k)$.

2. (Competitive Market) For every $k \in \{1, ..., K\}$, $\pi^k((M^k, \Theta^k), (\boldsymbol{M}^{-k}, \boldsymbol{\Theta}^{-k})) = 0$.

To resolve possible tie breaks, we make two assumptions. First, whenever indifferent, a consumer chooses each surplus maximizing option with positive probability. Second, there exists a smallest monetary unit $\delta > 0$, which we take to be positive but infinitesimally small.[12] This is equivalent to assuming that a firm, when best-responding, can resolve tie breaks in favor of the strictly more profitable product. We will exploit this equivalence when solving the model.

[11] Note, importantly, that due to the simultaneous choice of prices with other strategic variables any positive set-up costs will be covered by the sale price. The assumption of positive set-up costs will therefore *not* lead all firms to exit the market in the competitive equilibrium. An alternative assumption that yields the same results is that firms always choose the smallest profit-maximizing product line.

[12] Formally, let $\delta = \frac{1}{10^z}$ where $z \in \mathbb{Z}$ is an integer. Firms then choose qualities and prices from a discretized set of real numbers $R_z = \{r \in \mathbb{R} | (r \cdot 10^z) \in \mathbb{Z}\}$. In the limit $z \to \infty$ (i.e., $\delta \to 0^+$) this set is equal to $\mathbb{R}$.

3 Setting the Stage: Rational Benchmark and the Concept of Fooling

Rational Benchmark. When consumers are not sensitive to store context, our set-up yields a standard Bertrand outcome:

Lemma 1 (Rational Benchmark)**.** *Assume that consumers are* not *context-sensitive ($\beta = 1$). Then, in competitive equilibrium, at least two firms share the market. Each of these firms offer a single product with quality q^* priced at cost, $p^* = c(q^*)$, and do not engage in contextual distortion, $\Theta^k = (N)$. All other firms choose $(M^k, \Theta^k) = \emptyset$.*

When $\beta = 1$, preferences are stable and for both stages (outside and inside stores) uniquely defined by Equation (1). Neither the two-step choice of consumers nor potential naïveté is relevant in such a case because every consumer perfectly predicts her behavior in stage 2: The choice of a store is equivalent with the choice of a final product. A market so defined generates standard Bertrand incentives: A firm offering the highest (undistorted) surplus in the market wins all consumers. As a result, firms compete by marginally undercutting each other's sale price of a product with cost-efficient quality q^*. It follows that in equilibrium, at least one firm must offer a product with quality q^* at cost $c(q^*)$. The simultaneous choice of prices with other strategic variables implies that the sale price covers the positive set-up cost associated with offering the product (which is assumed infinitesimally small in our case and therefore does not show up explicitly), and thus, there is *no* incentive to exit the market in order to save this cost. At the same time, offering *more* than one product raises set-up costs without increasing profits and is also not part of a best response. The same reasoning implies that firms will not distort context: such a strategy would increase cost without affecting the preferences of rational consumers. While a second firm supplying q^* at cost $c(q^*)$ is necessary and sufficient to not admit deviations to higher profits, the definition of a competitive equilibrium admits any additional firms to not supply the market by choosing $(M^k, \Theta^k) = \emptyset$, thereby also yielding zero profits (at zero sales and zero cost).

Attraction and Fooling. Things change when $\beta > 1$ such that preferences are sensitive to the context in which products are presented at the point of purchase. Having consumers first select a store and then a product may now have important consequences for market supply. To see this, note that all consumers are attracted to a store by the product they *expect* to purchase. If a consumer is naïve regarding future preference changes, this product must not necessarily conform to the product the consumer will ultimately purchase. We therefore define:

Definition 1 (Attraction Product)**.** *We call product $j \in J^k$ the attraction product $a^k(\tilde{\beta})$ of firm k (for type $\tilde{\beta}$) if and only if a context-sensitive consumer with point-belief $\tilde{\beta}$ expects to purchase product j when entering store k.*

Definition 2 (Target)**.** *We call product $j \in J^k$ the target t^k of firm k if and only if a context-sensitive consumer ($\beta > 1$) who enters store k purchases product j.*

Of course, sophisticated context-sensitive consumers perfectly foresee their behavior at the store implying that for these consumers the firm's target is also the attraction product. In particular, these consumers enter store k if and only if target t^k is feasible ($p_{t^k} \leq b$ and $q_{t^k} \geq \underline{q}$) and provides at least as high (undistorted) surplus as any other firm's target. If a consumer is naïve, however, she may mispredict her choice at a store where preferences are distorted by local context. In this case, the consumer might be attracted to a store by a product that is not the target. If a firm designs a store that attracts type $\tilde{\beta} \neq \beta$ with a product that is not the target, we say that the firm *fools* the consumer:

Definition 3 (Fooling)**.** *Firm k* fools *type $\tilde{\beta}$ if and only if $a^k(\tilde{\beta}) \neq t^k$. If firm k fools type $\tilde{\beta}$,*

$$\text{(IC)} \qquad \hat{u}^k_{t^k} \geq \hat{u}^k_{a^k(\tilde{\beta})}$$

$$\text{(PCC)} \qquad E_{\tilde{\beta}}\left[\hat{u}^k_{t^k}\right] \leq E_{\tilde{\beta}}\left[\hat{u}^k_{a^k(\tilde{\beta})}\right]$$

with at least one of the inequalities being strict.

In this definition, condition (IC) is a standard incentive compatibility constraint: Inside store k, the consumer weakly prefers the target over the attraction product. Condition (PCC), on the other hand, is what we would call a *perceived choice* constraint: When entering store k, a consumer with expectation $E(\beta) = \tilde{\beta} \neq \beta$ (falsely) expects to weakly prefer the attraction product over the target. As will become clear over the course of our analysis, fooling is the sole function of in-store context in our framework. In other words, if a firm does not fool, that is, $a^k(\tilde{\beta}) = t^k$ for all $\tilde{\beta} \in supp[f(\tilde{\beta})]$, then the firm cannot improve by distorting preferences.

Profitable Fooling. It is clear that any naïve consumer can be fooled by a suitable choice of product line M^k and preference distortion Θ^k. However, the question arises under what circumstances fooling is profitable for a firm. This question is addressed by Lemma 2 below.

Lemma 2 (Profitable Fooling)**.** *Let $\beta > 1$. Assume w.l.o.g. that prices are unbounded, $b \rightarrow \infty$. Fix any target quality $q_{t^k} \geq \underline{q}$. If firm k does not fool, the maximum price at which* *the firm can sell quality q_{t^k} to consumers of type $\tilde{\beta}^0$ is*

$$p^0_{t^k} := q_{t^k} - \bar{u}(\tilde{\beta}^0),$$

where $\bar{u}(\tilde{\beta}^0) \geq 0$ is the highest undistorted surplus that a consumer with belief $\tilde{\beta}^0$ expects to receive when not *shopping at store k. Compare this to a fooling strategy where firm k attracts type $\tilde{\beta}^0$ with a different product $a^k \neq t^k$. Then it is true that:*

a) *If type $\tilde{\beta}^0$ over-estimates her sensitivity to context ($\tilde{\beta}^0 > \beta$), fooling her is* unprofitable*: Conditional on selling target $t^k \neq a^k$, the firm must charge a price p_{t^k} that is strictly lower than the price $p^0_{t^k}$ it can charge without fooling.*

b) *If type $\tilde{\beta}^0$ under-estimates her sensitivity to context ($\tilde{\beta}^0 < \beta$), fooling her is* profitable*:*

- *If in-store context inflates the qualities of the target and the attraction product, $(\theta_{a^k}, \theta_{t^k}) = (Q, Q)$* and *the target has a higher price and quality than the attraction product, $p_{t^k} > p_{a^k}$ and $q_{t^k} > q_{a^k}$ (i.e., the firm up-sells), or*
- *If in-store context inflates the prices of the target and the attraction product, $(\theta_{a^k}, \theta_{t^k}) = (P, P)$* and *the target has a lower price and quality than the attraction product, $p_{t^k} < p_{a^k}$ and $q_{t^k} < q_{a^k}$ (i.e., the firm down-sells), or*
- *If in-store context asymmetrically distorts surplus in favor of the target, $(\theta_a, \theta_t) \in \{(P, Q), (P, N), (N, Q)\}$ (allowing the firm to do both, up-* and *down-sell),*

the firm can sell target $t^k \neq a^k$ at a price p_{t^k} that is strictly higher than the price $p^0_{t^k}$ it can charge without fooling.

Moreover, fooling with other distortions is unprofitable. In particular, if in-store context asymmetrically distorts surplus in favor of the attraction product, $(\theta_{a^k}, \theta_{t^k}) \in \{(Q, P), (Q, N), (N, P)\}$, the consumer cannot be fooled to purchase target t^k at any $p_{t^k} \geq 0$.

It follows from part a) of Lemma 2 that a standard Bertrand strategy (without fooling the consumer) is more profitable than any fooling strategy when the firm sells to consumers who overestimate their sensitivity to context. Conversely, part b) establishes that fooling can yield *higher* profits than a Bertrand strategy when selling to consumers who are unaware of or under-estimate this sensitivity. Part b) also shows how particular forms of context-induced preference manipulation (shifting consumer perception of the quality and price of the firm's target and attraction product) relate to up-selling ($q_{t^k} > q_{a^k}$) or down-selling ($q_{t^k} < q_{a^k}$) strategies.

Before we explore these strategies in more detail in the next section, we end this one by characterizing the equilibrium for populations that consist entirely of consumers who do not lend themselves to (profitable) fooling:

Proposition 1 (Sophisticated and over-estimating consumers obtain the rational outcome)**.** *If consumers are context-sensitive ($\beta > 1$), but all of them (weakly) over-estimate their sensitivity to context ($\tilde{\beta} \geq \beta$ for all consumers), market supply is identical to the rational benchmark.*

4 Exploiting Naïve Context-Sensitivity: The Fooling Equilibrium and its Variations

We start by presenting a central yet—to some extent—auxiliary result of our paper below. Proposition 2 characterizes the equilibrium under the assumptions that *all* consumers under-estimate their context-sensitivity and that firms have an *unspecified* technology at hand that lets them choose the context of their store Θ^k. The result is central because—as we will show in later propositions—its major take-aways are generalizable: They are robust first, to making in-store context endogenous using the theories of Salience (Bordalo, Gennaioli and Shleifer, 2013), Focusing (Kőszegi and Szeidl, 2013) or Relative Thinking (Bushong, Rabin and Schwartzstein, 2016), and second, to the presence of rational and sophisticated consumers in the population. It is auxiliary because it helps us delineate the question of *how to* construct a specific context (using Salience, Focusing or Relative Thinking) from the strategic choice of *which context* to use.

Proposition 2 (Fooling Equilibrium)**.** *Assume that consumers are context-sensitive ($\beta > 1$) and all of them (strictly) under-estimate their sensitivity to context ($\tilde{\beta} < \beta$ for all consumers). Assume also that firms can choose Θ^k directly, either being restricted to* store-wide distortions*—for any two products j, j' at a given firm k, $\theta^k_j = \theta^k_{j'} = \theta^k$ and firms can choose $\theta^k \in \{Q, P, N\}$—or being able to choose* product-specific distortions*—firms can choose $\theta^k_j \in \{Q, P, N\}$ for each product $j \in J^k$ individually. Then, in competitive equilibrium, at least two firms share the market and each of these firms offers two products, t^k and $a^k \neq t^k$. Consumers are attracted to a firm by product a^k, which is offered below cost, $p_{a^k} < c(q_{a^k})$, but ultimately purchase product t^k which is priced at cost, $p_{t^k} = c(q_{t^k})$. All other firms choose $(M^k, \Theta^k) = \emptyset$. Moreover:*

a) **Store-Wide Distortions.** *Assume that for any two products j, j' at a given firm k, $\theta^k_j = \theta^k_{j'} = \theta^k$ and firms can choose $\theta^k \in \{Q, P, N\}$. Then firms with strictly positive* *demand choose either*

$$\theta^k = Q,\ q_{a^k} = \underline{q},\ q_{t^k} = q^Q > q^*,\ \textit{and up-sell } (q_{a^k} < q_{t^k}),$$
$$\textit{or}$$
$$\theta^k = P,\ p_{a^k} = b,\ q_{t^k} = q^P < q^*,\ \textit{and down-sell } (q_{a^k} > q_{t^k}).$$

Define

$$\nu^{(Q,Q)} := [q^Q - c(q^Q)] + (\beta - 1)(q^Q - \underline{q}),\ \textit{and}$$
$$\nu^{(P,P)} := [q^P - c(q^P)] + (\beta - 1)[b - c(q^P)].$$

Firms choose $\theta^k = Q$ and up-sell ($q_{a^k} < q_{t^k}$) if $\nu^{(Q,Q)} \geq \nu^{(P,P)}$, and choose $\theta^k = P$ and down-sell ($q_{a^k} > q_{t^k}$) if $\nu^{(Q,Q)} \leq \nu^{(P,P)}$.

b) **Product-Specific Distortions.** *Assume that firms can choose $\theta^k_j \in \{Q, P, N\}$ for each product $j \in J^k$ individually. Then firms with strictly positive demand choose*

$$(\theta^k_{a^k}, \theta^k_{t^k}) = (P, Q),\ p_{a^k} = b,\ q_{t^k} = q^Q > q^*,\ \textit{and down-sell } (q_{a^k} > q_{t^k}).$$

To develop an intuition, note first that Lemma 2 implies that the equilibrium must involve fooling: If consumers under-estimate context effects ($\tilde{\beta} < \beta$), fooling is more profitable than a classical Betrand undercutting strategy. When firms are restricted to store-wide distortions, they can choose $\theta^k = Q$ or $\theta^k = P$, leading to both target and attraction product being either overrated relative to the outside valuation with regard to quality, $(\theta_{a^k}, \theta_{t^k}) = (Q, Q)$, or with regard to price, $(\theta_{a^k}, \theta_{t^k}) = (P, P)$. According to Lemma 2, both types of context are more profitable than choosing a neutral frame $\theta^k = N$. When firms are able to choose product-specific distortions, it can be shown that inflating the quality of the target while at the same time inflating the price of the attraction product, $(\theta_{a^k}, \theta_{t^k}) = (P, Q)$, dominates all other choices. With either store-wide or product-specific distortions, selling quality q^* (as in the rational benchmark) is then no longer cost-efficient. Instead, firms sell a target t^k that caters to the consumers' (distorted) in-store preferences: $q^Q := \arg\max_q[\beta q - c(q)]$ and $q^P := \arg\max_q[q - \beta c(q)]$ define the cost-efficient qualities when selling a quality- or price-inflated target, respectively, leading to quality being either over- or under-provided relative to the rational benchmark.

To understand the rest of the equilibrium we need to concern ourselves with a firm's optimal choice of an attraction product a^k. Note first that a firm can attract *all* consumers who under-estimate their sensitivity to context ($\tilde{\beta} < \beta$) at maximum profit by using just *one* attraction product $a^k \neq t^k$: Given two products a^k and t^k, a firm maximizes the margin on the target by guaranteeing that the two products have identical perceived surplus at the store.[13] At $\hat{u}^k_{a^k} = \hat{u}^k_{t^k}$, however, *any* consumer who holds belief $\tilde{\beta} < \beta$ (falsely) expects to prefer product a^k over t^k at the store: $E_{\tilde{\beta}}\left[\hat{u}^k_{a^k}\right] > E_{\tilde{\beta}}\left[\hat{u}^k_{t^k}\right]$ for any $\tilde{\beta} < \beta$. This observation has two important consequences. First, because a firm is penalized for holding products that do not positively affect the profit margin, any firm that supplies the market holds exactly two products: the target t^k and just *one* attraction product a^k.[14] Second, because the best response does not generate heterogenous expectations among consumers with $\tilde{\beta} < \beta$, competition is Bertrand-like despite the extra fooling profits. As in the rational benchmark, the Bertrand undercutting dynamic comes to a stop only when the selling price of the target hits production cost $c(q_{t^k})$. The harsh competition leads to negative mark-ups on the attraction product, $p_{a^k} < c(q_{a^k})$, in equilibrium.

Rewriting $\hat{u}^k_a = \hat{u}^k_t$ as $u_{a^k} = u_{t^k} + [(\hat{u}^k_{t^k} - u_{t^k}) - (\hat{u}^k_{a^k} - u_{a^k})]$, or, equivalently, as

$$(4)\quad u_{a^k} \equiv \nu^{(\theta_{a^k},\theta_{t^k})} := \underbrace{u_{t^k}}_{\text{undistorted surplus of target}} + \underbrace{(\beta - 1)\begin{cases}(q_{t^k} - q_{a^k}) & \text{if } (\theta_{a^k},\theta_{t^k}) = (Q,Q),\\ (p_{a^k} - p_{t^k}) & \text{if } (\theta_{a^k},\theta_{t^k}) = (P,P),\\ (q_{t^k} + p_{a^k}) & \text{if } (\theta_{a^k},\theta_{t^k}) = (P,Q).\end{cases}}_{\text{context-dependent (virtual) surplus from switching to target}}$$

brings us closer to the heart of choosing an optimal attraction product with each context. Equation (4) re-interprets the surplus with which the consumer is attracted to the store ($= u_a$) as the *virtual* surplus of the target ($\equiv \nu^{(\theta_{a^k},\theta_{t^k})}$): When purchasing at firm k, a fooled consumer behaves *as if* she was maximizing the undistorted surplus u_{t^k} of the target plus a term that measures the context-dependent "sensation" of changing her mind when switching from product a^k to product t^k after entering the store. If context inflates qualities, $(\theta_{a^k},\theta_{t^k}) = (Q,Q)$, this sensation comes from switching to a product of higher than expected quality. If prices are inflated, $(\theta_{a^k},\theta_{t^k}) = (P,P)$, the consumer perceives a bargain effect when switching from the expensive product to the cheaper option. Finally, if the distortion is asymmetric, $(\theta_{a,^k}\theta_{t^k}) = (P,Q)$, the effect is a combination of a bargain effect from *not* purchasing the expensive product a^k and a value effect from purchasing a target with higher than expected quality. For a firm that fools, maximizing the sensation of switching products (i.e., the second summand in Equation (4)) means maximizing the effective margin on a

[13]If this was not the case, the firm could increase the target's price p_{t^k} or decrease its quality q_{t^k} (and thus, production cost $c(q_{t^k})$) and thereby increase profits.

[14]Note that the simultaneous choice of prices with other strategic variables implies that firms will cover the necessary set-up cost for these two products with the final sale price. For a firm that supplies the market, set-up costs thus do *not* provide a deviation incentive to exit the market. Because set-up costs are assumed infinitesimally small, they however do not show up *explicitly* in the price of the target.

given target and therefore drives the firm's choice of an optimal attraction product. Different contextual distortions therefore call for different marketing strategies: If a firm makes profit by up-selling the consumer, $(\theta_{a^k}, \theta_{t^k}) = (Q, Q)$, it achieves the strongest effect by attracting consumers with the lowest quality possible, $q_{a^k} = \underline{q}$, while competing with other firms on price p_{a^k}. If a firm down-sells, $(\theta_{a^k}, \theta_{t^k}) \in \{(P, P), (P, Q)\}$, the perception of making a bargain when switching products is maximized by fixing p_{a^k} at the maximum acceptable price b and competing with other firms on quality q_{a^k}.

Which context Θ^k firms finally choose in equilibrium depends on the firm's technological abilities as well as the profitability of feasible distortions. Condition (4) clearly shows that the asymmetric distortion $(\theta_{a^k}, \theta_{t^k}) = (P, Q)$ weakly dominates the symmetric distortions in profit.[15] When firms have access to a technology that allows for product-specific choice of θ_j^k, firms are best-off by constructing $(\theta_{a^k}, \theta_{t^k}) = (P, Q)$, which implies an *over*-provision of quality combined with a *down*-selling strategy in equilibrium. When firms are restricted to store-wide distortions, it depends on parameter values whether quality will be over- or under-provided and whether an up- or down-selling strategy prevails in equilibrium. Most intriguingly, the lower and upper bounds on acceptable quality and price, respectively, play an important role for this trade-off : Up-selling consumers by choosing $\theta^k = Q$ entails a high switching sensation and thus, high fooling profit when consumers can be attracted to the store with a product of rather low quality, that is, when $\underline{q}$ is low. Quality will be *over*-provided in this case. The same is true concerning the possibility to attract consumers with a product of relatively high price $p_{a^k} = b$ when firms choose $\theta^k = P$ and down-sell naïve consumers. Quality will be *under*-provided in this case. Limiting to the case of store-wide distortions, we would therefore expect to see more up-selling attempts in markets where consumers have a limited budget (b low) and can be attracted with low-quality products ($\underline{q}$ low), while down-selling attempts are more likely to be found in markets where consumers demand high quality ($\underline{q}$ high) and do not shy away from high prices (b high).

4.1 Salience, Focusing, and Relative Thinking

To make more specific predictions on *how* firms manipulate consumer preferences at their store, we now introduce the three theories of contextual distortion that have recently been provided by Bordalo, Gennaioli and Shleifer (2013), Kőszegi and Szeidl (2013), and Bushong, Rabin and Schwartzstein (2016) to our framework. These models make in-store context a function of the product line—assuming that consumers overweight attributes to

[15] Not shown in Expression 4 are other asymmetric distortions which are also dominated by $(\theta_{a^k}, \theta_{t^k}) = (P, Q)$. If $(\theta_{a^k}, \theta_{t^k}) = (N, Q)$, the context-dependent fooling profit is $(\beta - 1)q_t$, while for $(\theta_{a^k}, \theta_{t^k}) = (P, N)$, it is $(\beta - 1)p_a$.

which *relative comparisons* between products draw their attention.

Bordalo, Gennaioli and Shleifer (2013, henceforth BGS) assume that consumers attach disproportionally high weight to *salient* attributes, where "[a]n attribute is salient for a good when it stands out among the good's attribute relative to that attribute's average level in the choice set" (BGS, cited from the abstract, p. 803). We apply the original salience definition by BGS (BGS, Definition 1 and Assumption 1) to a choice set equal to the product line of store k:

Assumption S (Salience)**.** *Let z_R^k be the* average *level of attribute $z \in \{q, p\}$ at store k. The salience of attribute z_j, $z \in \{s, p\}$ at store k is given by a symmetric and continuous (real-valued) function $\sigma\left(z_j, z_R^k\right)$ that satisfies ordering and homogeneity of degree zero.*[16] *Then*

$$\theta_j^k = \begin{cases} Q & \textit{if and only if } \sigma\left(q_j, q_R^k\right) > \sigma\left(p_j, p_R^k\right) \\ P & \textit{if and only if } \sigma\left(q_j, q_R^k\right) < \sigma\left(p_j, p_R^k\right) \\ N & \textit{otherwise.} \end{cases}$$

Kőszegi and Szeidl (2013, henceforth KS) argue "that a person focuses more on, and hence overweights, attributes in which her options differ more" (KS, cited from the abstract, p. 53). We implement the central assumption of KS (Assumption 1) in the following way:

Assumption F (Focusing)**.** *Let Δ_z^k be the spread of attribute $z \in \{q, p\}$ at store k, $\Delta_z^k := \max_{j \in J^k} z_j - \min_{j \in J^k} z_j$, and let $\kappa_F \geq 0$ be some (exogenously defined) threshold. Then*

$$\theta_j^k = \begin{cases} Q & \textit{if and only if } \Delta_q^k - \Delta_p^k > \kappa_F \\ P & \textit{if and only if } \Delta_p^k - \Delta_q^k > \kappa_F \\ N & \textit{otherwise.} \end{cases}$$

Note that for any two products $\{j, i\} \in J^k$, $\theta_j^k = \theta_i^k = \theta^k$, i.e., distortions are store-wide.

Bushong, Rabin and Schwartzstein (2016, henceforth BRS) model the idea that "[f]ixed differences loom smaller when compared to large differences" (BRS, cited from the abstract, p.1). The consumer thus "weighs a given change along a consumption dimension by less when it is compared to bigger changes along that dimension" (ibid.). We base our implementation on the central norming assumptions N0-N2 in BRS.

[16] These are defined as follows. (1) Ordering: Let $\mu = sgn\left(z_k - z_R^k\right)$. Then, for any $\varepsilon, \varepsilon' \geq 0$ with $\varepsilon + \varepsilon' > 0$, $\sigma\left(z_j + \mu\varepsilon, z_R^k - \mu\varepsilon'\right) > \sigma\left(z_j, z_R^k\right)$. (2) Homogeneity of degree zero: $\sigma\left(\alpha z_j, \alpha z_R^k\right) = \sigma\left(z_j, z_R^k\right) \forall \alpha > 0$. These definitions are valid for $z_j > 0$ and $z_R^k > 0$. In order to work with nonpositive arguments, we would need to formulate additional properties, see BGS. For our results it is however sufficient to have salience defined in the positive domain.

Assumption RT (Relative Thinking). *Let Δ_z^k be the spread of attribute $z \in \{q, p\}$ at store k, $\Delta_z^k := \max_{j \in J^k} z_j - \min_{j \in J^k} z_j$, and let $\kappa_{RT} \geq \beta$ be some (exogenously defined) threshold. Then*

$$\theta_j^k = \begin{cases} Q & \text{if and only if } \dfrac{\Delta_p^k}{\Delta_q^k} > \kappa_{RT} \\ P & \text{if and only if } \dfrac{\Delta_q^k}{\Delta_p^k} > \kappa_{RT} \\ N & \text{otherwise.} \end{cases}$$

Note that for any two products $\{j, i\} \in J^k$, $\theta_j^k = \theta_i^k = \theta^k$, i.e., distortions are store-wide.[17]

When in-store context is defined according to one of the three theories, we can make the following observation.

Proposition 3 (Fooling with Decoys: Salience, Focusing, and Relative Thinking). *Assume that consumers are context-sensitive ($\beta > 1$) and all of them under-estimate their sensitivity to context ($\tilde{\beta} < \beta$ for all consumers). Let in-store context Θ^k be defined according to Salience (Assumption S), Focusing (Assumption F), or Relative Thinking (Assumption RT). Then, in competitive equilibrium, at least two firms share the market, each firm offering three products, t^k, $a^k \neq t$, and $d^k \notin \{t^k, a^k\}$. Firms sell product t^k at cost, $p_{t^k} = c(q_{t^k})$, but attract consumers with product a^k, offered at a price below cost, $p_{a^k} < c(q_{a^k})$. The sole function of product d^k is to manipulate preferences at the store: Product d^k is a* decoy. *Moreover:*

a) Under Assumption F or RT, preference distortions are store-wide: for any two products j, j' at a given firm k, $\theta_j^k = \theta_{j'}^k = \theta^k$. Firms implement $\theta^k \in \{Q, P\}$, (q_{t^k}, p_{t^k}), and (q_{a^k}, p_{a^k}) according to Proposition 2, Part a), using a single decoy d^k.

b) Under Assumption S, preference distortions may be product-specific. Firms implement $(\theta_{a^k}, \theta_{t^k}) = (P, Q)$, (q_{t^k}, p_{t^k}), and (q_{a^k}, p_{a^k}) according to Proposition 2, Part b), using a single decoy d^k.

[17] Assumption RT implements norming assumptions N0-N2 in the following way. Let $w(\cdot)$ denote the weight function that attaches weight $w_z^k \in \{1, \beta\}$ to attribute $z \in \{q, p\}$. N0 is simply the assumption that $w(\cdot)$ is a function of the attribute spread Δ_z^k. Now suppose that quality has a higher weight than price, i.e. $w_q^k = \beta$ and $w_p^k = 1$. According to our framework, $\theta_j^k = Q$ for all products $j \in J^k$. By N1, $w(\Delta_q^k) > w(\Delta_p^k) \Rightarrow \Delta_q^k < \Delta_p^k$. But N2 makes a more restrictive assumption, namely, $w(\Delta_q^k) > w(\Delta_p^k) \wedge \Delta_q^k < \Delta_p^k \Rightarrow w(\Delta_q^k)\Delta_s^k < w(\Delta_p^k)\Delta_p^k \Leftrightarrow \beta\Delta_q^k < \Delta_p^k$, which is identical to our implementation by Assumption RT if $\kappa_{RT} = \beta$. The possibility of $\kappa_{RT} > \beta$ captures cases where stronger stimulus is required. An analogous statement establishes the case of $\theta_j^k = P$.

The important take away from Proposition 3 is that the profit-maximizing distortions we have defined earlier (see Proposition 2) can be realized by adding just *one* additional product to the product line. As we illustrate in Figures 1 and 2 below, the location of this product in quality-price space depends on which specification is employed.

Figure 1 shows decoy-positions under Assumptions F and RT in equilibrium. Under these assumptions, firms are restricted to store-wide distortions. There are two cases: (1) the firm *up-sells*, $q_{t^k} > q_{a^k}$ and $p_{t^k} > p_{a^k}$, by inflating perceived qualities: $(\theta_{a^k}, \theta_{t^k}) = (Q, Q)$ (left panel), or (2) the firm *down-sells*, $q_{t^k} < q_{a^k}$ and $p_{t^k} < p_{a^k}$, by inflating perceived prices: $(\theta_{a^k}, \theta_{t^k}) = (P, P)$ (right panel). To achieve the profit-maximizing distortion without violating incentive compatibility, the firm has to add a decoy to the product line that resides within the boundaries of the grey shaded areas in Figure 1.[18] Note that the shaded areas for Assumption F and RT do not overlap, implying that decoys can help identifying the two models from data. Under both theories, a choice that resonates with experimental literature and anecdotal evidence on so-called *decoy effects* is to construct a decoy that copies the target in one attribute but is strictly worse along the other dimension.[19] The white markers in Figure 1 illustrate such a choice.

When attention is modeled according to BGS' model of Salience (Assumption S), contextual distortions of quality and price may be product-specific and choosing distortion $(\theta_{a^k}, \theta_{t^k}) = (P, Q)$ is profit-maximizing. Figure 2 illustrates how the firm can construct this distortion. The figure depicts the case when, as in equilibrium, $q_{a^k} > q_{t^k}$ and $p_{a^k} > p_{t^k}$ (the firm down-sells). The firm can implement the distortion $(\theta_{a^k}, \theta_{t^k}) = (P, Q)$ by constructing a reference point (q_R^k, p_R^k) that is either dominated by the target ($p_R^k = p_{t^k}$, but $q_R^k < q_{t^k}$) or by the attraction product ($q_R^k = q_{a^k}$, but $p_R^k > p_{a^k}$). Which of the two constructions is feasible depends on whether the target or the attraction product has a higher quality-to-price ratio (see the left panel and right panel of Figure 2, respectively). In both cases, such a reference point can always be constructed—using a single, unattractive decoy—without violating incentive compatibility.

4.2 Fooling with Mixed Populations

How is the predicted exploitation of naïve consumers affected by the co-existence of sophisticated or rational consumers? We show below that fooling survives in mixed populations. Firms react to context-sensitive yet more *sophisticated* consumers by providing

[18] The shaded areas show decoy positions for minimum thresholds $\kappa_F = 0$ and $\kappa_{RT} = \beta$. Larger thresholds demand decoys that are located further away from products t^k and a^k.

[19] See, e.g., Huber, Payne and Puto (1982); Doyle et al. (1999); Herne (1999) for experimental literature on asymmetrically dominated decoys.

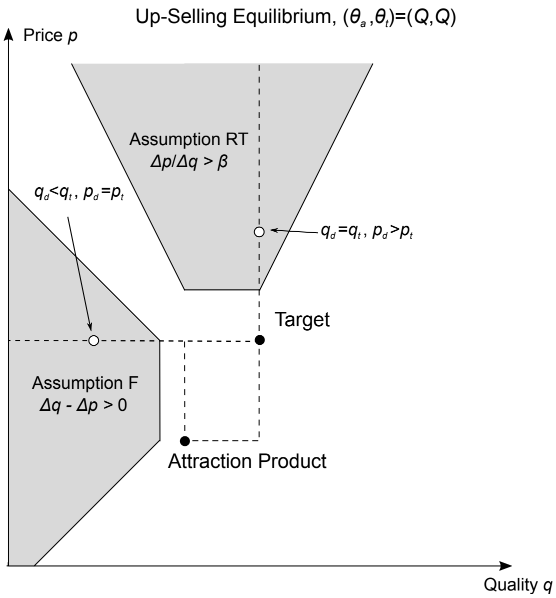

Figure 1: Decoy position (= within shaded areas) under Assumption F and Assumption RT

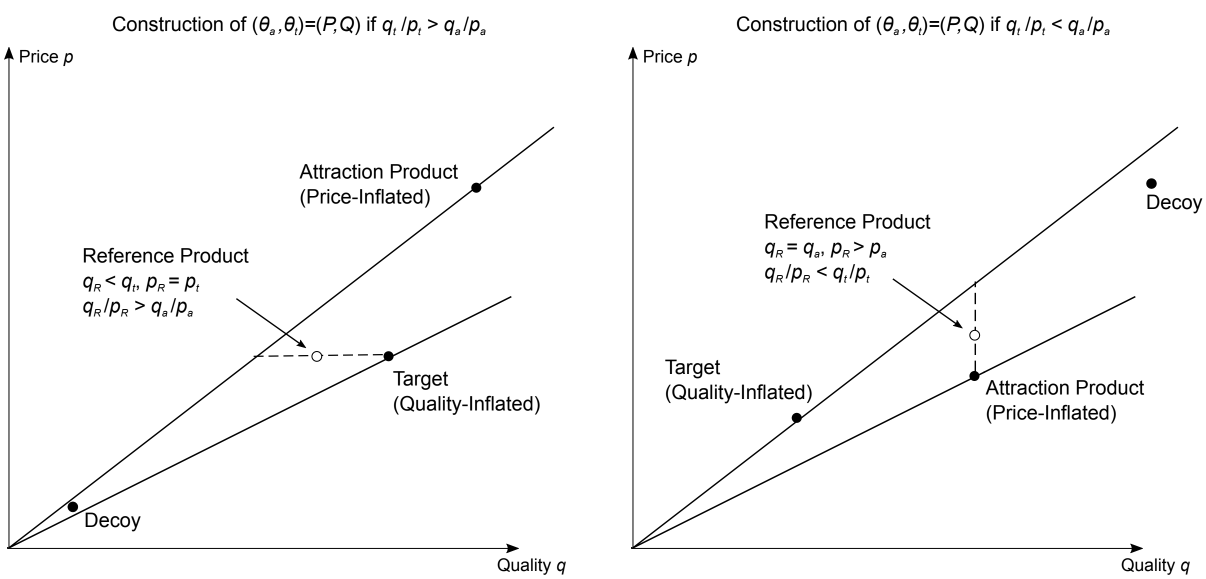

Figure 2: Construction of distortion $(\theta_a, \theta_t) = (P, Q)$ (with one decoy) under Assumption S. The construction exploits two central implications of the Salience framework: (1) If product $j \in J^k$ neither dominates nor is dominated by the reference point, i.e., $(q_j - q_R^k)(p_j - p_R^k) > 0$, then the "advantageous" attribute of product j—higher quality *or* lower price relative to the average—is overweighted if and only if the product has better-than-average quality-to-price ratio, that is, $(q_j/p_j) > (q_R^k/p_R^k)$. (2) If one attribute of product $j \in J^k$ is average while the other is not (e.g., $q_j = q_R^k$, but $p_j \neq p_R^k$), then the latter is overweighted.

additional, non-distortionary stores that allow consumers to self-commit to the ex-ante efficient product (mirroring market supply in the rational benchmark). Because all consumers who are aware of or over-estimate their bias sort into these additional stores, the market supply and exploitation of naïve under-estimating types is completely unaffected by the presence of more sophisticated types. The presence of *rational* consumers ($\beta = 1$), on the other hand, affects the "degree" to which firms can fool naïves: Having no commitment problem, rational consumers will always re-exploit a fooling firm by purchasing its (non-profitable) attraction product instead of the (profitable) target. In order to not loose too much money on rational consumers, firms will make the attraction product less of a bargain, moving it closer to the rational benchmark. However, the incentive to use context effects to up- or down-sell naïve consumers is *not* lessened. Fooling survives with the result being a trade-off between the profit lost on rational consumers ($p_{a^k} < c(q_{a^k})$) and the profit made on up- or downsold naïves ($p_{t^k} > c(q_{t^k})$). Because rational agents gain from the presence of naïves (the bargain of the former being subsidized by the latter), the exploitation of naïves even *increases* compared to the original fooling equilibrium.

For the following two propositions, let firms either directly choose θ_j^k (with store-wide or product-specific distortions, following the assumptions in Proposition 2), or let Assumption S, F, or RT be satisfied (firms can manipulate θ_j^k indirectly using decoy products).

Proposition 4 (Co-Existence of Sophisticated and Naïve Agents)**.** *Assume that all consumers are context sensitive ($\beta > 1$) with arbitrary distribution of naïveté $F(\tilde{\beta})$ in the population. In competitive equilibrium, product supply for (naïve) consumers who under-estimate their sensitivity to context ($\tilde{\beta} < \beta$) is unaffected by the existence of consumers who are sophisticated or over-estimate their sensitivity to context ($\tilde{\beta} \geq \beta$). In particular:*

a) If there exist consumers who are sophisticated or over-estimate their sensitivity to context ($\tilde{\beta} \geq \beta$), then there exist at least two non-fooling ('truthful') firms that each offer a single, undistorted product with quality q^* at marginal cost $c(q^*)$ and all consumers with $\tilde{\beta} \geq \beta$ buy at one of these firms.

b) If there exist consumers who under-estimate their sensitivity to context ($\tilde{\beta} < \beta$), then there exist at least two fooling firms that each offer two or three products, according to Propositions 2 and 3, respectively, and all consumers with $\tilde{\beta} < \beta$ buy at one of these firms.

Proposition 5 (Co-Existence of Rational and Naïve Agents)**.** *Assume that a share $\eta > 0$ of consumers are context-sensitive ($\beta > 1$) and naïve (with belief $\tilde{\beta} < \beta$), while the remaining* *share $(1-\eta) > 0$ of consumers are rational (β=1). Assume that b is sufficiently large to allow for interior solutions (w.l.o.g., let $b \to \infty$). In competitive equilibrium, at least 2 firms share the market, each choosing a distortionary in-store context Θ^k and fooling naïve consumers by attracting them with product a^k offered at price $p_{a^k} < c(q_{a^k})$, but up- or down-selling them to product $t^k \neq a^k$ at price $p_{t^k} > c(q_{t^k})$. The exploitation of naïves is larger than without the existence of rational consumers. Rational consumers re-exploit firms by purchasing the attraction product a^k: Total profits of each firm are zero. As $\eta \to 0$, product supply for rational consumers converges to the rational benchmark, $\lim_{\eta\to 0}(q_{a^k}, p_{a^k}) = (q^*, c(q^*))$, but the exploitation of naïves persists. In particular, with any equilibrium in-store context Θ^k, $\lim_{\eta\to 0} q_t^k \neq q^*$ and $\lim_{\eta\to 0} p_t^k > c(q_t^k)$. Moreover:*

a) **Store-Wide Distortions.** *Assume that for any two products j, j' at a given firm k, $\theta_j^k = \theta_{j'}^k = \theta^k$ and firms can choose $\theta^k \in \{Q, P, N\}$. Then firms with strictly positive choose either*

$$\theta^k = Q,\ q_{a^k} = \underline{q}_a := \max\{\underline{q}, q|_{c'(q)=1-\frac{\eta}{1-\eta}(\beta-1)}\} < q^*,\ q_{t^k} = q^Q > q^*, \text{and up-sell } (q_{a^k} < q_{t^k}),$$

or

$$\theta^k = P,\ q_{a^k} = \bar{q}_a := q|_{c'(q)=1+\frac{\eta}{1-\eta}\left(1-\frac{1}{\beta}\right)} > q^*,\ q_{t^k} = q^P < q^*, \quad \text{and down-sell } (q_{a^k} > q_{t^k}).$$

Define

$$\nu^{(Q,Q)} := \eta\left[[q^Q - c(q^Q)] + (\beta-1)\left(q^Q - \underline{q}_a\right)\right] + (1-\eta)\,[\underline{q}_a - c(\underline{q}_a)],\ \text{and}$$
$$\nu^{(P,P)} := \eta\left[[q^P - c(q^P)] + \left(1-\frac{1}{\beta}\right)\left(\bar{q}_a - q^P\right)\right] + (1-\eta)\,[\bar{q}_a - c(\bar{q}_a)].$$

Firms choose $\theta^k = Q$ and up-sell ($q_{a^k} < q_{t^k}$) if $\nu^{(Q,Q)} \geq \nu^{(P,P)}$ and choose $\theta^k = P$ and down-sell ($q_{a^k} > q_{t^k}$) if $\nu^{(Q,Q)} \leq \nu^{(P,P)}$.

b) **Product-Specific Distortions.** *Assume that firms can choose $\theta_j^k \in \{Q, P, N\}$ for each product $j \in J^k$ individually. Then firms with strictly positive demand choose*

$$(\theta_{a^k}^k, \theta_{t^k}^k) = (P, Q),\ q_{a^k} = q|_{c'(q)=1+\frac{\eta}{1-\eta}(\beta-1)} > q^*,\ q_{t^k} = q^Q > q^*.$$

Firms up-sell ($q_{a^k} \leq q_{t^k}$) if $\eta \leq \frac{1}{2}$ and down-sell ($q_{a^k} > q_{t^k}$) if $\eta > \frac{1}{2}$.

Proposition 5 highlights an important comparative static: For a rational consumer, the surplus from re-exploiting firms disappears as the share of naïves goes to zero: In the limit, as $\eta \to 0$, they are provided with the exact same product as in the rational benchmark. A similarly comforting conclusion can however not be drawn for the naïves: No matter how small their share η in the consumer population, the exploitation never goes to zero. In fact, while the quality they receive ($q_t \in \{q^Q, q^P\}$) is independent of η, they pay a price strictly above cost, $p_t > c(q_t)$, whenever their share is below unity. This finding shows that fooling may be an important, welfare-relevant phenomenon even when the mass of victims falling prey to such practices is small.

5 Conclusion

We conclude by discussing two modeling assumptions, namely (1) the assumption that consumers can only visit one store and (2) the assumption that firms pay an infinitesimally small set-up cost for each product, and by relating our results to earlier findings in the literature on market competition with biased consumers.

Discussion of modeling assumptions. The impossibility of consumers to visit multiple stores may seem too restrictive at first glance. For the qualitative results and conclusions of our paper, the consequences of this assumption are in fact very mild. To see this note first that—in comparison to standard models of consumer search—the consumer in our framework has *full* information regarding her choice set when making the entry decision in stage 1: Because firms commit to perfectly observable product lines ex-ante, there is no information to gain from visiting multiple stores. The commitment to a fixed, i.e., deterministic product line distances the fooling equilibrium also from extensively studied forms of "bait-and-switch" where firms limit the stock of the attraction product and then rely on positive switching cost to sell a profitable target to those customers who missed the limited "bait offer" (see, e.g., Lazear, 1995). As we will now argue, the exploitation we describe in this paper does *not* rely on switching cost. The assumption that consumers visit only one store for this matter does not conceal a possible store-switching incentive on the side of consumers. The first to note is that the full information set-up in our framework implies that the target must be a competitive offer in equilibrium. Because firms cannot withdraw the bait offer made to consumers ex-ante, competition is transferred into the store via the option to buy the attraction product. As in a model of direct product choice, the mark-up on the target is competed away in equilibrium. Clearly, sophisticated and rational consumers have no incentive to visit more than one store—knowing ex-ante that the choices available elsewhere do not increase their surplus. In order to study naïve consumers in a setting where switching stores is possible, one needs to define how these consumers value the product lines of other firms when preferences (unexpectedly) change due to being exposed to context Θ^k. Two possible assumptions come to mind. The first—in our view, the more natural interpretation of context-sensitivity—is that preferences reflect a general "state of mind" that applies to any options the consumer might consider when exposed to context Θ^k. In such a state of mind, options at other stores that are identical to those available at store k will be quality- or price-inflated in the exact same way as products at store k. For instance, in-store context might induce a "quality-salient" (or "price-salient") state of mind, making the consumer generally willing to spend more (or less) money on a given unit of quality—regardless of where the product is located. Fixing any equilibrium we have defined in this paper, a naïve consumer would then never want to visit a second store as she does not gain a product of higher surplus elsewhere. Another possible assumption—which we find less compelling—is that context Θ^k affects only the preferences over products at store k, leaving the valuation of all other products (even identical ones) unaffected. A naïve consumer might then not buy a price-inflated target ($\theta_{t^k} = P$), because she suddenly perceives the (undistorted) attraction products and targets at other stores as more valuable. If switching costs are not too high, she will want to visit more than one store. When a firm sells a *quality*-inflated target ($\theta_{t^k} = Q$), however, the result that consumers only visit one store (where they are fooled) is robust without imposing switching costs. Because quality-inflated targets are not restricted to up-selling equilibria, up-selling (with $\theta_{a^k} = \theta_{t^k} = Q$) and down-selling (with $(\theta_{a^k}, \theta_{t^k}) = (P, Q)$) predictions survive.[20]

We have assumed that there exists an infinitesimally small cost for setting-up a product. This implies that firms will not unnecessarily inflate the product line. One could argue that in reality, set-up costs are either zero (in online markets) or sizable (in bricks-and-mortar markets). When set-up costs are zero, all of our results go through except that firms are now indifferent between setting-up profit-maximizing product lines of minimal size (which are identical with the product lines we have defined) and larger product lines that include products that have zero marginal effect on profit. Consumer choice is unaffected. We think that even without explicit set-up costs, there are enough reasons for firms to not inflate the product line with options that do not affect consumer choice.[21] Of course, if set-up costs

[20] Of course, things become more complicated if we consider the possibility that the information of a preference change leads naïve consumers to learn something about their bias. This is an assumption that is rarely made in the literature, with Ali (2011) being a notable exception. Experiments show that people perform badly in updating beliefs about their own biases, leading us to conjecture that such effects are unlikely to make consumers fully rational. If consumers simply become more sophisticated without increasing the ability to control themselves, none of our results changes. If some consumers suddenly become rational, our results survive as long as a positive share of consumers remains naïve (see Proposition 5). A study of more involved updating procedures lies outside of the scope of this paper and is relegated to future research.

[21] Note that decoys and attraction products in this paper are not unnecessary products. These products have strictly positive marginal effect on profit by enabling the fooling outcome, even in the case where no consumer purchases these products. For this reason, the minimal size of profit-maximizing product lines in

are positive and sizable, fooling becomes more difficult to sustain. In this case, there will be a minimal degree of context-sensitivity necessary for firms to recover the additional set-up cost for the un-sold attraction product (and, potentially, a decoy) with the additional fooling profit made on naïve consumers. Note that positive set-up costs do not in general provide a strategic incentive to exit the market (even when profits are zero): Because the size of the product line is chosen simultaneously with other strategic variables such as qualities and prices, firms that supply the market will recover (positive but sufficiently low) set-up costs with the sale price.

Related theory and findings in behavioral I.O. We begin by laying out differences between the strategic use of context we describe in this paper and the role of "salience effects" for product-choice in models of "direct" competition as studied by Bordalo, Gennaioli and Shleifer (2016). With or without a second phase of consumer choice, context-sensitivity may lead firms to over- or under-provide quality (relative to the rational benchmark) in equilibrium. However, the forces driving this distortion are very different. In Bordalo, Gennaioli and Shleifer (2016), the decision to over- or under-provide quality is driven by salience effects *between* firms: Firms over-provide quality when competing on quality is more likely to draw consumer attention to the firm than a price-competition. In our paper, the decision is driven by context effects *within* firms: Firms over-provide quality when a positive shock to quality preferences at the store generates a stronger behavioral reaction than sudden shocks to price sensitivity. On the observable side, the most obvious difference in outcomes relates to product line choices: When consumers cannot switch products after they have selected a firm, offering pure attraction products as firms do in our paper cannot be part of a best response.[22] Because consumers in direct markets (as considered by Bordalo, Gennaioli and Shleifer, 2016) are therefore not "fooled" into thinking that they buy products other than they end up consuming and do not value products in a distorted manner when purchasing them, the additional welfare effects of the exploitation we describe are likely to be considerable.

There are other papers in behavioral I.O. that feature a two-phase choice procedure by which consumers first select a firm and then a product, but no study has so far considered the design of choice environments to be a source of preference distortions. Related to us by the idea that "marketing devices" play a role in attracting consumers to a firm is Eliaz and

the case of fooling is two (without decoys, Proposition 2) or three (with decoys, Proposition 3), respectively.

[22]Although Bordalo, Gennaioli and Shleifer (2016) do not study the possibility of firms to offer decoy products (in their model, each firm is restricted to offering a single product), they may in principle also play a role in models of direct competition. Whether decoys would be used in a similar way as in our model is an interesting question for future research.

Spiegler (2011*b*). The authors study the role of zero-utility products for attracting consumers to a firm with a larger product line. At first glance, these so-called "attention grabbers" seem to be very much related to what we call the attraction product of a firm. However, the relationship is only remote. In our case, (naïve) consumers attend to the attraction product because they (falsely) expect to consume it. In Eliaz and Spiegler (2011*b*), people follow attention grabbers for exogenous reasons such as sensationalism or similarity to familiar products. Because attention grabbers do not reflect consumer preferences, firms in Eliaz and Spiegler (2011*b*) use them to attract consumers toward products that are better suited to fulfill the consumer's underlying preference than the product she would be otherwise consuming. This is of course the exact reverse to what we predict, namely that the use of a separate attraction product is always associated with a firm that fools consumers into buying a product of *lesser* value.[23] The testable implication that arises from this difference is that according to Eliaz and Spiegler (2011*b*), no consumer would ever consume the attention grabber, while we predict that rational or self-controlled consumers (or, for that matter, any consumer in expectation) would always buy the attraction product. Note further that a decoy, which firms in our model may produce and no consumer is keen to purchase, is markedly different from the attention grabber as well. Decoys are unattractive at any stage of the decision process and therefore cannot be used to attract consumers to the firm. Moreover, a decoy is used by firms for its ability to affect the preference relation of more attractive options in the product line. A product that leaves preferences entirely unaffected—as attention grabbers in Eliaz and Spiegler (2011*b*) do—would never be used in our framework. Our paper compares similarly to Eliaz and Spiegler (2011*a*) and Piccione and Spiegler (2012), where—as in Eliaz and Spiegler (2011*b*) and in stark comparison to our paper—the distortive mechanism is relevant only in the first phase of consumer choice and operates over manipulating the consideration set rather than the preferences: At first glance, the two papers relate to ours by the idea that "frames" can influence consumer choice. At second glance, however, the mechanism of the bias is entirely different to our understanding and modeling of context effects. Similar to their use of attention grabbers in Eliaz and Spiegler (2011*b*), firms in Eliaz and Spiegler (2011*a*) and Piccione and Spiegler (2012) use exogenously defined "frames" to attract consumers away from status-quo products and toward products of higher value. This is the exact opposite to how firms use context in

[23] In Eliaz and Spiegler (2011*b*), the distortive mechanism operates over manipulating the consideration set rather than the preferences. This difference in approaches to consumer bias seems to be driving the prediction whether firms use a "psychology-based" strategic variable (a.k.a. "salience effects") to improve outcomes for the biased consumer (Eliaz and Spiegler 2011*b*, for similar results see also Eliaz and Spiegler 2011*a* and Piccione and Spiegler 2012) or to generate possibilities to exploit them (our paper, for similar results see also Gabaix and Laibson 2006, Heidhues and Kőszegi 2010 and many more). A more in-depth analysis of this, admittedly, very interesting comparison lies however outside of the scope of this paper.

our paper.

More related to the exploitation we describe are the results of models that—while at first glance being unrelated to the idea of context-sensitive choice—also combine a two-phase choice procedure with some form of "naïve" preference-distortion. These include studies of markets where firms sell a bundled product that consists of a base product and a costly, unavoidable add-on (e.g., Gabaix and Laibson, 2006; Ellison, 2005), the related "hidden price" literature (e.g., Heidhues, Kőszegi and Murooka, 2017), and the literature on contracting with time-inconsistent consumers (e.g., DellaVigna and Malmendier, 2004; Heidhues and Kőszegi, 2010). In all of these papers, naïve consumers mispredict their demand (or, equivalently, the prices) at a given firm k when selecting between different suppliers. Profit-maximizing firms exploit this naïveté by acting as "aftermarket monopolists" for those consumers who experience an unexpected change to their preferences. In these models, similar to our paper, (1) competition over consumers (in the first stage) does not solve the exploitation problem, (2) the co-existence of rational and profitable-to-exploit consumers increases the problem for the exploited instead of mitigating it,[24] and (3) bias-*over*estimating consumers, while also naïve, cannot be profitably exploited (see, for this particular point, Heidhues and Kőszegi, 2010). Our paper extends these findings to a new form of bias that predicts and explains the exploitation of naïve consumers in markets and circumstances that are not covered by the existing literature. Moreover, because our study of context effects allows time-inconsistency to be endogenously triggered and directed by firms instead of simply *assuming* a specific form of it, we provide an extended explanation of how such biases may be formed and exploited by firms. Doing so we find that firms may in fact find it optimal to increase price-sensitivity instead of reducing it (as has been implicitly assumed by the papers cited above), generating the novel prediction of down-selling phenomena. While our model focuses on product line effects, similar incentives to design the choice environment of consumers might hold for the markets studied in other papers. In contract environments, for example, whether consumers are more or less present-biased is likely to be affected by how the terms of a contract are presented. Exploiting naïve consumers by varying the presentation of contract terms over the consumption schedule would then be very close to the context-related fooling strategies that we have described in this paper. Studying this possibility in further detail is an interesting topic for future research.

[24] Armstrong (2015) has recently surveyed models that make this prediction, a characteristic he calls "ripoff externalities".

References

Ali, S. Nageeb. 2011. "Learning Self-Control." *Quarterly Journal of Economics*, 126: 857–893.

Ariely, Dan, George Loewenstein, and Drazen Prelec. 2003. "Coherent Arbitrariness: Stable Demand Curves Without Stable Preferences." *Quarterly Journal of Economics*, 118(1): 74–105.

Armstrong, Mark. 2015. "Search and Ripoff Externalities." *Review of Industrial Organization*, 47: 273–302.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2013. "Salience and Consumer Choice." *Journal of Political Economy*, 121(5): 803–843.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2016. "Competition for Attention." *Review of Economic Studies*, 83: 481–513.

Bushong, Benjamin, Matthew Rabin, and Joshua Schwartzstein. 2016. "A Model of Relative Thinking." mimeo (This version: March 31, 2016).

DellaVigna, Stefano, and Ulrike Malmendier. 2004. "Contract Design and Self-Control: Theory and Evidence." *The Quarterly Journal of Economics*, 119(2): 353–402.

Doyle, John R., David J. O'Connor, Gareth M. Reynolds, and Paul A. Bottomley. 1999. "The Robustness of the Asymmetrically Dominated Effect: Buying Frames, Phantom Alternatives, and In-Store Purchases." *Psychology and Marketing*, 16(3): 225–243.

Eliaz, Kfir, and Ran Spiegler. 2011*a*. "Consideration Sets and Competitive Marketing." *Review of Economic Studies*, 78: 235–262.

Eliaz, Kfir, and Ran Spiegler. 2011*b*. "On the Strategic Use of Attention Grabbers." *Theoretical Economics*, 6: 127–155.

Ellison, Glen. 2005. "A Model of Add-On Pricing." *Quarterly Journal of Economics*, 120(2): 585–637.

Ellison, Glen, and Sarah Fisher Ellison. 2009. "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica*, 77(2): 427–452.

Gabaix, Xavier, and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia and Information Suppression in Competitive Markets." *The Quarterly Journal of Economics*, 121(2): 505–540.

Heidhues, Paul, and Botond Kőszegi. 2008. "Competition and Price-Variation When Consumers are Loss Averse." *American Economic Review*, 98(4): 1245–1268.

Heidhues, Paul, and Botond Kőszegi. 2010. "Exploiting Naivete about Self-Control in the Credit Market." *American Economic Review*, 100(5): 2279–2303.

Heidhues, Paul, Botond Kőszegi, and Takeshi Murooka. 2017. "Inferior Products and Profitable Deception." *Review of Economic Studies*, 84(1): 323–356.

Herne, Kaisa. 1999. "The Effects of Decoy Gambles on Individual Choice." *Experimental Economics*, 2: 31–40.

Huber, Joel, John W. Payne, and Christopher Puto. 1982. "Adding Asymmetrically Dominated Alternatives: Violations of Regularity and the Similarity Hypothesis." *Journal of Consumer Research*, 9(1): 90–98.

Jahedi, Salar. 2011. "A Taste for Bargains." mimeo.

Kamenica, Emir. 2008. "Contextual Inference in Markets: On the Informational Content of Product Lines." *American Economic Review*, 98(5): 2127–2149.

Kőszegi, Botond, and Adam Szeidl. 2013. "A Model of Focusing in Economic Choice." *The Quarterly Journal of Economics*, 128(1): 53–104.

Lazear, Edward P. 1995. "Bait and Switch." *Journal of Political Economy*, 103(4): 813–830.

Mazar, Nina, Botond Kőszegi, and Dan Ariely. 2014. "True Context-Dependent Preferences? The Causes of Market-Dependent Valuations." *Journal of Behavioral Decision Making*, 27(4): 200–208.

Piccione, Michele, and Ran Spiegler. 2012. "Price Competition under Limited Comparability." *Quarterly Journal of Economics*, 127(1): 97–135.

Simonson, Itamar. 1989. "Choice Based on Reasons: The Case of Attraction and Compromise Effects." *Journal of Consumer Research*, 16(2): 158–174.

Appendix

A.1 Auxiliary Results

Some of the results in the main text build on the following auxiliary result.

Lemma A.1 (Naïveté and Fooling)**.** *Let $\beta > 1$. Let product t^k be the target of firm k. Assume that the firm fools consumers of type $\tilde{\beta}^0$ by attracting them with product $a^k \neq t^k$. Then it is true that*

- *Store-context distorts valuations of at least one of the two products: $(\theta_{a^k}, \theta_{t^k}) \neq (N, N)$.*
- *The fooled consumer is (partially) naïve regarding contextual distortions: $\tilde{\beta}^0 \neq \beta$.*

a) *Over-estimators and under-estimators cannot be fooled by the same pair of products (a^k, t^k):*

   If the fooled consumer is over-estimating *($\tilde{\beta}^0 > \beta$), other over-estimating types may be attracted by product $a^k \neq t^k$, but under-estimating consumers always correctly expect to prefer the target. In particular, $\forall \tilde{\beta} < \beta$, $E_{\tilde{\beta}}\left[\hat{u}^k_{t^k}\right] > E_{\tilde{\beta}}\left[\hat{u}^k_{a^k}\right]$.*

   If the fooled consumer is under-estimating *($\tilde{\beta}^0 < \beta$), other under-estimating types may be attracted by product $a^k \neq t^k$, but over-estimating consumers always correctly expect to prefer the target. In particular, $\forall \tilde{\beta} > \beta$, $E_{\tilde{\beta}}\left[\hat{u}^k_{t^k}\right] > E_{\tilde{\beta}}\left[\hat{u}^k_{a^k}\right]$.*

b) *Being fooled increases undistorted surplus for over-estimators, but lowers it for under-estimators:*

   If the fooled consumer is over-estimating *($\tilde{\beta}^0 > \beta$), she receives* higher *undistorted surplus than expected. In particular, $u_{t^k} > u_{a^k}$.*

   If the fooled consumer is under-estimating *($\tilde{\beta}^0 < \beta$), she receives* lower *undistorted surplus than expected. In particular, $u_{t^k} < u_{a^k}$.*

Proof. For ease of notation, we drop the superscript k on products a^k and t^k. Assume that the firm fools the consumer, selling product $t \in J^k$, but attracting the consumer with another product $a \in J^k$, $a \neq t$. Then

(by IC) $$\hat{u}^k_t \geq \hat{u}^k_a$$

(by PCC) $$E_{\tilde{\beta}^0}\left[\hat{u}^k_t\right] \leq E_{\tilde{\beta}^0}\left[\hat{u}^k_a\right]$$

with at least one inequality strict: incentive compatibility (IC) requires that the consumer weakly prefers the target at the store, the perceived choice constraint (PCC) requires that the consumer *expects* to weakly prefer the attraction product at the store. To induce the consumer to switch products, at least one inequality needs to be strict. Fix θ_a and θ_t. If $(\theta_a, \theta_a) = (N, N)$, $E_{\tilde{\beta}^0}\left[\hat{u}_j^k\right] = \hat{u}_j^k = u_j$ for $j \in \{a, b\}$ and any $\tilde{\beta}^0$ such that both, IC and PCC hold with equality, a contradiction. For at least one inequality to be strict, $(\theta_a, \theta_a) \neq (N, N)$: if the consumer is fooled, in-store context distorts the valuation of at least one of the products. Also, if $\tilde{\beta}^0 = \beta$, then $E_{\tilde{\beta}^0}\left[\hat{u}_j^k\right] = \hat{u}_j^k$ for any distortion (θ_a, θ_t). Again, IC and PCC cannot hold with one inequality being strict. Thus, if the consumer is fooled, she must be naïve regarding contextual distortions, that is, $\tilde{\beta}^0 \neq \beta$.

Assume for the rest of the proof that $(\theta_a, \theta_t) \neq (N, N)$ and $\tilde{\beta}^0 \neq \beta$. Define the function $v_j(\gamma) := \hat{u}_j^k|_{\beta=\gamma}$. Note that $v_j(1) = u_j$, $v_j(\beta) = \hat{u}_j^k$ and $v_j(\tilde{\beta}^0) = E_{\tilde{\beta}^0}\left[\hat{u}_j^k\right]$. We can rewrite the fooling constraints as

$$\text{(by IC)} \qquad \hat{u}_t^k \geq \hat{u}_a^k \Leftrightarrow v_t(\beta) \geq v_a(\beta)$$

$$\text{(by PCC)} \qquad E_{\tilde{\beta}^0}\left[\hat{u}_t^k\right] \leq E_{\tilde{\beta}^0}\left[\hat{u}_a^k\right] \Leftrightarrow v_t(\tilde{\beta}^0) \leq v_a(\tilde{\beta}^0).$$

The two conditions imply that if the consumer is fooled, $\exists \gamma^0 \in [\min\{\tilde{\beta}^0, \beta\}, \max\{\tilde{\beta}^0, \beta\}]$ s.t. $v_t(\gamma^0) - v_a(\gamma^0) = 0$ (a point where products a and t generate identical surplus). We first want to show that this crossing is unique. For this, note that for a given distortion θ_j, $\partial v_j / \partial \gamma = const$. Thus, $\partial[v_t(\gamma) - v_a(\gamma)]/\partial \gamma = const$. Because at least one inequality is strict, $\partial[v_t(\gamma) - v_a(\gamma)]/\partial \gamma \neq 0$. Thus, if a crossing exists, it must be unique. It follows:

a) If the consumer is *over-estimating* ($\tilde{\beta}^0 > \beta$) and fooled, then $\exists! \gamma^0 \in [\beta, \tilde{\beta}^0]$ s.t. $v_t(\gamma^0) - v_a(\gamma^0) = 0$. By IC, $v_t(\beta) - v_a(\beta) \geq 0$ and thus, $v_t(\gamma^0) - v_a(\gamma^0) < 0 \; \forall \gamma > \gamma^0$ and $v_t(\gamma^0) - v_a(\gamma^0) > 0 \; \forall \gamma < \gamma^0$.

   This implies that

   - $\forall \tilde{\beta} > \gamma^0$, $v_t(\tilde{\beta}) - v_a(\tilde{\beta}) < 0 \Leftrightarrow E_{\tilde{\beta}}\left[\hat{u}_t^k\right] - E_{\tilde{\beta}}\left[\hat{u}_a^k\right] < 0$: all *over-estimating* agents with $\tilde{\beta} > \gamma^0$ (falsely) expect to prefer product a over product t at the store and are therefore *also* fooled by the pair (a, t). If $\gamma^0 = \beta$, *all* over-estimating agents are fooled.
   - $\forall \tilde{\beta} < \beta \leq \gamma^0$, $v_t(\tilde{\beta}) - v_a(\tilde{\beta}) > 0 \Leftrightarrow E_{\tilde{\beta}}\left[\hat{u}_t^k\right] - E_{\tilde{\beta}}\left[\hat{u}_a^k\right] > 0$: all *under-estimating* agents (correctly) expect to prefer product t over product a at the store and are therefore *not* fooled by the pair (a, t).
   - $v_t(1) - v_a(1) > 0 \Leftrightarrow u_t - u_a > 0$ by $\gamma^0 > 1$: the target generates higher undistorted surplus than the attraction product.

b) If the consumer is *under-estimating* ($\tilde{\beta}^0 < \beta$) and fooled, then $\exists! \gamma^0 \in [\tilde{\beta}^0, \beta]$ s.t.

$v_t(\gamma^0) - v_a(\gamma^0) = 0$. By IC, $v_t(\beta) - v_a(\beta) \geq 0$, and thus, $v_t(\gamma^0) - v_a(\gamma^0) > 0 \;\forall \gamma > \gamma^0$ and $v_t(\gamma^0) - v_a(\gamma^0) < 0 \;\forall \gamma < \gamma^0$.

This implies that

- $\forall \tilde{\beta} < \gamma^0$, $v_t(\tilde{\beta}) - v_a(\tilde{\beta}) < 0 \Leftrightarrow E_{\tilde{\beta}}\left[\hat{u}_t^k\right] - E_{\tilde{\beta}}\left[\hat{u}_a^k\right] < 0$: all under-estimating agents with $\tilde{\beta} < \gamma^0$ (falsely) expect to prefer product a over t at the store and are therefore *also* fooled by the pair (a, t). If $\gamma^0 = \beta$, *all* under-estimating agents are fooled.
- $\forall \tilde{\beta} > \beta \geq \gamma^0$, $v_t(\tilde{\beta}) - v_a(\tilde{\beta}) > 0 \Leftrightarrow E_{\tilde{\beta}}\left[\hat{u}_t^k\right] - E_{\tilde{\beta}}\left[\hat{u}_a^k\right] > 0$: all *over-estimating* agents (correctly) expect to prefer product t over product a at the store and are therefore *not* fooled by the pair (a, t).
- $v_t(1) - v_a(1) < 0 \Leftrightarrow u_t - u_a < 0$ by $\gamma^0 > 1$: the target generates lower undistorted surplus than the attraction product.

□

A.2 Proofs of the Results in the Main Text

We use the following method throughout all proofs to find market supply in the competitive equilibrium: First, we derive the best response of some firm k to a fixed competitor offer $(\boldsymbol{M}^{-k}, \boldsymbol{\Theta}^{-k})$ *conditional* on attracting a positive share of consumers under the assumption that the maximum price b consumers are able to pay is arbitrarily large, i.e., $b \to \infty$. In general, this best response will be unique and continuous in $(\boldsymbol{M}^{-k}, \boldsymbol{\Theta}^{-k})$. Due to this characteristic, in a second step, we can find the competitive market supply by searching for the competitor offer $(\boldsymbol{M}^{-k}, \boldsymbol{\Theta}^{-k})$ that equates the profits of this response to zero. At this point, firms that supply the market will sell a cost-efficient quality (q^*, q^Q, or q^P) at cost, making zero profit. When we drop the assumption $b \to \infty$, consumers will always buy such a product if $b \geq c(q^Q) > c(q^*) > c(q^P)$, which holds by our assumptions on the cost function (see section 2). The (interior) solution we define using this method is thus valid without the assumption $b \to \infty$. Moreover, firms who do not supply the market must always choose $(M^k, \Theta^k) = \emptyset$, because this is the only response that avoids any costs and yields nonnegative profits. While supplying the market at cost and choosing $(M^k, \Theta^k) = \emptyset$ both yield zero profits and are thus best responses, in equilibrium, at least 2 firms must choose to supply the market. Otherwise there would exist some firm k that faced only competitors choosing $(M^k, \Theta^k) = \emptyset$, making a deviation to monopoly profits possible. In general, we therefore have a range of competitive equilibria that all result in the same market supply: At least 2 firms share the market and sell at cost, while all other firms choose $(M^k, \Theta^k) = \emptyset$.

Proof of Lemma 1 (Rational Benchmark). Let $\beta = 1$. This implies that consumers are homogeneous and have time-consistent surplus function $u_j = q_j - p_j$. Context leaves valuations unaffected, $\theta_j^k = N$ for all j and k.

Consider some firm k and fix the competitor offer $\boldsymbol{M}^{-k}$. Let $\bar{u} \geq 0$ be the maximum surplus attainable outside of firm k (this surplus is implicitly defined by $\boldsymbol{M}^{-k}$ and the outside option of no purchase). Let $b \to \infty$ and consider the best response *conditional* on attracting a positive share of consumers. Fix some quality $q_j \geq \underline{q}$. The firm can sell q_j to all consumers at price $p_j = \lim_{\delta \to 0}(q_j - \bar{u} - \delta) = q_j - \bar{u}$, where $\delta > 0$ is the smallest monetary unit. At this price, the firm offers just enough surplus to let consumers marginally improve over the highest surplus available elsewhere, thereby winning all consumers. For given quality q_j, no other price can achieve higher profits: A higher price implies the loss of all consumers, a lower price cannot attract more. This price implies profit $\pi^k = q_j - \bar{u} - c(q_j)$ and thus, the profit-maximizing quality to sell is $q^* := \arg\max[q - c(q)]$, or $c'(q^*) = 1$. Note that $q^* > \underline{q}$ by assumption, making this interior solution valid. Offering additional products is costly and cannot increase profits. It follows: Conditional on attracting a positive share of consumers, the unique best response is the product line $M^k = ((q^*, q^* - \bar{u}))$. Note that the best response so defined is unique and continuous in $\bar{u}$. Market supply in the competitive equilibrium can thus be found by searching for $\bar{u}$ where this response yields zero profits. This unique point exists at $\bar{u} = q^* - c(q^*)$, implying marginal cost pricing, $p_j = c(q^*)$ and the product line $M^* = ((q^*, c(q^*))$. This solution is valid by our model assumption $b > c(q^*)$, such that we can drop the assumption $b \to \infty$.

Given that some firm offers $M^* = ((q^*, c(q^*))$, other firms face $\bar{u} = q^* - c(q^*)$. There are two best responses: (1) Sell $M^* = ((q^*, c(q^*))$ as well, which yields zero profits, (2) Offer nothing, $M^k = \emptyset$, which is the only response avoiding all costs and also yields zero profits. In any equilibrium, at least 2 firms must offer the product line M^*: If no firm offered M^*, then any firm would face an outside option $\bar{u} = 0 < q^* - c(q^*)$ and there would exist a deviation incentive to monopoly profits. If only one firm offered M^*, then, similarly, this firm could earn monopoly profits by deviating. We thus have a range of competitive equilibria that all result in the same market supply: At least 2 firms share the market and offer M^*, while all other firms choose $M^k = \emptyset$.

$\square$

Proof of Lemma 2 (Profitable Fooling). We consider a unique firm k throughout. For ease of notation, we drop the superscript k on products a^k and t^k. Fix $(\boldsymbol{M}^{-k}, \boldsymbol{\Theta}^{-k})$. $(\boldsymbol{M}^{-k}, \boldsymbol{\Theta}^{-k})$ implies an outside option with surplus $\bar{u}(\tilde{\beta}^0) \geq 0$ for a consumer of type $\tilde{\beta}^0$. Assume $\underline{q} \to 0$ and $b \to \infty$. Fix any target quality $q_t = q_t^0 \geq 0$. If the firm does not fool, the consumer correctly expects to purchase target t when entering firm k. Thus, conditional on not fooling, the maximum selling price for quality q_t^0 is

$$p_t^0 := q_t^0 - \bar{u}(\tilde{\beta}^0).$$

For example, the firm could only offer product t (and no other product). Then consumers of type $\tilde{\beta}^0$ enter the store of firm k if $u_t \geq \bar{u}(\tilde{\beta}^0)$. Given quality $q_t = q_t^0$ and price p_t^0, this condition holds with equality. Formally, with $\delta \to 0^+$ being the smallest monetary unit, the firm can achieve that type $\tilde{\beta}^0$ enters the store with certainty by choosing $p_t = \lim_{\delta \to 0}[q_t^0 - \bar{u}(\tilde{\beta}^0) - \delta] = p_t^0$.

Part a) (If type $\tilde{\beta}^0$ is over-estimating ($\tilde{\beta}^0 > \beta$), fooling her is *unprofitable*). Assume that firm k fools type $\tilde{\beta}^0$. We will first show that fooling an over-estimating type ($\tilde{\beta}^0 > \beta$) is unprofitable. For this, we will derive an upper bound on the price for a given target quality q_t^0, $\bar{p}_t(q_t^0)$ (conditional on fooling the consumer and selling her target $t \neq a$), and show that this bound is lower than the price p_t^0.

Consider stage 2, i.e., the decision of the consumer of what product $j \in J^k$ to purchase after she has entered the store of firm k. A lower bound on the (context-dependent) surplus of the target is given by $\hat{u}_t^k = \hat{u}_a^k$: Lowering $\hat{u}_t^k$ by charging a higher price p_t or offering a lower quality q_t will make the consumer choose product a over t, violating incentive compatibility. Rewriting $\hat{u}_t^k = \hat{u}_a^k$ as $(\hat{u}_t^k - u_t) + u_t = (\hat{u}_a^k - u_a) + u_a \Leftrightarrow (\hat{u}_t^k - u_t) + q_t - p_t = (\hat{u}_a^k - u_a) + u_a$ and solving this expression for p_t and q_t, respectively, yields as an upper bound on price:

$$p_t = q_t - u_a + (u_a - u_t), \tag{5}$$

Now consider stage 1. Condition $\hat{u}_t^k = \hat{u}_a^k$ implies that if the consumer is fooled, $E_{\tilde{\beta}^0}\left[\hat{u}_t^k\right] < E_{\tilde{\beta}^0}\left[\hat{u}_a^k\right]$: the consumer expects to *strictly* prefer product a over t at store k. She enters the store if and only if the *expected* purchase—i.e., product a—generates undistorted surplus that is as least as high as her outside option, i.e., $u_a \geq \bar{u}(\tilde{\beta}^0)$. Consider the bound on p_t as defined in Equations (5): Clearly, this bound is maximized if the participation constraint $u_a \geq \bar{u}(\tilde{\beta}^0)$ binds, i.e., if $u_a = \bar{u}(\tilde{\beta}^0)$. We conclude: conditional on fooling and selling target $t \neq a$ to a consumer of type $\tilde{\beta}^0$, an upper bound on the price for given target quality q_t^0 is given by

$$\bar{p}_t(q_t) := q_t - \bar{u}(\tilde{\beta}^0) + (u_a - u_t).$$

It is now easy to see that fooling over-estimating types is not profitable. By Lemma A.1, if an over-estimating type is fooled, $u_a < u_t$: It follows from $\bar{p}_t(q_t)$ and $\underline{q}_t(p_t)$ that if the firm sells $t \neq a$, then it must charge a lower price $p_t < p_t^0$ for quality q_t^0. This concludes the proof for part a).

Part b) (If type $\tilde{\beta}^0$ is under-estimating ($\tilde{\beta}^0 < \beta$), fooling her is *profitable*). Turning to the case of an under-estimating type, note first that if $\tilde{\beta}^0 < \beta$, then by Lemma A.1, $(u_t - u_a) < 0 \Leftrightarrow (u_a - u_t) > 0$ and thus, $\bar{p}_t(q_t^0) > p_t^0$. This suggests that fooling an under-estimating type may be profitable. We will now show that this is indeed the case if and only if in-store context distorts the surplus of products a and t as described in the lemma. Given a distortion (θ_a, θ_t), we can rewrite IC and PCC as conditions on the attributes of products a and t:

- Assume that $(\theta_a, \theta_t) = (Q, Q)$. If $\tilde{\beta}^0 < \beta$, the following two statements are equivalent:
  1. If the consumer is fooled,

     $$\text{(by IC)} \qquad \hat{u}_t^k \geq \hat{u}_a^k \Leftrightarrow \beta(q_t - q_a) \geq p_t - p_a$$
     $$\text{(by PCC)} \qquad E_{\tilde{\beta}^0}\left[\hat{u}_t^k\right] \leq E_{\tilde{\beta}^0}\left[\hat{u}_a^k\right] \Leftrightarrow \tilde{\beta}^0(q_t - q_a) \leq p_t - p_a,$$

     with at least one inequality strict.
  2. If the consumer is fooled, $q_t > q_a$ and $p_t > p_a$ (the firm up-sells).

- Assume that $(\theta_a, \theta_t) = (P, P)$. If $\tilde{\beta}^0 < \beta$, the following two statements are equivalent:
  1. If the consumer is fooled,

     $$\text{(by IC)} \qquad \hat{u}_t^k \geq \hat{u}_a^k \Leftrightarrow \beta(p_a - p_t) \geq q_a - q_t$$
     $$\text{(by PCC)} \qquad E_{\tilde{\beta}^0}\left[\hat{u}_t^k\right] \leq E_{\tilde{\beta}^0}\left[\hat{u}_a^k\right] \Leftrightarrow \tilde{\beta}^0(p_a - p_t) \leq q_a - q_t,$$

     with at least one inequality strict.
  2. If the consumer is fooled, $q_t < q_a$ and $p_t < p_a$ (the firm down-sells).

- Assume that $(\theta_a, \theta_t) = (P, Q)$. If $\tilde{\beta}^0 < \beta$, the following two statements are equivalent:
  1. If the consumer is fooled,

     $$\text{(by IC)} \qquad \hat{u}_t^k \geq \hat{u}_a^k \Leftrightarrow \beta(q_t + p_a) \geq q_a + p_t$$
     $$\text{(by PCC)} \qquad E_{\tilde{\beta}^0}\left[\hat{u}_t^k\right] \leq E_{\tilde{\beta}^0}\left[\hat{u}_a^k\right] \Leftrightarrow \tilde{\beta}^0(q_t + p_a) \leq q_a + p_t,$$

     with at least one inequality strict.
  2. If the consumer is fooled, $q_t + p_a > 0$ and $q_a + p_t > 0$.

- Assume that $(\theta_a, \theta_t) = (P, N)$. If $\tilde{\beta}^0 < \beta$, the following two statements are equivalent:

  1. If the consumer is fooled,

  $$\begin{aligned} &\text{(by IC)} && \hat{u}_t^k \geq \hat{u}_a^k \Leftrightarrow \beta p_a \geq q_a - q_t + p_t \\ &\text{(by PCC)} && E_{\tilde{\beta}^0}\left[\hat{u}_t^k\right] \leq E_{\tilde{\beta}^0}\left[\hat{u}_a^k\right] \Leftrightarrow \tilde{\beta}^0 p_a \leq q_a - q_t + p_t, \end{aligned}$$

  with at least one inequality strict.

  2. If the consumer is fooled, $p_a > 0$ and $q_a - q_t + p_t > 0$.

- Assume that $(\theta_a, \theta_t) = (Q, P)$. If $\tilde{\beta}^0 < \beta$, the following two statements are equivalent:

  1. If the consumer is fooled,

  $$\begin{aligned} &\text{(by IC)} && \hat{u}_t^k \geq \hat{u}_a^k \Leftrightarrow \beta(q_a + p_t) \leq q_t + p_a \\ &\text{(by PCC)} && E_{\tilde{\beta}^0}\left[\hat{u}_t^k\right] \leq E_{\tilde{\beta}^0}\left[\hat{u}_a^k\right] \Leftrightarrow \tilde{\beta}^0(q_a + p_t) \geq q_t + p_a, \end{aligned}$$

  with at least one inequality strict.

  2. If the consumer is fooled, $q_a + p_t < 0$ and $q_t + p_a < 0$.

- Assume that $(\theta_a, \theta_t) = (Q, N)$. If $\tilde{\beta}^0 < \beta$, the following two statements are equivalent:

  1. If the consumer is fooled,

  $$\begin{aligned} &\text{(by IC)} && \hat{u}_t^k \geq \hat{u}_a^k \Leftrightarrow \beta q_a \leq q_t - p_t + p_a \\ &\text{(by PCC)} && E_{\tilde{\beta}^0}\left[\hat{u}_t^k\right] \leq E_{\tilde{\beta}^0}\left[\hat{u}_a^k\right] \Leftrightarrow \tilde{\beta}^0 q_a \geq q_t - p_t + p_a, \end{aligned}$$

  with at least one inequality strict.

  2. If the consumer is fooled, $q_a < 0$ and $q_t - p_t + p_a < 0$.

- Assume that $(\theta_a, \theta_t) = (N, P)$. If $\tilde{\beta}^0 < \beta$, the following two statements are equivalent:

  1. If the consumer is fooled,

  $$\begin{aligned} &\text{(by IC)} && \hat{u}_t^k \geq \hat{u}_a^k \Leftrightarrow \beta p_t \leq q_t - q_a + p_a \\ &\text{(by PCC)} && E_{\tilde{\beta}^0}\left[\hat{u}_t^k\right] \leq E_{\tilde{\beta}^0}\left[\hat{u}_a^k\right] \Leftrightarrow \tilde{\beta}^0 p_t \geq q_t - q_a + p_a, \end{aligned}$$

  with at least one inequality strict.

  2. If the consumer is fooled, $p_t < 0$ and $q_t - q_a + p_a < 0$.

It is obvious from the conditions derived above that any distortion $(\theta_a, \theta_t) \in \{(Q,P),(Q,N),(N,P)\}$ cannot lead to a profitable fooling outcome. In particular, fooling with any one of these distortions requires that either, $q_a < 0$, or $p_t < 0$, or both. But if $q_a < 0$, the attraction product has below minimum quality $\underline{q}$ and can therefore not attract consumers to the store, while if $p_t < 0$, the firm would sell the target strictly below marginal cost and make negative profit. We conclude: If in-store context asymmetrically distorts context in favor of the attraction product, $(\theta_a, \theta_t) \in \{(Q,P),(Q,N),(N,P)\}$, the consumer cannot be fooled to purchase target t at any $p_t \geq 0$.

It remains to be shown that fooling is profitable if $(\theta_a, \theta_t) \in \{(Q,Q),(P,P),(P,Q),(P,N),(N,Q)\}$. In particular, we will show that with any one of these distortions, the firm can indeed sell quality q_t^0 at price $\bar{p}_t(q_t^0) \geq p_t^0$. W.l.o.g, we assume for the rest of the proof that either a and t are the only products that firm k offers, or that other existing products do not violate IC and PCC, that is, $\forall j \in J^k, j \notin \{a,t\}$, $\hat{u}_j^k < \hat{u}_t^k$ and $E_{\tilde{\beta}^0}\left[\hat{u}_j^k\right] < E_{\tilde{\beta}^0}\left[\hat{u}_a^k\right]$. Fix some $q_t^0 > 0$ and $p_t = \bar{p}_t(q_t^0)$. Recall that the construction of $\bar{p}_t(q_t^0)$ implies that IC is satisfied with equality, i.e, $\hat{u}_t^k = \hat{u}_a^k$. Formally, with $\delta \to 0^+$ being the smallest monetary unit, the firm sets $\hat{u}_t^k$ arbitrarily close but above $\hat{u}_a^k$, i.e., $\hat{u}_t^k = \lim_{\delta \to 0}(\hat{u}_a^k + \delta) = \hat{u}_a^k$. Similarly, the construction of $\bar{p}_t^0$ implies that $u_a = \bar{u}(\tilde{\beta}^0)$: If the consumer is fooled, she expects to receive surplus identical to her outside option $\bar{u}(\tilde{\beta}^0)$. Again, by choosing u_a arbitrarily close but above $\bar{u}(\tilde{\beta}^0)$, $u_a = \lim_{\delta \to 0}(u_a + \delta) = \bar{u}(\tilde{\beta})$, the firm can guarantee that the consumer enters its store with certainty. To prove that the firm can sell q_t at $p_t = \bar{p}_t(q_t^0)$ it remains be shown that—given distortion (θ_a, θ_t)—there exists an attraction product with $q_a > 0$ s.t. $E_{\tilde{\beta}^0}\left[\hat{u}_t^k\right] < E_{\tilde{\beta}^0}\left[\hat{u}_a^k\right]$: the consumer expects to strictly prefer product a over t. Note that with q_t and p_t being fixed, $\hat{u}_t^k = \hat{u}_a^k$ determines an one-to-one function between q_a and p_a. We are thus left with only one degree of freedom. Consider the three possible cases listed in the lemma:

- Assume that $(\theta_a, \theta_t) = (Q,Q)$. PCC holds with strict inequality if and only if $q_t > q_a$ and $p_t > p_a$. Pick $q_a \in (0, q_t^0)$, which exists by construction. For example, choose $q_a = \underline{q}$. Then $q_a < q_t$ and by $\hat{u}_t^k = \hat{u}_a^k \Leftrightarrow p_a = p_t - \beta(q_t - q_a) < p_t$: PCC holds with strict inequality. (q.e.d.)

- Assume that $(\theta_a, \theta_t) = (P,P)$. PCC holds with strict inequality if and only if $p_t < p_a$ and $q_t < q_a$. Pick $p_a > \bar{p}_t(q_t^0)$, which exists by construction. For example, choose $p_a = b$. Then $p_a > p_t$ and by $\hat{u}_t^k = \hat{u}_a^k \Leftrightarrow q_a = q_t + \beta(p_a - p_t) > q_t$: PCC holds with strict inequality. (q.e.d.)

- Assume that $(\theta_a, \theta_t) \in \{(P,Q),(P,N),(N,Q)\}$.

1. If $(\theta_a, \theta_t) = (P, Q)$, PCC holds with strict inequality if and only if $q_t + p_a > 0$ and $q_a + p_t > 0$. Pick $p_a > 0$ sufficiently large, s.t. by $\hat{u}_t^k = \hat{u}_a^k \Leftrightarrow q_a = \beta(q_t^0 + p_a) - \bar{p}_t(q_t^0) > 0$. For example, choose $p_a = b$. Then $q_t + p_a > 0$ and $q_a + p_t > 0$ by construction: PCC holds with strict inequality. (q.e.d.)
2. If $(\theta_a, \theta_t) = (P, N)$, PCC holds with strict inequality if and only if $p_a > 0$ and $q_a - q_t + p_t > 0$. Pick $p_a > 0$ sufficiently large, s.t. by $\hat{u}_t^k = \hat{u}_a^k \Leftrightarrow q_a = \beta p_a + q_t - \bar{p}_t(q_t^0) > 0$. For example, choose $p_a = b$. Then $p_a > 0$ and $q_a - q_t + p_t = \beta p_a > 0$ by construction: PCC holds with strict inequality. (q.e.d.)
3. If $(\theta_a, \theta_t) = (N, Q)$, PCC holds with strict inequality if and only if $q_t > 0$ and $q_a - p_a + p_t > 0$. Pick $p_a > 0$ sufficiently large, s.t. by $\hat{u}_t^k = \hat{u}_a^k \Leftrightarrow q_a = p_a + \beta q_t - \bar{p}_t(q_t^0) > 0$. For example, choose $p_a = b$. Then $q_t > 0$ and $q_a - p_a + p_t = \beta q_t > 0$ by construction: PCC holds with strict inequality. (q.e.d.)

This concludes the proof for part b). □

Proof of Proposition 1 (Sophistication/pessimism induces the rational outcome). Let $\beta > 1$ (consumers are context-sensitive). Assume that $\tilde{\beta} \geq \beta$ for all consumers.

We first derive the (unique) best response for a generic firm k *conditional* on attracting a positive share of consumers under the assumption that $b \to \infty$. Fix the competitor offer $(\boldsymbol{M}^{-k}, \boldsymbol{\Theta}^{-k})$ and let $\bar{u}(\tilde{\beta}) \geq 0$ be type $\tilde{\beta}$'s *expected* maximum surplus attainable outside of firm k. By Lemma 3, fooling over-estimating types $\tilde{\beta} > \beta$ is not profitable: The firm can sell any quality $q_t \geq \underline{q}$ at a strictly higher price if $a^k(\tilde{\beta}) = t^k$ for all $\tilde{\beta}$. Also, sophisticated consumers cannot be fooled. It follows that if $\tilde{\beta} \geq \beta$ for all consumers, a firm can sell any target t^k at a strictly higher profit if it does not fool. So assume that the firm does not fool. Drop the superscript k on product t^k for ease of notation. All consumers correctly expect to buy the target when entering firm k and thus, the demand of firm k depends entirely on the characteristics of this target, q_t and p_t. Let $D(q_t, p_t) \in [0, 1]$ be the corresponding demand function of firm k. The profit of firm k is then $\pi^k(q_t, p_t) = D(q_t, p_t)[p_t - c(q_t)]$ and depends only on the characteristics of product t: Offering more than this product is unnecessary yet costly and cannot be part of the best response. Fix quality q_t and p_t at a strictly positive demand $\bar{D} = D(\bar{q}_t, \bar{p}_t)$. Note that $\bar{D} = D(\bar{q}_t, \bar{p}_t) = D(\bar{q}_t + \Delta q, \bar{p}_t + \Delta q) > 0$, for any increment in quality $\Delta q \in \mathbb{R}$, because $\bar{u}_t = \bar{q}_t + \Delta q - (\bar{p}_t + \Delta q) = \bar{q}_t - \bar{p}_t$. Solving for the profit-maximizing

Δq, $\arg\max_{\Delta q} \pi^k(\Delta q) = \bar{D}[\bar{p}_t + \Delta q - c(\bar{q}_t + \Delta q)]$, yields the condition $c'(\bar{q}_t + \Delta q) = 1$. In other words, for any positive demand $\bar{D}$, the profit-maximizing quality to sell is defined by $c(q_t) = 1$, i.e., $q_t = q^*$. This interior solution is valid by assumption that $\underline{q} < q^*$. Because profits (conditional on not fooling) depend on the outside valuation (undistorted surplus) u_t, distorting context at the store is unnecessary yet costly and cannot be part of the best response. It follows that the unique best response is to offer one undistorted product with quality $q_t = q^*$, $\theta_t = N$ at price $p_t = \arg\max_{p \in \mathbb{R}} D(q^* - p)[p - c(q^*)]$ (and no other products). Market supply in any equilibrium must follow this rule: If a firm with a positive market share would choose differently,—by the uniqueness of the best response derived above—there would exist a deviation incentive. The only other response that can be profit-maximizing is to choose $(M^k, \Theta^k) = \emptyset$, i.e., to not supply any products, which yields zero profits. Note that best response behavior is *near-identical* to the rational benchmark: Firms behave *as if* consumers were rational but (possibly) heterogeneous in their outside options $\bar{u}(\tilde{\beta})$. However, at any point of mutual best response, $\bar{u}(\tilde{\beta}) = \bar{u}\ \forall \tilde{\beta} \geq \beta$: If no firm fools, all consumers must expect to receive the same maximum surplus. Once $\bar{u}$ is unique, the unique best response conditional on attracting a positive share of consumers collapses to $(M^k, \Theta^k) = ((q^*, q^* - \bar{u}), (\theta_t^k = N))$—identical to the best response in the rational benchmark. Hence, the competitive equilibrium must conform to the equilibrium derived in Lemma 1. The remainder of the proof is identical to the second part of the proof of Lemma 1 and is therefore omitted. □

Proof of Proposition 2 (Fooling Equilibrium). We derive the equilibrium from the best response of a given firm k to a generic market situation. For ease of notation, we drop the superscript k from products t^k and a^k.

Part a) (Store-Wide Distortions). We begin the proof by considering a perfectly homogeneous, under-estimating consumer population with unique type $\tilde{\beta}^0 < \beta$. Consider a generic firm k. Fix the competitor offer $(\boldsymbol{M}^{-k}, \boldsymbol{\Theta}^{-k})$ and let $\bar{u} = \bar{u}(\tilde{\beta}^0) \geq 0$ be type $\tilde{\beta}^0$'s *expected* maximum surplus attainable outside of firm k. Assume that for any two products j, i at firm k, $\theta_j^k = \theta_i^k = \theta^k$ and the firm chooses $\theta^k \in \{Q, P, N\}$. Assume (for now) that $b \to \infty$. Consider the best response *conditional* on attracting a positive share of consumers. By Lemma 3, the best response will involve fooling and the distortion of context. This yields strictly higher profits than not fooling and choosing $\theta^k = N$. Hence, the best response will involve choosing either $\theta^k = Q$ or $\theta^k = P$. We will now derive the two equilibrium candidates that derive from assuming either $\theta^k = Q$ or $\theta^k = P$.

- Assume that $\theta^k = Q$. The maximum price the firm can sell any target quality q_t is given by the upper bound $\bar{p}_t(q_t)$ which we have derived in the proof of Lemma 2. To achieve $\bar{p}_t(q_t)$, offering a second product $a \neq t$ is necessary and sufficient. Holding more than 2 products is unnecessary yet costly and can thus not be part of the best response. If $\theta^k = Q \Rightarrow (\theta_a, \theta_t) = (Q, Q)$, by (the proof of) Lemma 2, the firm sells q_t at $p_t = \bar{p}_t(q_t)$ if and only if it chooses $q_a < q_t$ and $p_a < p_t$. If $(\theta_a, \theta_t) = (Q, Q)$, $\bar{p}_t(q_t)$ can be rewritten as

$$\bar{p}_t(q_t, q_a, p_a) = \beta(q_t - q_a) + p_a,$$

under the condition $q_a - p_a = \bar{u}$ (the participation constraint binds). To find the best response, we need to choose q_t and (q_a, p_a) such that profit at this price is maximized. Consider the choice of q_t first. Because quality q_t is inflated by a factor β when $(\theta_a, \theta_t) = (Q, Q)$, it is easy to see that the cost-efficient quality to sell is

$$q_t = q^Q := \arg\max_q [\beta q - c(q)] \Leftrightarrow c'(q^Q) = \beta.$$

This interior solution is valid by assumption $q^Q > \underline{q}$. We are left with the choice of the attraction product (q_a, p_a). Maximizing profit for any q_t implies maximizing $\bar{p}_t(q_t, q_a, p_a)$ under the constraint $q - p_a = \bar{u}$. There are 2 opposing forces: Minimizing q_a and maximizing p_a. The profit-maximizing choice is to minimize q_a: Because quality q_a is inflated at the store, the positive effect on profits of decreasing quality q_a is larger than the positive effect of increasing price p_a. The unique profit-maximizing choice is therefore to choose $q_a = \underline{q}$, which implies $p_a = \underline{q} - \bar{u}$. Note that this choice satisfies the fooling conditions $q_a < q_t$ and $p_a < p_t$ for any $q_t > 0$. For later reference, note the marketing implications of this best response: To attract consumers, the firm fixes attraction quality $q_a = \underline{q}$ and competes with other firms on the price of this low-quality product. We conclude: Conditional on $\theta^k = Q$, the best response in the domain of positive profits is unique and continuous: the firm offers 2 products, t and $a \neq t$, with $(q_t, p_t) = (q^Q, \bar{p}_t(q^Q))$ and $(q_a, p_a) = (\underline{q}, \underline{q} - \bar{u})$.

- Assume that $\theta^k = P$. Analogously to the case of $\theta^k = Q$, we find the best response by maximizing profit at price $\bar{p}_t(q_t)$ which we can now express as

$$\bar{p}_t(q_t, q_a, p_a) = p_a - \frac{1}{\beta}(q_a - q_t)$$

under the condition $q_a - p_a = \bar{u}$ (the participation constraint binds). 2 products, t and $a \neq t$ are necessary and sufficient to yield this maximum price for any target quality q_t.

Holding more products cannot be part of a best response. By (the proof of) Lemma 2, the firm sells q_t at $p_t = \bar{p}_t(q_t)$ if and only if it chooses $q_a > q_t$ and $p_a > p_t$. With price of the target being inflated at the store, the cost-efficient quality to sell is

$$q_t = q^P := \arg\max_q \left[q - \beta c(q)\right] \Leftrightarrow c'(q^P) = \frac{1}{\beta}.$$

This interior solution is valid by assumption $q^P \geq \underline{q}$. Maximizing profit for any q_t implies maximizing $\bar{p}_t(q_t, q_a, p_a)$ under the constraint $q_a - p_a = \bar{u}$. There are 2 opposing forces: Minimizing q_a and maximizing p_a. Contrary to the case of $\theta^k = Q$, the profit-maximizing choice now is to maximize p_a: Because price p_a is inflated at the store, the positive effect on profits of increasing price p_a is larger than the positive effect of decreasing quality q_a. The unique profit-maximizing choice is therefore to choose $p_a = b$, which implies $q_a = b + \bar{u}$. Note that this choice satisfies the fooling conditions $p_a > p_t$ and $q_a > q_t$ for any $p_t < b$. For later reference, note the marketing implications of this best response: To attract consumers, the firm fixes attraction price $p_a = b$ and competes with other firms on the quality of this high-price product. We conclude: Conditional on $\theta^k = P$, the best response in the domain of positive profits is unique and continuous: the firm offers 2 products, t and $a \neq t$, with $(q_t, p_t) = (q^P, \bar{p}_t(q^P))$ and $(q_a, p_a) = (b + \bar{u}, b)$.

Note that the best response in both cases is independent of the degree of naïveté of type $\tilde{\beta}^0 < \beta$: Due to the optimality condition $\hat{u}_t^k = \hat{u}_a^k$ (the IC binds), any consumer with belief $\tilde{\beta} < \beta$ (falsely) believes to purchase product a with certainty. The best response does *not* generate heterogeneous expectations among a purely under-estimating consumer population. If firms play mutual best responses, any heterogeneity in types $\tilde{\beta}$ is therefore rendered unimportant for market supply: Uniqueness of the best response (given a distortion θ^k) implies that firms generating positive demand must choose according to it; otherwise, there would exists a strict deviation incentive. This response does not generate heterogeneous expectations. Firms not generating positive demand, on the other hand, choose $(M^k, \Theta^k) = \emptyset$ to avoid positive costs and thus negative profits. These firms do not generate heterogeneous expectations either. It follows that in any equilibrium, $\bar{u}(\tilde{\beta}) = \bar{u}\ \forall \tilde{\beta} < \beta$: the outside option is a unique value. We can find market supply in the competitive equilibrium by searching for $\bar{u}$ that equates the best response profits to zero. This yields the following two candidates for equilibrium market supply:

$$(Q^*) \qquad \theta^k = Q,\ (q_t, p_t) = (q^Q, c(q^Q)),\ (q_a, p_a) = (\underline{q}, c(q^Q) - \beta(q^Q - \underline{q}))$$

$$(P^*) \qquad \theta^k = P,\ (q_t, p_t) = (q^P, c(q^P)),\ (q_a, p_a) = (q^P + [b - c(q^P)], b)$$

By reasoning analogous to the second part of the proof of Lemma 1, at least 2 firms must provide a product line according to (Q^*) or (P^*). These firms share the market. All other firms choose $(M^k, \Theta^k) = \emptyset$. Fix an equilibrium where at least one firm chooses (M^k, θ^k) according to (Q^*). Then there must be at least one other firm that provides the same expected surplus $\bar{u} = u_a = \underline{q} - c(q^Q) + \beta(q^Q - \underline{q})$. Otherwise, the firm would have a deviation incentive to strictly positive profits. What remains to be checked is a deviation towards the other regime $\theta^k = P$, where the maximum profit is given by the unique best response defined above. In other words, the firm may offer 2 products, t and $a \neq t$, with $(q_t, p_t) = (q^P, \bar{p}_t(q_t))$ and $(q_a, p_a) = (b + \bar{u}, b)$. There exists a strict deviation incentive if and only if, under this formulation, $q_t - p_t > 0$. Rearranging, this is the case if and only if $\nu^{(Q,Q)} < \nu^{(P,P)}$, where

$$\begin{aligned}
\nu^{(Q,Q)} &:= (q^Q - c(q^Q)) + (\beta - 1)(q^Q - \underline{q}), \text{ and} \\
\nu^{(P,P)} &:= (q^P - c(q^P)) + (\beta - 1)(b - c(q^P)).
\end{aligned}$$

Analogously, in an equilibrium where at least one firm plays according to (P^*), firms have a deviation incentive towards $\theta^k = Q$ if and only if $\nu^{(Q,Q)} > \nu^{(P,P)}$.

We conclude: A competitive equilibrium exists. In any such equilibrium, at least 2 firms share the market. These firms offer 2 products, t and $a \neq t$. All other firms choose $(M^k, \Theta^k) = \emptyset$. The characteristics of t and a as well as θ^k are uniquely defined by (Q^*) if and only if $\nu^{(Q,Q)} > \nu^{(P,P)}$ and by (P^*) if and only if $\nu^{(Q,Q)} < \nu^{(P,P)}$. If $\nu^{(Q,Q)} = \nu^{(P,P)}$, any firm that supplies the market chooses t and a according to either (Q^*) or (P^*). As a final step, we can drop the assumption that $b \to \infty$. In particular, our characterization is valid for any $b \geq c(q^Q) > c(q^P)$ as assumed in the model section of this paper. This concludes the proof for part a) (Store-Wide-Distortions).

Part b) (Product-Specific Distortions). Assume that firms choose $\theta_j^k \in \{Q, P, N\}$ for each product $j \in J^k$ individually. The proof works similarly as the proof for part a). We start again with the assumption of a homogeneous population with unique type $\tilde{\beta}^0 < \beta$ and determine the best response *conditional* on attracting a positive market share under the assumption that $b \to \infty$. Fix the competitor offer $(\boldsymbol{M}^{-k}, \boldsymbol{\Theta}^{-k})$ and let $\bar{u} = \bar{u}(\tilde{\beta}^0) \geq 0$ be type $\tilde{\beta}^0$'s *expected* maximum surplus attainable outside of firm k. Fix some quality $q_t \geq \underline{q}$. By the proof of Lemma 2, the maximum price the firm can sell any q_t is $\bar{p}_t(q_t)$. To sell at this price, a second product $a \neq t$ is necessary and sufficient. Offering more products is unnecessary yet costly and hence, cannot be part of the best response. Moreover, the (in-store) valuation of at least one product $j \in \{a, t\}$ must be distorted, in particular, with any distortion $(\theta_a, \theta_t) \in \{(Q, Q), (P, P), (P, Q), (P, N), (N, Q)\}$, a strictly higher price than without fooling can be realized. It follows that the best response must involve one of these distortions. It is easy to see that choosing $(\theta_a, \theta_t) = (P, Q)$ strictly dominates any other choice of (θ_a, θ_t): No other distortion yields an overvaluation of the target relative to the attraction product that is as extreme. This is of course reflected in $\bar{p}_t(q_t)$, which under the condition that $q_a - p_a = \bar{u}$ (the participation constraint binds) can be rewritten as

$$\bar{p}_t(q_t, q_a, p_a) = \begin{cases} \beta q_t - q_a + \beta p_a & \text{if } (\theta_a, \theta_t) = (P, Q) \\ q_t - q_a + \beta p_a & \text{if } (\theta_a, \theta_t) = (P, N) \\ \beta q_t - q_a + p_a & \text{if } (\theta_a, \theta_t) = (N, Q) \\ \beta q_t - \beta q_a + p_a & \text{if } (\theta_a, \theta_t) = (Q, Q) \\ \frac{1}{\beta} q_t - \frac{1}{\beta} q_a + p_a & \text{if } (\theta_a, \theta_t) = (P, P). \end{cases}$$

If $(\theta_a, \theta_t) \in \{(P, P), (P, Q), (P, N), (N, Q)\}$ (all except $(\theta_a, \theta_t) = (Q, Q)$), $\bar{p}_t(q_t, q_a, p_a)$ is maximized by choosing $p_a = b$ (which implies $q_a = b + \bar{u}$). This choice satisfies all fooling constraints: the consumer indeed enters the store of firm k and buys q_t (see also the proof of Lemma 2). Clearly, $(\theta_a, \theta_t) = (P, Q)$ yields the highest price. To see the dominance of $(\theta_a, \theta_t) = (P, Q)$ over $(\theta_a, \theta_t) = (Q, Q)$, note that $\bar{p}_t(q_t, q_a, p_a)$ is maximized under $(\theta_a, \theta_t) = (Q, Q)$ by choosing $q_a = \underline{q}$ (which implies $p_a = q_a - \bar{u}$). This is a feasible choice also when $(\theta_a, \theta_t) = (P, Q)$, which yields a strictly higher price. Hence, only $(\theta_a, \theta_t) = (P, Q)$ can be part of the best response.

Given 2 products t and $a \neq t$ as well as distortion $(\theta_a, \theta_t) = (P, Q)$, we need to define the profit-maximizing choice of (q_t, p_t) and (q_a, p_a). For any q_t, the profit-maximizing price is $p_t = \bar{p}_t(q_t, q_a, p_a)$ as defined above. With quality being inflated at the store, the cost-efficient choice of q_t is

$$q_t = q^Q := \arg\max_q [\beta q - c(q)] \Leftrightarrow c'(q^Q) = \beta.$$

This interior solution is valid by assumption $q^Q > \underline{q}$. We have already noted above that $\bar{p}_t(q_t, q_a, p_a)$ is maximized by choosing $p_a = b$ and thus, $q_a = b + \bar{u}$. While there are 2 opposing forces when maximizing $\bar{p}_t(q_t, q_a, p_a)$—minimizing q_a and maximizing p_a—, maximizing p_a is the dominant choice: Due to price p_a being inflated at the store, a marginal increase in price (accompanied by a marginal increase in quality) always yields a higher marginal effect on profits than the equivalent decrease in quality. The marketing implication of this choice is identical to the case of a purely price-inflated store (see case a), $\theta^k = P$): Competition outside the store is on quality q_a and not on price. It follows: Conditional on attracting a positive share of consumers, the unique best response of firm k is to offer 2 products, t and $a \neq t$, choose distortion $\Theta^k = (\theta_a, \theta_t) = (P, Q)$, and product-characteristics $(q_t, p_t) = (q^Q, \bar{p}_t(q^Q))$ and $(q_a, p_a) = (b + \bar{u}, b)$.

Identical to part a) of this proof, the best response is independent of the degree of naïveté of type $\tilde{\beta}^0 < \beta$. Also, firms not generating positive demand choose $(M^k, \Theta^k) = \emptyset$ to avoid positive costs and thus negative profits. By analogous statements as those in part a) it follows that at any point of mutual best response, any heterogeneity in types $\tilde{\beta}$ is rendered unimportant for market supply. It follows that in any equilibrium, $\bar{u}(\tilde{\beta}) = \bar{u}\ \forall \tilde{\beta} < \beta$: the outside option is a unique value. We can find market supply in the competitive equilibrium by searching for $\bar{u}$ that equates the profits in the best response defined above to zero. We conclude: A competitive equilibrium exists. In any such equilibrium, at least 2 firms share the market. These firms offer 2 products, t and $a \neq t$. All other firms choose $(M^k, \Theta^k) = \emptyset$. For the firms that share the market, Θ^k and the characteristics of products a and t are uniquely defined by

$$(PQ^*) \qquad \Theta^k = (\theta_a, \theta_t) = (P, Q),\ (q_t, p_t) = (q^Q, c(q^Q)),\ (q_a, p_a) = (\beta(q^Q + b) - c(q^Q), b).$$

As a final step, we can drop the assumption that $b \to \infty$. In particular, our characterization is valid for any $b \geq c(q^Q)$ as assumed in the model section of this paper. This concludes the proof for part b). □

Proof of Proposition 3 (Fooling with Salience, Focusing, or Relative Thinking). The proof is constructed as follows: We will first consider Assumption F (Focusing) and Assumption RT (Relative Thinking). Both of these assumptions imply that distortions are store-wide, i.e., for any two products j, i in J^k, $\theta_j^k = \theta_i^k = \theta^k$. Following Proposition 2, if firms have a technology that allows for store-wide distortions, they will either want to fool with $(\theta_a, \theta_t) = (Q, Q)$ or $(\theta_a, \theta_t) = (P, P)$. We will show that if context is a function of the product line and follows Assumption F or Assumption RT, a firm can construct $(\theta_a, \theta_t) = (Q, Q)$ and $(\theta_a, \theta_t) = (P, P)$ and fool according to the best response defined in the proof of Proposition 2 if and only if it introduces a third product to the product line. In other words, one *decoy* is necessary and sufficient to fool according to Proposition 2, part a). We will then turn to Assumption S (Salience). This assumption allows firms to construct *product-specific* distortions. Proposition 2, part b) has shown that if firms can choose θ_j^k for each product individually, they will want to fool with $(\theta_a, \theta_t) = (P, Q)$. Again, we will show that under Assumption S, the firm can construct such distortion and fool according to the best response defined in the proof of Proposition 2 if and only if it introduces a third product to the product line. In other words, one *decoy* is necessary and sufficient to fool according to Proposition 2, part b).

We make the proof by concentrating on a unique firm k throughout, allowing us to drop the superscript k on most variables such as the target t, the attraction product a, the decoy

d and the surplus function at store k, $\hat{u}_j$.

Assumption F (Focusing).

Step 1: Fooling is not possible without a third product (a decoy is necessary). Assume that Assumption F holds and that firm k offers only two products a and t, $a \neq t$. The firm may either fool with $(\theta_a, \theta_t) = (Q, Q)$ or $(\theta_a, \theta_t) = (P, P)$.

- Assume that the firm fools with $(\theta_a, \theta_t) = (Q, Q)$. We have shown in the proofs of Lemma A.1 and Lemma 2 that fooling an under-estimating consumer ($\tilde{\beta} < \beta$) with $(\theta_a, \theta_t) = (Q, Q)$ implies $u_t < u_a$, $q_t > q_a$ and $p_t > p_a$ (the firm up-sells). Note that $u_t < u_a \Leftrightarrow q_t - q_a < p_t - p_a$. By Assumption F this implies that $\Delta_q^k < \Delta_p^k$, and thus, $\theta^k \in \{P, N\}$, a contradiction.

- Assume that the firm fools with $(\theta_a, \theta_t) = (P, P)$. We have shown in the proofs of Lemma A.1 and Lemma 2 that fooling an under-estimating consumer ($\tilde{\beta} < \beta$) with $(\theta_a, \theta_t) = (P, P)$ implies $u_t < u_a$, $q_t < q_a$ and $p_t < p_a$ (the firm down-sells). Note that $u_t < u_a \Leftrightarrow p_a - p_t < q_a - q_t$. By Assumption F this implies that $\Delta_q^k > \Delta_p^k$, and thus, $\theta^k \in \{Q, N\}$, a contradiction.

This concludes the proof of step 1. Note that this result is not an artefact of our rank-based implementation of KS, but a generic characteristic of the Focusing framework, which requires, by Assumption 1 in KS, that whenever preferences are shifted towards a product that is dominant in one attribute (z), but not in the other ($-z$), $\Delta_z^k > \Delta_{-z}^k$, which is in contradiction to fooling condition $u_t < u_a$.

Step 2: Fooling is always possible with a third product (a single decoy is sufficient). Assume that Assumption F holds and that firm k offers three products, a, t, and d. The firm may either fool with $(\theta_a, \theta_t) = (Q, Q)$ or $(\theta_a, \theta_t) = (P, P)$.

- Assume that the firm wants to fool with $(\theta_a, \theta_t) = (Q, Q)$. Fix any characteristics (q_a, p_a) and (q_t, p_t) that imply that the consumer is fooled if $(\theta_a, \theta_t) = (Q, Q)$. Then $u_t < u_a$, $q_t > q_a$, and $p_t > p_a$. (For example, choose the characteristics defining the best response in the proof of Proposition 2.) Choose $p_d = p_t$ and $q_d < q_t - (p_t - p_a) - \kappa_F$.[25] Then by Assumption F, $\Delta_q^k - \Delta_p^k > \kappa_F \Leftrightarrow \theta^k = Q \Rightarrow (\theta_a, \theta_t) = (Q, Q)$. Note that product d is strictly dominated by target t and thus, does not violate incentive compatibility of the fooling regime: The firm can sell product t according to the best response defined in the proof of Proposition 2.

[25] Recall from Assumption F that $\kappa_F \geq 0$ is some (exogenously defined) threshold that measures the level of stimulus necessary for a preference distortion.

- Assume that the firm wants to fool with $(\theta_a, \theta_t) = (P, P)$. Fix any characteristics (q_a, p_a) and (q_t, p_t) that imply that the consumer is fooled if $(\theta_a, \theta_t) = (P, P)$. Then $u_t < u_a$, $q_t > q_a$, and $p_t > p_a$. (For example, choose the characteristics defining the best response in the proof of Proposition 2.) Choose $q_d = q_t$ and $p_d > p_t + (q_a - q_t) + \kappa_F$.[26] Then by Assumption F, $\Delta_p^k - \Delta_q^k > \kappa_F \Leftrightarrow \theta^k = P \Rightarrow (\theta_a, \theta_t) = (P, P)$. Note that product d is strictly dominated by target t and thus, does not violate incentive compatibility of the fooling regime: The firm can sell product t according to the best response defined in the proof of Proposition 2.

This concludes the proof of step 2. Again, notice that this a result generic to the Focusing framework and does not depend on the rank-based formulation of preferences that we have assumed for our model. The Focusing framework assumes utility weights to be a function of the attribute spread Δ_z^k. Most naturally, such spreads are open to manipulation by a single option, i.e., a single decoy.

It follows that under Assumption F, the characterization of products a and t corresponds to the equilibrium defined in Proposition 2, part a). Holding more than three products is unnecessary yet costly which implies that the fooling equilibrium of Proposition 2 will be realized with exactly three products of which one is a decoy.

Assumption RT (Relative Thinking).

Step 1: Fooling is not possible without a third product (a decoy is necessary). We show that norming assumptions N1 and N2 in BRS imply that fooling is impossible with only 2 products. The result then readily extends to Assumption RT. Attention weights in BRS are a function of the spread of an attribute in the choice set, Δ_z, $z = q, p$; we call the weight function $w(\Delta_z)$. By N1, $w(\Delta_t)$ is strictly decreasing in Δ_z. By N2, $w(\Delta_z)\Delta_z$ is strictly increasing in Δ_z.

Suppose that firm k offers only two products, a and t. Fooling requires that $u_a > u_t$ (see Lemma A.1) while inside the store $\hat{u}_t \geq \hat{u}_a$ by incentive compatibility (IC). We show that the norming assumptions in BRS rule out such a preference change if a and t are the only products in the product line.

Assume $u_a > u_t \Leftrightarrow q_a - p_a > q_t - p_t$. Then either (1) $q_a > q_t$ and $p_a > p_t$, or (2) $q_a < q_t$ and $p_a < p_t$, or (3) $q_a > q_t$ and $p_a < p_t$. If (1) is true, then $u_a > u_t \Leftrightarrow \Delta_q > \Delta_p$. N2 then implies $w(\Delta_q)\Delta_q > w(\Delta_p)\Delta_p$, which is equivalent to $w(\Delta_q)s_a - w(\Delta_p)p_a > w(\Delta_q)q_t - w(\Delta_p)p_t$. Thus, the same product is preferred outside and inside the store and fooling is not possible. If (2) is true, $u_a > u_t \Leftrightarrow \Delta_q < \Delta_p$ leads to a similar contradiction. Then N2 implies

[26] See the previous footnote.

$w(\Delta_q)\Delta_q < w(\Delta_p)\Delta_p$ and product a (now being the less qualitative option) will be preferred both inside and outside the store. Finally, if (3) is true, product a dominates product t in both attributes implying that a is strictly preferred over t for any (positive) attribute weights. Again, product a is preferred both inside and outside the store and fooling is not possible.

These results readily extend to Assumption RT. Case (3) is immediate. For case (1), note that t is the less-qualitative product. A preference change towards t would thus require $\theta^k = P$, that is, $\frac{\Delta_q}{\Delta_p} > \kappa_{RT} \geq \beta$.[27] But with only two products spanning the attribute range Δ_q, $\Delta_q > \beta\Delta_p$ implies that the product with higher quality is preferred at the store. That is, $\hat{u}_a^k > \hat{u}_t^k$, which contradicts incentive compatibility (IC). With case (2) we get a similar contradiction. This concludes the proof of step 1.

Step 2: Fooling is always possible with a third product (a single decoy is sufficient). Assume that Assumption RT holds and that firm k offers three products, a, t, and d. The firm may either fool with $(\theta_a, \theta_t) = (Q, Q)$ or $(\theta_a, \theta_t) = (P, P)$.

- Assume that the firm wants to fool with $(\theta_a, \theta_t) = (Q, Q)$. Fix any characteristics (q_a, p_a) and (q_t, p_t) that imply that the consumer is fooled if $(\theta_a, \theta_t) = (Q, Q)$. Then $u_t < u_a$, $q_t > q_a$, and $p_t > p_a$. (For example, choose the characteristics defining the best response in the proof of Proposition 2.) Choose $q_d = q_t$ and $p_d > p_a + \kappa_{RT}(q_t - q_a) > p_t$.[28] Then by Assumption RT, $\frac{\Delta_p^k}{\Delta_q^k} > \kappa_{RT} \Leftrightarrow \theta^k = Q \Rightarrow (\theta_a, \theta_t) = (Q, Q)$. Note that product d is strictly dominated by target t and thus, does not violate incentive compatibility of the fooling regime: The firm can sell product t according to the best response defined in the proof of Proposition 2.

- Assume that the firm wants to fool with $(\theta_a, \theta_t) = (P, P)$. Fix any characteristics (q_a, p_a) and (q_t, p_t) that imply that the consumer is fooled if $(\theta_a, \theta_t) = (P, P)$. Then $u_t < u_a$, $q_t < q_a$, and $p_t < p_a$. (For example, choose the characteristics defining the best response in the proof of Proposition 2.) Choose $p_d = p_t$ and $q_d < q_a - \kappa_{RT}(p_a - p_t)$, which implies $q_d < q_t$. Then by Assumption RT, $\frac{\Delta_q^k}{\Delta_p^k} > \kappa_{RT} \Leftrightarrow \theta^k = P \Rightarrow (\theta_a, \theta_t) = (P, P)$. Note that product d is strictly dominated by target t and thus, does not violate incentive compatibility of the fooling regime: The firm can sell product t according to the best response defined in the proof of Proposition 2.

This concludes the proof of step 2. Similar to the Focusing framework, this result is generic to the model by BRS and does not depend on our rank-based implementation. The

[27]Recall from Assumption RT that $\kappa_{RT} \geq \beta$ is some (exogenously defined) threshold that measures the level of stimulus necessary for a preference distortion.

[28]See the previous footnote.

framework of Relative Thinking assumes utility weights to be a function of the attribute spread Δ_a^k. Most naturally, such spreads are open to manipulation by a single option, i.e., a single decoy.

It follows that under Assumption RT, the characterization of products a and t corresponds to the equilibrium defined in Proposition 2, part a). Holding more than three products is unnecessary yet costly which implies that the fooling equilibrium of Proposition 2 will be realized with exactly three products of which one is a decoy.

Assumption S (Salience)

We have shown in the proof of Proposition 2 that if firms can choose θ_j^k for each product individually, the unique weakly undominated best response is to fool with $a \neq t$ and choose $(\theta_a, \theta_t) = (P, Q)$. We show that under Assumption S, this choice is possible if and only if the firm adds a third product (i.e., a single decoy) to the product line.

Step 1: A decoy is necessary. The specifications of products a and t that a best-responding firm will choose are given in the proof of Proposition 2. We show that a distortion $(\theta_a, \theta_t) = (P, Q)$ with these product specifications cannot be constructed without the help of additional (decoy) products. Note first that the specification Proposition 2 implies $q_a > q_t$ and $p_a > p_t$. Thus, none of the two products is dominated. Suppose that the firm only holds these two products. Then the reference quality is given by $z_R^k = \frac{(q_a+q_t)}{2}$ and the reference price is given by $p_R^k = \frac{(p_a+p_t)}{2}$. Because $(q_j - q_R^k)(p_j - p_R^k) > 0$ for $j \in \{a, f\}$, we can exploit Proposition 1 in BGS: The "advantageous" attribute of product j—higher quality *or* lower price relative to the reference—is overweighted if and only if $\frac{q_j}{p_j} > \frac{q_R^k}{p_R^k}$. Also, if and only if $\frac{q_j}{p_j} < \frac{q_R^k}{p_R^k}$, then the "disadvantageous" attribute of product j is overweighted, while if and only if $\frac{q_j}{p_j} = \frac{q_R^k}{p_R^k}$, consumers weigh both attributes equally.

Assume towards a contradiction that the firm can construct $(\theta_a, \theta_t) = (P, Q)$. For t being quality-salient, by $q_t < q_R^k$ and Proposition 1 in BGS,

$$\frac{q_t}{p_t} < \frac{q_R^k}{p_R^k} \Leftrightarrow \frac{q_t}{p_t} < \frac{q_a}{p_a}.$$

But for a being price-salient, by $q_a > q_R^k$ and Proposition 1 in BGS,

$$\frac{q_a}{p_a} < \frac{q_R^k}{p_R^k} \Leftrightarrow \frac{q_t}{p_t} > \frac{q_a}{p_a},$$

a contradiction. This concludes the proof of step 1.

Step 2: A single decoy is sufficient. Assume that firm k wants to fool using distortion $(\theta_a, \theta_t) = (P, Q)$ and chooses the specifications of product a and t according to the best response defined in the proof of Proposition 2 part b). Note that by this specification $q_a > q_t \geq \underline{q} > 0$ and $p_a > p_t > 0$.

- Assume that $\frac{q_t}{p_t} > \frac{q_a}{p_a}$. We construct a reference point using one additional product d that satisfies the following properties: (1) $p_R^k = p_t$, (2) $q_R^k < q_t$ and (3) $\frac{q_a}{p_a} < \frac{q_R^k}{p_R^k} < \frac{q_t}{p_t}$. The construction is illustrated in Figure 2. With such a reference point,

    1. Product t is quality-salient: By $p_R^k = p_t$, the salience of p_t is $\sigma(p_t, p_t)$. By homogeneity of degree zero, $\sigma(\alpha p_t, \alpha p_t) = \sigma(p_t, p_t)$ for any $\alpha > 0$. Let $\alpha = \frac{q_t}{p_t} > 0$, then $\sigma(p_t, p_t) = \sigma(q_t, q_t)$. By ordering, $\sigma(q_t, q_t) < \sigma(q_t, q_R^k)$ because $q_R^k < q_t$. Thus, $\sigma(q_t, q_R^k) > \sigma(p_t, p_R^k)$: product t is quality-salient.
    2. Product a is price-salient: By $q_R^k < q_t < q_a$ and $p_R^k = p_t < p_a$, $(q_a - q_R^k)(p_a - p_R^k) > 0$, and product a neither dominates nor is dominated by the reference good. Thus, Proposition 1 in BGS applies. Because $q_a > q_R^k$, by $\frac{q_R^k}{p_R^k} > \frac{q_a}{p_a}$, product a is price-salient.

  To satisfy property (1), choose $p_d = 2p_t - p_a$, which implies $p_d < p_t$. To satisfy property (2) and (3), choose $q_d < 2q_t - q_a$, which implies $q_d < q_t$. It remains to be shown that the decoy d does not violate fooling conditions. Note that $q_d - p_d < 2q_t - q - a - (2p_t - p_a) \Leftrightarrow u_d < 2u_t - u_a$. Because $u_t < u_a$ by the specifications of a and t, this implies that $u_d < u_t < u_a$. We first show that IC is not violated: Because t is quality-salient, $\hat{u}_t^k = \beta q_t - p_t > u_t$. But then, if (i) $\theta_d^k = N$, $\hat{u}_t^k > \hat{u}_d^k$ follows from $\hat{u}_t^k > u_t > u_d = \hat{u}_d^k$, if (ii) $\theta_d^k = Q$, $\hat{u}_t^k > \hat{u}_d^k$ follows from $q_d < q_t$, $p_d < p_t$ and $u_t > u_d$, if (iii) $\theta_d^k = P$, then $\hat{u}_t^k > \hat{u}_d^k$ if and only if $\hat{u}_a^k > \hat{u}_d^k \Leftrightarrow q_a - q_d > \beta(p_a - p_d)$ by $\hat{u}_t^k = \hat{u}_a^k$. To prove that $q_a - q_d > \beta(p_a - p_d)$, note that $q_a - q_d > q_a - (2q_t - q_a) = 2(q_t - q_a)$ by $q_d < 2q_t - q_a$ and $p_a - p_d = p_a - (2p_t - p_a)$ by $p_d = 2p_t - p_a$. Thus $q_a - q_d > \beta(p_a - p_d)$ if $2(q_a - q_t) > 2\beta(p_a - p_t) \Leftrightarrow (q_a - q_t) > \beta(p_a - p_t)$. But the latter inequality is true by $\hat{u}_t^k = \hat{u}_a^k \Leftrightarrow q_a - \beta q_t = \beta p_a - p_t$. Thus, $\hat{u}_t^k > \hat{u}_d^k$. Finally, we have to show that PCC is not violated, i.e., that $E_{\tilde{\beta}}\left[\hat{u}_a^k\right] > E_{\tilde{\beta}}\left[\hat{u}_d^k\right]$. To see that this is true note that we have shown that $u_a > u_t > u_d$ and $\hat{u}_a^k = \hat{u}_t^k > \hat{u}_d^k$. Because $E_{\tilde{\beta}}\left[\hat{u}_a^k\right]$ is between $\hat{u}_a^k$ and u_a and $E_{\tilde{\beta}}\left[\hat{u}_d^k\right]$ is between $\hat{u}_d^k$ and u_d (both by $\tilde{\beta} < \beta$) it follows that $E_{\tilde{\beta}}\left[\hat{u}_a^k\right] > E_{\tilde{\beta}}\left[\hat{u}_d^k\right]$.

- Assume that $\frac{q_t}{p_t} < \frac{q_a}{p_a}$. We construct a reference point using one additional product d that satisfies the following properties: (1) $q_R^k = q_a$, (2) $p_R^k > p_a$ and (3) $\frac{q_a}{p_a} > \frac{q_R^k}{p_R^k} > \frac{q_t}{p_t}$. The construction is illustrated in Figure 2. With such a reference point,

1. Product t is quality-salient: By $q_R^k > q_t$ and $p_R^k > q_t$, $(q_t - q_R^k)(p_t - p_R^k) > 0$, and product t neither dominates nor is dominated by the reference good. Thus, Proposition 1 in BGS applies. Because $q_t < q_R^k$, by $\frac{q_R^k}{p_R^k} > \frac{q_t}{p_t}$, product t is quality-salient.
2. Product a is price-salient: By $q_R^k = q_a$, the salience of q_a is $\sigma(q_a, q_a)$. By homogeneity of degree zero, $\sigma(\alpha q_a, \alpha q_a) = \sigma(q_a, q_a)$ for any $\alpha > 0$. Let $\alpha = \frac{p_a}{q_a} > 0$, then $\sigma(q_a, q_a) = \sigma(p_a, p_a)$. By ordering, $\sigma(p_a, p_a) < \sigma(p_a, p_R^k)$ because $p_R^k > q_t$. Thus, $\sigma(q_a, q_R^k) < \sigma(p_a, p_R^k)$: product a is price-salient.

To satisfy property (1) choose $q_d = 2q_a - q_t > q_a$. To satisfy property (2) and (3), choose $p_d > 2p_a - p_t$. It remains to be shown that the decoy d does not violate fooling conditions. But note that $p_d > p_a = b$: The decoy has a price above the maximum willingness to pay and thus, will never be chosen (and can therefore not violate fooling conditions).

This concludes the proof of step 2. We conclude: Under Assumption S, the characterization of products a and t corresponds to the equilibrium defined in Proposition 2, part b). Holding more than three products is unnecessary yet costly which implies that the fooling equilibrium of Proposition 2 will be realized with exactly three products of which one is a decoy.

□

Proof of Proposition 4 (Co-Existence of Sophisticated and Naïve Agents). The result is trivial if either *all* consumers are under-estimating or *all* consumers are sophisticated/over-estimating. These cases were covered by our earlier Propositions. So assume that there exists a positive mass of consumers with beliefs $\tilde{\beta} < \beta$ and a positive mass of consumers with beliefs $\tilde{\beta} \geq \beta$.

Let $\beta > 1$. Fix a market offering according to the Proposition. There exists two types of stores with strictly positive demand, k^L and k^H. Type k^L is a fooling firm that supplies products according to the equilibrium defined in Proposition 2 and k^H is a non-fooling firm that supplies products according to the equilibrium defined in Proposition 1. There exists at least 2 firms of each type. All other firms choose $(M^k, \Theta^k) = \emptyset$. All firms make zero profits. Note that conditional on purchasing at type k^H, all consumers expect to purchase q^* at price $p^* = c^*$. At the same time, conditional on purchasing at type k^L, all sophisticated and over-estimating consumers (correctly) expect to purchase the target at the fooling firms (see Lemma A.1 for a proof), while all under-estimators (falsely) expects to purchase the attraction product. We prove that a competitive equilibrium with this market supply exists and that it defines the unique competitive market supply.

Existence. Assume that we have an equilibrium. Firms of type k^L fool and sell quality $q_t \neq q^*$ at $p_t = c(q_t)$ to the under-estimators, while firms of type k^H are truthful and sell q^* at $p^* = c^*$ to the sophisticates/over-estimators. We have to check whether consumers or firms want to deviate. Consider first the under-estimating population that are assumed to purchase at k^L. They have the alternative to purchase (q^*, c^*) at k^H instead of (q_a, p_a) at k^L (of course, they only *expect* to buy product a, while they really buy the target t). However, because $u_a > u^*$ in the candidate equilibrium defined above, purchasing at k^L always promises a higher payoff and under-estimators will not switch to k^H:

1. Consider an equilibrium according to Proposition 2 part b). Then $u_a = q_a - p_a = \beta(q^Q + b) - c^Q - b$. Note that $(\beta - 1)b > 0$ by $\beta > 1$. Strict convexity of the cost function then implies $\beta q^Q - q^Q > \beta q^* - c^* > q^* - c^*$ and thus, $u_a > u^*$.

2. Consider an equilibrium according to Proposition 2 part a) where $(\theta_a, \theta_t) = (Q, Q)$. Then $u_a = q_a - p_a = \underline{q} - c^Q + \beta(q^Q - \underline{q})$. Note that by assumption, $\underline{q} < q^*$ and thus, $u_a > q^* - c^Q + \beta(q^Q - q^*)$ by $\beta > 1$. It follows that $u_a > u^*$ because $q^* - c^Q + \beta(q^Q - q^*) > q^* - c^* \Leftrightarrow \beta q^Q - c^Q > \beta q^* - c^*$ by strict convexity of the cost function.

3. Consider an equilibrium according to Proposition 2 part a) where $(\theta_a, \theta_t) = (P, P)$. Then $u_a = q_a - p_a = q^P + (\beta b - c^P) - b$ and $u_a > u^* \Leftrightarrow q^P - \beta c^P + \beta b - b > q^* - c^*$ $\Leftrightarrow q^P/\beta - c^P + (1 - 1/\beta)b > q^*/\beta - c^* + (1 - 1/\beta)c^*$. Note first that $b > c^*$, so $(1 - 1/\beta)b > (1 - 1/\beta)c^*$. Strict convexity of the cost function further implies that $q^P/\beta - c^P > s^*/\beta - c^*$. So $u_a > u^*$.

Over-estimators also do not want to switch to k^L. They correctly expect to buy product t at k^L (for a proof see Lemma A.1) which generates surplus $u_t = q_t - c(q_t)$. Because $q^* = \arg\max(q - c(s))$ and $q_t \neq q^*$ by strict convexity of $c(q)$, $u^* = q^* - c^* > u_t$ and shopping at type k^H generates higher surplus. Finally, firms of either type have no incentive to deviate. By Proposition 2, no firm can find a more profitable strategy when serving under-estimating agents if there are at least 2 firms of type k^L. By Proposition 1, no firm can find a more profitable strategy when serving over-estimating agents if there are at least 2 firms of type k^H. The only strategy that yields non-negative profits when not generating demand is $(M^k, \Theta^k) = \emptyset$. This strategy yields zero profits as well and thus, does not constitute a deviation incentive. Hence, this is an equilibrium. Note also that any firm that is not of type k^L or k^H must choose $(M^k, \Theta^k) = \emptyset$ by above reasoning. (q.e.d.)

Uniqueness. The proofs of Propositions 1 and 2, respectively, show that unless there exist at least 2 firms supplying products according to Proposition 1 as well as at least 2 firms supplying products according to Proposition 2, there exists a deviation incentive to a strategy with strictly positive profits. In particular, by the uniqueness and continuity of the best response conditional on attracting only sophisticated/over-estimating consumers (Proposition 1), there must exist at least 2 firms supplying a product with expected surplus $\bar{u}^H \geq u^* = q^* - c(q^*)$ to consumers of type $\tilde{\beta} \geq \beta$. Otherwise, at least one firm could attract the entire population of types $\tilde{\beta} \geq \beta$ at strictly positive profit. Similarly, there must exist at least 2 firms supplying a product with expected surplus $\bar{u}^L \geq u_a = q_a - p_a$ to consumers of type $\tilde{\beta} < \beta$, where q_a and p_a are defined by the equilibrium characterized in Proposition 2. Otherwise, at least one firm could attract the entire population of types $\tilde{\beta} < \beta$ at strictly positive profit. By the strict difference of u_a and u^* (in particular, $u_a > u^*$, see the existence proof above), 1 firm cannot satisfy both of these conditions at the same time (attracting both groups of consumers with positive probability), even if it would play a mixed strategy: Such a firm would either have to make negative profits in expectation (to attract both groups without generating a deviation incentive for other firms) or generate an offer that (for at least one of the two groups of consumers) could be profitably undercut by other firms. It follows that at least 2 firms satisfying the respective condition must exist for each group *separately*. Because each firm only serves one group of consumers, the only possibility to satisfy the respective condition without making negative profit is for each firm to choose market supply according to Propositions 1 and 2, respectively. It follows that any competitive equilibrium must have the characteristics listed in the Proposition. (q.e.d.) □

Proof of Proposition 5 (Co-Existence of Rational and Naïve Agents). Fix a consumer population of unit mass with a share $\eta > 0$ being context-sensitive ($\beta > 1$) and under-estimating of this sensitivity ($\tilde{\beta} < \beta$) and the remaining share $(1 - \eta) > 0$ being rational ($\beta = 1$). For ease of notation, we refer to the first group simply as naïves. We continue concentrating on interior solutions (regarding the choice of target quality q_{t^k} and price p_{t^k}) by assuming, throughout, that $b \to \infty$.

Fix any Nash equilibrium. By homogeneity, naïves and rationals share the same preferences in stage 1, i.e., outside stores. We now show that they also share the same expectations about which product they will purchase in stage 2. This implies that both consumer groups will enter the same firm (with probability one if there is one firm that offers the highest surplus in expectation and with strictly positive probability if there are multiple firms that offer the highest surplus in expectation). Consider any firm k. There are two cases: (1) If the firm does not fool, all consumer types correctly expect to purchase the target t^k and thus, have the same expectations. (2) If the firm fools, context-sensitive consumers purchase target t^k, but, by the definition of fooling (Definition 3), there exists some naïve, under-estimating type who expects to purchase some other product $a^k \neq t^k$. By Lemma A.1, $u_{a^k} > u_{t^k}$ in this case, implying that all rational consumers are attracted by the attraction product a^k as well, which (in comparison to context-sensitive consumers) they also purchase. Profit-maximization implies further that the firm sets $\hat{u}^k_{t^k} = \hat{u}^k_{a^k}$: At $\hat{u}^k_{t^k} < \hat{u}^k_{a^k}$, context-sensitive consumers would purchase the attraction product and there would be no fooling. At $\hat{u}^k_{t^k} > \hat{u}^k_{a^k}$, a ceteris paribus increase in the price of the target p_{t^k} (or a decrease in its quality q_{t^k}) would increase the profit-margin on fooled consumers as well as (weakly) increasing the share of naïve consumers who are being fooled (by decreasing $\hat{u}^k_{t^k}$, the firm makes people with increasingly smaller deviations from sophistication also believe that they will prefer the attraction product at the store). At $\hat{u}^k_{t^k} = \hat{u}^k_{a^k}$, however, $E_{\tilde{\beta}}[\hat{u}^k_{t^k}] < E_{\tilde{\beta}}[\hat{u}^k_{a^k}]$ for all $\tilde{\beta} < \beta$: *Any* under-estimating, context-sensitive consumer expects to purchase (and is attracted by) the attraction product a^k. Thus, all consumers share the same expectations. It follows that at any point of mutual best response, there is a unique maximum surplus $\bar{u} > 0$ that both rational and naïve consumers expect to receive and are attracted by. In any equilibrium then, all consumers purchase at the same firms. Moreover, if a firm attracts all consumers of one group, it also attracts all consumers of the other group.

We now consider the best response of some firm k to a given competitor offer conditional on attracting a positive share of consumers. Denote the expected utility that all consumers expect to receive outside of firm k, $\bar{u} > 0$. For ease of notation, we drop the superscript k on all variables of firm k. The firm can either choose to to *not* fool, selling some product j at price $p_j = q_j - \bar{u}$ (generating surplus $u_j = \bar{u}$) to all consumers and yielding profit $\pi = p_j - c(q_j)$, or it can choose to fool, in which case the firm sells two different products to naïves (target t) and rationals (attraction product a), yielding profit $\pi = \eta(p_t - c(q_t)) + (1-\eta)(p_a - c(q_a))$. If the firm fools, profit maximization implies that (1) $\hat{u}_t = \hat{u}_a$, and (2) $u_a = \bar{u}$.[29]

It is clear that fooling yields higher profit than not fooling. Without fooling, the firm maximizes profit by selling $q_j = q^*$ at $p_j = q^* - \bar{u}$. If the firm fools, it could still attract with a product of the same characteristics, sell it at unchanged profit to rationals, while increasing profits on the naïves by inducing them to buy another product at the store (see the proof of Lemma A.1 for a formal proof of this claim). We will now determine the *optimal* fooling strategy, that is, the optimal choice of the attraction product and the target. Assume, w.l.o.g., that the firm offers only two products, the attraction product a and the target t.

We begin with store-wide distortions, that is, for all $\{i, j\} \subseteq J^k$, $\theta^k_i = \theta^k_j = \theta^k \in \{Q, P\}$,

[29]Otherwise, the firm could increase the price of the target (1) or the price of the attraction product (2) without affecting demand, violating the profit-maximum.

and, as a first step, define the optimal choice of (q_a, p_a) and (q_t, p_t) for a given context θ^k.

Assume $\theta^k = Q$. From the two optimality conditions, $\hat{u}_t = \hat{u}_a$ and $u_a = \bar{u}$, we find $p_t = \beta(q_t - q_a) + p_a$ and $p_a = q_a - \bar{u}$. Profit is

$$\pi(q_t, q_a) = \eta\left[\beta q_t - (\beta - 1)\, q_a - c(q_t)\right] + (1-\eta)\left[q_a - c(q_a)\right] - \bar{u}.$$

First-order conditions $\frac{\partial \pi}{\partial q_t} = 0$ and $\frac{\partial \pi}{\partial q_a} = 0$ yield $c'(q_t) = \beta \Leftrightarrow q_t = q^Q$ and $c'(q_a) = 1 - \frac{\eta}{1-\eta}(\beta - 1)$, respectively. Second-order conditions hold by strict convexity of $c(q)$. Quality q_a so defined is valid if and only if it yields $q_a \geq \underline{q}$, so we have $q_a = \underline{q}_a := \max\{\underline{q}, q|_{c'(q_a)=1-\frac{\eta}{1-\eta}(\beta-1)}\}$. Note that for any positive share of naïves, $\eta > 0$, $q_a < q^* < q_t$ (the firm up-sells naïve consumers). As $\eta \to 0$, q_a approaches the rational benchmark, $q_a \to q^*$, from below. Fixing $\theta^k = Q$, we can find equilibrium market prices by setting $\pi = 0$. This yields

$$\begin{aligned} p_a &= \eta c(q^Q) + (1-\eta)c(q_a) - \eta \cdot\beta(q^Q - q_a) \\ p_t &= \eta c(q^Q) + (1-\eta)c(q_a) + (1-\eta)\cdot\beta(q^Q - q_a). \end{aligned}$$

In such an equilibrium, $p_t > c(q_t)$ and $p_a < c(q_a)$ if and only if $\beta q^Q - c(q^Q) > \beta q_a - c(q_a)$, which holds by strict convexity of $c(q)$ and by $q^Q = \arg\max[\beta q - c(q)]$. As $\eta \to 0$, product-supply for the rational consumers approaches the rational benchmark ($q_a \to q^*$, $p_a \to c(q^*)$), while the exploitation of naïve consumers persists ($q_t = q^Q \neq q^*$ and $p_t \to c(q^*) + \beta(q^Q - q^*) > c(q_t)$).

Assume $\theta^k = P$. From the two optimality conditions, $\hat{u}_t = \hat{u}_a$ and $u_a = \bar{u}$, we find $p_t = p_a - \frac{1}{\beta}(q_a - q_t)$ and $p_a = q_a - \bar{u}$. Profit is

$$\pi(q_t, q_a) = \eta\left[\frac{1}{\beta}\cdot q_t + \left(1 - \frac{1}{\beta}\right)\cdot q_a - c(q_t)\right] + (1-\eta)\left[q_a - c(q_a)\right] - \bar{u}.$$

First-order conditions $\frac{\partial \pi}{\partial q_t} = 0$ and $\frac{\partial \pi}{\partial q_a} = 0$ yield $c'(q_t) = \frac{1}{\beta} \Leftrightarrow q_t = q^P$ and $c'(q_a) = 1 + \frac{\eta}{1-\eta}\cdot\left(1 - \frac{1}{\beta}\right) \Leftrightarrow q_a = \bar{q}_a := q|_{c'(q)=1+\frac{\eta}{1-\eta}\cdot\left(1-\frac{1}{\beta}\right)}$, respectively. Second-order conditions hold by strict convexity of $c(q)$. Note that for any positive share of naïves, $\eta > 0$, $q_a > q^* > q_t$ (the firm down-sells naïve consumers). As $\eta \to 0$, q_a approaches the rational benchmark, $q_a \to q^*$, from above. Fixing $\theta^k = P$, we can find equilibrium market prices by setting $\pi = 0$. This yields

$$\begin{aligned} p_a &= \eta c(q^P) + (1-\eta)c(q_a) + \eta \cdot\frac{1}{\beta}(q_a - q^P) \\ p_t &= \eta c(q^P) + (1-\eta)c(q_a) - (1-\eta)\cdot\frac{1}{\beta}(q_a - q^P). \end{aligned}$$

In such an equilibrium, $p_t > c(q_t)$ and $p_a < c(q_a)$ if and only if $\frac{1}{\beta} \cdot q^P - c(q^P) > \frac{1}{\beta} \cdot q_a - c(q_a)$, which holds by strict convexity of $c(q)$ and by $q^P = \arg\max [q - \beta c(q)]$. As $\eta \to 0$, product-supply for the rational consumers approaches the rational benchmark ($q_a \to q^*$, $p_a \to c(q^*)$), while the exploitation of naïve consumers persists ($q_t = q^P \neq q^*$ and $p_t \to c(q^*) - \frac{1}{\beta}(q^* - q^P) > c(q_t)$).

To derive the choice of $\theta^k \in \{Q, P\}$ in equilibrium, fix an equilibrium with $\theta^k = Q$ to find

$$\bar{u} = q_{a^{-k}} - p_{a^{-k}} = \eta \left[[q^Q - c(q^Q)] + (\beta - 1) \left(q^Q - \underline{q}_a \right) \right] + (1 - \eta) [\underline{q}_a - c(\underline{q}_a)] =: \nu^{(Q,Q)}.$$

Substitute $\bar{u}$ in the (best response) profit function when choosing the opposite context $\theta^k = P$,

$$\pi^k = \eta \left[\frac{1}{\beta} \cdot q^P + \left(1 - \frac{1}{\beta} \right) \cdot \bar{q}_a - c(q^P) \right] + (1 - \eta) [\bar{q}_a - c(\bar{q}_a)] - \bar{u} =: \nu^{(P,P)} - \nu^{(Q,Q)}.$$

If $\pi^k > 0 \Leftrightarrow \nu^{(Q,Q)} < \nu^{(P,P)}$, equilibrium choice of in-store context is $\theta^k = P$, if $\pi^k < 0 \Leftrightarrow \nu^{(Q,Q)} > \nu^{(P,P)}$, it is $\theta^k = Q$, and in the knife-edge case of $\pi^k = 0 \Leftrightarrow \nu^{(Q,Q)} = \nu^{(P,P)}$, firms may choose either of the two in equilibrium.

Now consider product-specific distortions, that is, the possibility of constructing different distortions for products a and t. The proof of Proposition 2, part b) confirms the intuition that the firm is best-off choosing $(\theta_a^k, \theta_t^k) = (P, Q)$. From the two optimality conditions, $\hat{u}_t = \hat{u}_a$ and $u_a = \bar{u}$, we find $p_t = \beta q_t + \beta p_a - q_a$ and $p_a = q_a - \bar{u}$. Profit is

$$\pi(q_t, q_a) = \eta [\beta q_t - (\beta - 1) q_a - \beta \bar{u} - c(q_t)] + (1 - \eta) [q_a - \bar{u} - c(q_a)].$$

First-order conditions $\frac{\partial \pi}{\partial q_t} = 0$ and $\frac{\partial \pi}{\partial q_a} = 0$ yield $c'(q_t) = \beta \Leftrightarrow q_t = q^Q$ and $c'(q_a) = 1 + \frac{\eta}{1-\eta}(\beta - 1) \Leftrightarrow q_a = q|_{c'(q)=1+\frac{\eta}{1-\eta}(\beta-1)}$, respectively. Second-order conditions hold by strict convexity of $c(q)$. Note that for any positive share of naïves, $\eta > 0$, $q_a > q^*$ and as $\eta \to 0$, q_a apporaches the rational benchmark, $q_a \to q^*$, from above. Whether the firm up- or down-sells, however, now depends on the share of naïves in the population: If the majority of consumers is rational $\eta \leq \frac{1}{2}$, the firm up-sells ($q_a \leq q_t = q^Q$), and if $\eta > \frac{1}{2}$, it down-sells ($q_a > q_t = q^Q$). We can find equilibrium market prices by setting $\pi = 0$. This yields

$$
\begin{aligned}
p_a &= [\eta c(q^Q) + (1-\eta)c(q_a) \qquad\qquad\qquad + \eta \cdot (q_a - \beta q^Q)] \cdot \frac{1}{1+\eta(\beta-1)} \\
p_t &= [\beta \cdot \eta c(q^Q) + \beta \cdot (1-\eta)c(q_a) - (1-\eta) \cdot (q_a - \beta q^Q)] \cdot \frac{1}{1+\eta(\beta-1)}
\end{aligned}
$$

In such an equilibrium, $p_t > c(q_t)$ and $p_a < c(q_a)$ if and only if $\beta q^Q - c(q^Q) > q_a - \beta c(q_a)$, which holds by strict convexity of $c(q)$ and by $q^Q = \arg\max[\beta q - c(q)]$. As $\eta \to 0$, product-supply for the rational consumers approaches the rational benchmark ($q_a \to q^*$, $p_a \to c(q^*)$), while the exploitation of naïve consumers persists ($q_t = q^Q \neq q^*$ and $p_t \to \beta c(q^*) + \beta q^Q - q^* > c(q_t)$).

$\square$