Biased Beliefs in Search Markets

Tobias Gamp (HU Berlin)
Daniel Krähmer (University of Bonn)

Discussion Paper No. 365

December 30, 2022

Collaborative Research Center Transregio 190 | www.rationality-and-competition.de
Ludwig-Maximilians-Universität München | Humboldt-Universität zu Berlin
Spokesperson: Prof. Georg Weizsäcker, Ph.D., Humboldt University Berlin, 10117 Berlin, Germany
info@rationality-and-competition.de

Biased Beliefs in Search Markets

Tobias Gamp[†] and Daniel Krähmer[‡]

May 2022

Abstract

We study the implications of biased consumer beliefs for search market outcomes in the seminal framework due to Diamond (1971). Biased consumers base their search strategy on a belief function which specifies for any (true) distribution of utility offers in the market a possibly incorrect distribution of utility offers. If biased consumers overestimate the best offer in the market, a novel type of equilibrium may emerge in which firms make exceptionally favourable offers in order to meet biased consumers' unreasonable high expectations which then become partially self-fulfilling. Consequently, the presence of biased consumers may improve the welfare of all consumers.

Keywords: Consumer search, Bounded Rationality, Cursed Beliefs.
JEL Classification No.: D18, D21, D43, D83

We would like to thank Yair Antler, Paul Heidhues, Maarten Janssen, Stephan Lauermann, Dmitry Lubensky, Eeva Mauring, Moritz Meyer-ter-Vehn, Jose-Luis Moraga Gonzalez and Andrew Rhodes as well as seminar audiences at Humboldt University Berlin, Bonn University, Vienna University, Toulouse University, the Berlin IO Day, the Search and Matching Workshop in Bristol, the Workshop on Correlation Neglect and Belief Formation in Berlin, the 7th Workshop on Behavioral Game Theory at the University of East Anglia, the Consumer Search and Switching Costs Workshop in Los Angeles, the EARIE Conference in Barcelona and the Consumer Search Digital Seminar Series for helpful comments and suggestions. We gratefully acknowledge financial support from the German Research Foundation (DFG) through Germany's Excellence Strategy EXC 2126/1–390838866 and CRC TR 224 (Krähmer) and CRC TR 190 (Gamp).

[†]Humboldt University Berlin, Spandauer Straße 1, D-10178 Berlin; email: *tobias.gamp@hu-berlin.de*

[‡]University of Bonn, Institute for Microeconomics, Adenauer Allee 24-42, D-53113 Bonn (Germany); email: *kraehmer@hcm.uni-bonn.de*

1 Introduction

Consumer search is an important element in safeguarding competition and consumer welfare in markets. A key determinant of consumers' search incentives is their beliefs about the benefits of search. In equilibrium models of consumer search, consumers can correctly assess their search benefits, because in equilibrium they know the aggregate distribution of offers in the market. Evidence suggests, however, that in practice, this might not always be the case, especially if a product is purchased rarely.[1] To illustrate that forming correct beliefs might be difficult, consider a consumer who seeks to buy a product that she purchases rarely such as equipment for a leisure activity, a home appliance or new furniture. To assess her search benefits, she not only has to know the (marginal) distribution of a number of relevant product characteristics such as prices and qualities on the market, but also their joint distribution in order to assess how likely it is that she actually finds a product that is both cheap and good. Without experience, it is unlikely that all consumers know all relevant information or are able to correctly process it.

In this article, we study implications for search market outcomes when some "biased" consumers' beliefs only partially reflect the true market conditions. To capture this, we introduce a belief function which specifies for any (true) distribution of utility offers in the market a biased distribution of utility offers upon which biased consumers base their search strategy. This entails that the beliefs of biased consumers in our model depend on other agents' strategies (and thus the true market conditions) but will typically be incorrect in equilibrium.[2] We impose minimal regularity assumptions on the belief function requiring that small or monotone (in the first order sense) changes in the true distribution of offers result in small or monotone changes in beliefs[3], but otherwise allow for arbitrary belief biases.

Our article makes two main points. First, we show how biased consumer beliefs shape market outcomes, and second, we address welfare implications. To bring out the role of consumer beliefs most clearly, we consider the seminal search model by Diamond (1971) with the novelty

[1] For example, Matsumoto and Spence (2016) show that students, especially inexperienced ones, underestimate the benefits of searching an online market place for textbook prices. **?** elicit consumer beliefs about prices for a popular home appliance and show that while consumers underestimate the average price, they overestimate the dispersion of prices prior to their search. More generally, Price and Zhu (2016) illustrate the importance of consumers' beliefs for their search behaviour. They show for a variety of service markets such as electricity, broadband internet and current bank accounts that a consumer's belief about how much money she could save by shopping around, correlates strongly with her reported search and switching behaviour.

[2] Our belief function shares this feature with other theories on how boundedly rational agents form beliefs such as analogy-based equilibrium (Jehiel (2005)), cursed equilibrium (Eyster and Rabin (2005)), or Berk-Nash equilibrium (Esponda and Pouzo (2016)). Similar to our model, Heller and Winter (2020) endow agents with an abstract belief function, but, unlike us, allow them to choose their belief function strategically.

[3] We impose a third, more technical condition, as explained later.

that a fraction of consumers has biased beliefs. In this model, finitely many firms offer a homogeneous product and consumers can search firms sequentially and in random order so as to learn how much utility each firm offers before acquiring a product. When all consumers have correct beliefs, the Diamond paradox where all firms make the monopoly offer is the only equilibrium outcome: because in equilibrium, all consumers anticipate that all firms make the same offer, search becomes pointless inducing consumers to accept the offer from the first firm they visit which, in turn, provides firms with monopoly power.

We show that *non-pessimism* is the critical property of the belief function that determines whether other market outcome can occur relative to the benchmark. We refer to a belief function as non-pessimistic if there is a utility distribution whose best offer a biased consumer overestimates. Intuitively, non-pessimism captures that a consumer overestimates the tails of the utility distribution which is a key factor in determining her search benefits. Note that non-pessimism is a fairly weak property, because the existence of already a single utility distribution for which a naive consumer overestimates the value of the best offer guarantees that the belief function is non-pessimistic.[4]

We find that novel equilibrium types can occur if and only if biased consumers have non-pessimistic beliefs, and we fully characterize these equilibria as a function of search costs.[5] More specifically, under non-pessimism, for sufficiently large search costs, even unrealistically hopeful biased consumers consider search too costly, and unsurprisingly, the Diamond outcome occurs. However, for intermediate search costs, the Diamond outcome can no longer be sustained, because as a consequence of non-pessimism, biased consumers overestimate the value of the best equilibrium offer and would find it worthwhile to search for a golden egg: an offer which in reality does not exist. For example, in the introductory examples above, non-pessimistic beliefs arise when the consumer underestimates the correlation between price and quality in the market, leading her to overestimate the occurrence of cheap high quality products. The resulting equilibria display a form of utility dispersion that, intuitively, reflects periods in which firms make what we refer to as *penny sales offers*, that is, with positive probability they make offers which supply distinctly more utility than a regular offer. From the perspective of firms, penny sales offers serve the purpose to cater to biased consumers who demand better offers than standard consumers to stop their search and buy. In this sense, a biased consumer's belief that the market supplies golden eggs is partially confirmed in equilibrium. Finally, for sufficiently small search costs, biased consumers become so demanding that it becomes unprofitable for firms to attract them with special offers, and biased consumers, in a futile search effort, visit all firms before returning to the best

[4]To be precise, our regularity conditions imply, however, that if there is one such distribution, then there is a small class of distributions for which this must be the case.

[5]If the belief function is pessimistic, then the Diamond outcome is the only equilibrium outcome.

one. The equilibrium outcome is then the same as in Stahl (1989) in which so-called *shoppers*, a consumer group with zero search costs, visit all firms. Thus, non-pessimism provides a somewhat surprising foundation for the existence of shoppers in an environment where all consumers have positive search costs.

The second main goal of our article is to study welfare implications of biased beliefs. We address this question in the framework where biased consumers hold *cursed* beliefs as in Eyster and Rabin (2005). This corresponds to a form of correlation neglect where, as in the introductory examples above, a consumer underestimates the correlation between the cons (price) and the pros (quality) of the offers in the market.[6] The basic observation is that a more strongly cursed consumer holds a more dispersed belief about the utility distribution in the market in the mean-preserving spread sense. Thus, higher cursedness translates into higher benefits of search and thus into a higher reservation value.

We show that the effects of higher cursedness on consumer welfare are ambiguous and depend on the absolute degree of cursedness. Starting from low levels, increasing cursedness may stimulate search by naive consumers and break the Diamond outcome, because firms start to compete to attract the searching biased consumers. As cursedness increases further, however, biased consumers become more demanding in their willingness to accept an offer, and firms' trade-off between maximizing market share and surplus share tilts towards the latter. As a result, firms increasingly focus on extracting surplus from the less demanding standard consumers, and the average utility they offer to consumers goes down after a certain level of cursedness is reached. Hence, if one interprets cursedness as a policy variable that a regulator can influence by education or transparency campaigns, these findings suggest that an intermediate level of cursedness maximizes consumer welfare.[7]

Related literature

Our article is related to a recent literature that studies the role of consumer beliefs in search settings. Antler and Bachi (2021) study a marriage market where agents entertain wrong beliefs about other agent's propensity to consider them to be an acceptable match. Like in our model with non-pessimistic beliefs, some agents then overestimate the benefits of search, engage in excessive search and may end up being eternal singles, in particular, when search frictions are small. In their marriage market, utility is non-transferable and no market side sets prices, a key focus of our article. Mauring (2020) studies a labour market where workers wrongly believe that past

[6]See, e.g., Enke and Zimmermann (2019).

[7]The comparative static results with respect to the share of biased consumers or the number of firms are analytically intractable. We provide numerical results which, in particular, show that biased consumer welfare might be non-monotonic in the share of biased consumers or the number of firms.

wages predict current wages. These wrong beliefs determine the worker's order of search, but not the trade-off between accepting an offer or continuing search, as in this article. Gamp and Krähmer (2021) study a search market where, as a consequence of having mistaken beliefs, some consumers misjudge a firm's true quality. By contrast, in this article, naive consumers observe the true utility of any product they encounter. Somewhat orthogonal to our approach, Janssen and Shelegia (2020) highlight the role of beliefs in a model with rational consumers by showing that consumer beliefs after a price deviation off the equilibrium path can have a key impact on market outcomes.[8]

Our article is also related to a literature in behavioural industrial organization (for a review, see Heidhues and Kőszegi (2018)).[9] A key question is how the presence of naive consumers affects consumer welfare. Our point that with cursed consumer beliefs an intermediate level of naivete is optimal for consumer welfare shares similarities with the findings in Ispano and Schwardmann (2020) but contrasts with Armstrong and Chen (2009) where an intermediate fraction of naive consumers minimizes consumer welfare. In both articles, naive consumers effectively underestimate quality differences between firms. In Ispano and Schwardmann (2020) this may improve consumer welfare, as it reduces the mark-up of high quality firms that otherwise would not have to compete with low quality ones. In Armstrong et al. (2009), it may relax competition and harm all consumers, as it allows firms to differentiate in high quality firms targeting sophisticated consumers and (inefficiently) low quality ones targeting naive consumers with low prices. In contrast, in our setting naive consumers observe the utility of the products they encounter, and price competition is stimulated when they overestimate the differences between firms, as this creates incentives to search.

Finally, our article is broadly related to a literature which, like us, is interested in market outcomes when consumers do not know the true distribution of market offers but, unlike us, assumes that there is aggregate uncertainty.[10]

---

[8]Related, Parakhonyak and Sobolev (2015) study search market outcomes when consumers have no prior beliefs at all but rather search to minimize regret in a (hypothetical) worst case scenario. In this case, in contrast to this article, a consumer's optimal search rule does not depend on the true market conditions.

[9]In a search context, Heidhues et al. (2021) argue that regulating complexly priced products may induce consumers to search more firms, enhancing competition and improve consumer welfare, as they spend less time studying offers carefully. Karle et al. (2020) show that loss aversion induces naive consumers to reduce the number of firms they visit, benefiting in particular low quality ones.

[10]See, e.g., Benabou and Gertner (1993), Dana (1994), or, more recently, Janssen et al. (2017), Garcia and Shelegia (2018) and Lauermann et al. (2018).

2 Model

The objective of our paper is to study the role of biased consumer beliefs in search markets. To bring out this role most clearly, we extend the basic search model due to Diamond (1971) by allowing for general consumer beliefs.[11] The market consists of N firms indexed by $n = 1, \dots, N$ which sell the same homogeneous product that they can produce at zero marginal cost. There is a unit mass of consumers. All consumers have the same valuation for the product and consume at most one unit of it. Let ω denote the surplus from trade. At the outset, each firm chooses a utility level u_n to offer to consumers. When trade takes place, the consumer obtains u_n, and the firm's profit is $\pi_n = \omega - u_n$.[12] A mixed strategy for firm n is thus a cdf κ_n that describes the distribution of consumer utility. Let CDF be the set of all cdfs over $\mathbb{R}$.

Ex ante, consumers do not know the utility offered by firms but can engage in costly, sequential search, undirected and with recall. When visiting firm n, a consumer perfectly observes the utility u_n offered by firm n. In addition, there is an outside option which supplies zero utility to the consumer. Search entails a marginal search cost $s > 0$, except for the first search which is free.[13]

The novelty of our article is that we allow a fraction $\gamma \in (0, 1)$ of biased consumers to have beliefs that are inconsistent with what the market actually supplies. Given a firm's true utility distribution κ, a biased consumer believes the utility distribution to be $\beta_B(\kappa)$, where the belief function

$$\beta_B : CDF \to CDF \tag{1}$$

captures the extent to which beliefs are biased. We assume that the remaining fraction $1 - \gamma$ of consumers is "standard" and has correct beliefs which are characterized by the belief function β_0 with $\beta_0(\kappa) = \kappa$.

Benchmarks: Diamond paradox and Stahl equilibrium

We begin by reviewing two well-known benchmarks that play an important role in our analysis.

The first benchmark, due to Diamond (1971), is the rational consumer benchmark when $\gamma = 0$. In this case, the famous Diamond paradox occurs: for any level of search costs, in equilibrium, each firm chooses its monopoly strategy, and each consumer buys from the first firm she visits. In our setup, this means that each firm supplies zero utility with probability 1 and obtains the

[11] In the Discussion section, we discuss how our insights carry over to richer search models.

[12] Equivalently, each firm charges a price p_n, and when a consumer buys, she gets $\omega - p_n = u_n$, and the firm gets $\pi_n = p_n = \omega - u_n$.

[13] This common assumption ensures that there are equilibria with trade. For a discussion, see Stiglitz (1979).

entire surplus from trade.[14] Intuitively, anticipating (correctly) that all firms make the same offer, search is pointless from a consumer's point of view, which, in turn, provides firms with monopoly power.

The second benchmark, due to Stahl (1989), obtains when instead of biased consumers, there is a fraction γ of so-called shoppers: standard consumers with zero search costs who visit all firms. Then there is no equilibrium in pure firm strategies, and in the unique symmetric equilibrium, firms offer a utility distribution with interval support. In particular, in equilibrium, consumers with positive search costs find all offers acceptable and thus accept the offer from the first firm they visit. By contrast, shoppers visit all firms and pick the offer with the highest utility among all firms. Intuitively, the equilibrium distribution of offers balances the trade-off between attracting shoppers with high offers and "exploiting" non-shoppers with low offers.

Example: cursed beliefs
Before we begin with the full analysis of biased consumer beliefs, we provide a microfoundation for a specific belief function and illustrate its role in shaping market outcomes. Consider a model where each firm is equally likely to produce a product of high quality q_H at marginal cost c_H, and a product of low quality $q_L < q_H$ at marginal cost $c_L < c_H$, and where all firms generate the same surplus $\omega = q_H - c_H = q_L - c_L$. Let $v = q_H - q_L$. At the outset, each firm with quality q_θ once and for all sets a price p_θ, $\theta \in \{H, L\}$. A consumer's utility from consuming quality q at price p is $u = q - p$, and when a consumer visits a firm, she uncovers both the firm's quality and price.

When all consumers are standard, the aforementioned Diamond paradox occurs (compare Footnote 14) and the only equilibrium outcome is that each firm charges its monopoly price which is equal to the consumer's willingness to pay: $p_\theta = q_\theta$.

We argue now that the Diamond paradox may break down if we introduce a fraction of biased consumers who hold cursed beliefs in the sense of Eyster and Rabin (2005). This means that a biased consumer, while correctly understanding the marginal distribution of price and quality in the market, fails to make the correct connection between price and quality, effectively underestimating the correlation between them. Specifically, a biased consumer believes that with probability $\chi \in [0, 1/4]$ a product is of high quality and offered at the price of a low quality product and vice versa.[15] The parameter χ captures the extent to which biased consumers are cursed. Table

[14] The Diamond paradox applies more generally when all consumers face possibly heterogeneous, strictly positive search costs (bounded away from zero), and either for each consumer all products offer the same utility, or each product offers the same utility to all consumers. Under our unit demand assumption consumers are pushed to their outside option under the monopoly outcome but this is not substantial for our results.

[15] More precisely, in the formulation of Eyster and Rabin (2005), a biased consumer believes that, with probability $1 - \xi$, a firm of quality q_θ choose its true price $p_\theta = q_\theta$, and with probability ξ, it randomizes between p_H and p_L

2 illustrates the true and the cursed joint distribution of price and quality under the Diamond outcome.

	q_H	q_L
$p_H = q_H$	1/2	0
$p_L = q_L$	0	1/2

	q_H	q_L
$p_H = q_H$	$1/2-\chi$	χ
$p_L = q_L$	χ	$1/2-\chi$

Figure 1: True (left) and cursed (right) joint price quality distribution under the Diamond outcome.

If search costs are not too large, their cursed beliefs induce biased consumers to visit more than one firm, given the prices under the Diamond outcome. Indeed, by assumption, the first firm supplies the consumer with zero utility. On the other hand, the consumer could adopt the (not necessarily optimal) search strategy to visit one more firm and buy from that firm among the two which offers more utility. A biased consumer wrongly believes that there is a chance of χ that the next firm she will visit offers a *golden egg*, that is, high quality at the low price, and thus supplies the fictitious utility $q_H - p_L = \upsilon$. Therefore, she believes that visiting one additional firm yields the additional utility

$$\chi \upsilon - s. \tag{2}$$

Hence, if s is sufficiently small, she believes that visiting more than one firm is better than buying at the first firm, and so the Diamond outcome breaks down if s is sufficiently small.

We now illustrate how the example can be viewed as a special case of our general setting. Notice first that all products generate the same surplus from trade as in our general setup and consider again the utility distribution induced by the Diamond outcome in Table 2. A consumer's utility is $u = 0$ if an outcome on the diagonal obtains, is $u = \upsilon$ if the bottom-left outcome obtains, and is $u = -\upsilon$ if the top-right outcome obtains. Thus, while the true utility distribution places mass 1 on the utility level $u = 0$, a biased consumer's belief is a three-point distribution that takes away the mass 2χ from the true utility $u = 0$ and spreads it symmetrically on the fictitious utility levels $0 - \upsilon$ and $0 + \upsilon$.

This property carries over to the case when firms adopt any (mixed) pricing strategies which result in utility distributions, κ_θ, which are symmetric in the utility dimension across qualities: $\kappa_H = \kappa_L \equiv \kappa$. This is the case in symmetric equilibrium where a firm's utility offer only depends on the surplus from trade, and thus not on its quality. Then any firm offers the true utility distribution

according to the (true) marginal price distribution, resulting in $\chi = 1/4 \cdot \xi$. In general, if the probability for high quality is not 1/2 but σ, then $\chi = \sigma \cdot (1-\sigma) \cdot \xi$. We omit the details.

κ, and it is not hard to see that now the cursed belief takes away the mass 2χ from any utility u in the support of κ and spreads it symmetrically on $u-\upsilon$ and $u+\upsilon$. Formally, this results in the belief function[16]

$$\beta_B(\kappa)(u) \equiv \chi \cdot \kappa(u+\upsilon) + (1-2\chi) \cdot \kappa(u) + \chi \cdot \kappa(u-\upsilon), \quad \text{with } \chi \in [0, 1/4]. \tag{3}$$

In the remainder of the article, we extend this example to general belief specifications and characterize precisely when the Diamond outcome does or does not break down and derive the equilibrium outcomes that arise instead.[17] For our welfare analysis in Section 5, we return to specification (3).

3 Equilibrium definition and equilibrium structure

We restrict attention to symmetric equilibria in which all firms choose the same utility distribution. In this case, both standard and biased consumers believe that they sample utility offers iid. Therefore, it follows from Kohn and Shavell (1974) that a consumer's optimal search rule is myopic and fully characterized by a reservation utility. Thus, while the set of consumer search strategies is, in principle, very large, when restricting attention to symmetric equilibria, we can equate a search strategy with a reservation utility.

More precisely, the reservation utility is the minimal utility that a consumer's best option must supply so that she stops her search and takes this option (possibly her outside option of 0). Otherwise she continues her search. If the consumer has visited all firms, she returns to her best option. Let $U(\beta)$ denote a consumer's reservation utility given her belief $\beta \in CDF$. From Kohn and Shavell (1974), $U(\beta)$ is given as the unique solution to the equation

$$U(\beta) = R(U(\beta), \beta), \tag{4}$$

[16]The belief function in (3) is a mean-preserving spread of κ. In particular, biased consumers overestimate the dispersion of offers in the market which is consistent with some empirical evidence (see **?**). Specification (3) also applies in settings where a product's surplus is the aggregate of two attributes, which are inherently negatively correlated, for example, a car with the attributes size and fuel consumption. A biased consumer displaying correlation neglect, might wrongly believe that with some probability (χ), she may encounter a product that excels in both (resp. neither) attributes, entailing a utility gain of υ (resp. $-\upsilon$). For evidence on correlation neglect in general, see Enke and Zimmermann (2019).

[17]To be sure, biased beliefs are not the only reason why the Diamond Paradox may fail. For example, this also happens when products are differentiated (Wolinsky (1986)), consumers search simultaneously (Salop and Stiglitz (1977), Burdett and Judd (1983)) or some consumers, by assumption, visit several firms (Varian (1980), Stahl (1989)).

where

$$R(U,\beta) \equiv -s + \int \max\{u,U\}\, d\beta(u). \tag{5}$$

Note that the reservation value $U(\beta) = U(\beta, s)$ also depends on s. To save notation, we omit the dependency on s unless necessary.

We can now define a symmetric equilibrium with biased consumer beliefs.[18,19]

Definition 1 *A symmetric equilibrium (with biased consumer beliefs) is a triple* (κ^*, U_0^*, U_B^*) *such that*
(i) κ^* *maximizes a firm's profit, given that all other firms adopt* κ^* *and consumers search according to* U_0^* *and* U_B^**;*
(ii) U_0^* *and* U_B^* *are optimal reservation values, given standard and biased consumer beliefs induced by the true utility distribution* κ^**:*

$$U_0^* = U(\beta_0(\kappa^*)) \quad \textit{and} \quad U_B^* = U(\beta_B(\kappa^*)). \tag{6}$$

Equilibrium structure
We begin by providing necessary conditions that pin down the types of symmetric equilibria that can arise. Below, we then characterize which equilibrium arises as a function of β_B. To state the lemmata, we denote the bounds of the support of a utility distribution κ by

$$\underline{u} = \min[Supp(\kappa)], \quad \overline{u} = \max[Supp(\kappa)]. \tag{7}$$

We indicate the bounds by a star if they belong to an equilibrium utility distribution, and refer to the utility level $\underline{u}$ and $\overline{u}$ as the worst respectively best offer in the market. Our first lemma provides necessary conditions for worst and best offers in equilibrium. We say, an equilibrium displays utility dispersion if $\underline{u}^* < \overline{u}^*$.

Lemma 1 *In any symmetric equilibrium with utility dispersion, we have:*
(i) $\underline{u}^* = \max\{0, \min\{U_0^*, U_B^*\}\}$*,*
(ii) $\overline{u}^* \leq \max\{0, \max\{U_0^*, U_B^*\}\}$*.*

To understand the lemma, note that, in any equilibrium, the minimal utility that induces a visiting consumer of type $\tau \in \{0, B\}$ to stop and buy is

$$\max\{0, U_\tau^*\}. \tag{8}$$

[18]Our equilibrium definition entails that a consumer's belief is passive and does not change when she observes an offer which is an impossible event given her beliefs about the true distribution of offers. See Janssen and Shelegia (2020) for a recent article that departs from the common assumption of passive beliefs.

[19]To derive and characterize equilibrium, we assume that a consumer accepts an offer if she is indifferent between accepting and continuing her search, and that she never chooses her outside option when she is indifferent.

The reason is that only when a firm offers $u \geq U_\tau^*$, the consumer stops. Moreover, all other firms that she might have visited before must supply less utility than U_τ^*, as she otherwise would have stopped before. Therefore, if the firm offers, in addition, more utility than her outside option, then its product is the consumer's best offer and she buys it upon stopping.

With this in mind, part (i) says that the worst equilibrium offer $\underline{u}^*$ is the minimal utility that induces at least one consumer type to stop and buy. Intuitively, offering more utility is not optimal because a firm that makes the worst offer is not competing for returning consumers. On the other hand, offering less utility generates no demand, because no consumer would stop and buy or return. Part (i) implies that in any equilibrium with utility dispersion, at least one consumer type stops and buys at the first firm she visits.

Part (ii) shows that the best equilibrium offer $\overline{u}^*$ does not exceed the minimal utility that induces both consumer types to stop and buy: $\overline{u}^* \leq \max\{0, \max\{U_0^*, U_B^*\}\}$. The reason is that offering more utility is never profit maximizing, because the firm could offer marginally less utility so as to increase its share of the surplus without losing demand, as any visiting consumer stops and buys if it offers more utility than $\max\{0, \max\{U_0^*, U_B^*\}\}$ from (8).

Building on Lemma 1, the next lemma shows that there are only three types of candidates for symmetric equilibria.

Lemma 2 *In a symmetric equilibrium only the following three constellations can arise:*
(i) A Diamond-type equilibrium*, that is, all firms offer the same single utility level, and both standard and biased consumers buy from the first firm they visit. Formally, there is $v^* \geq 0$ so that*

$$\underline{u}^* = \overline{u}^* = v^*, \quad \max\{U_0^*, U_B^*\} \leq v^*. \tag{9}$$

(ii) A Stahl-type equilibrium*, that is, the equilibrium features "smooth" utility dispersion, biased consumers visit all firms and then return to the highest offer, whereas standard consumers buy from the first firm they visit. Formally,*

$$Supp(\kappa^*) = [\underline{u}^*, \overline{u}^*], \quad U_0^* \leq \underline{u}^* < \overline{u}^* \leq U_B^*. \tag{10}$$

(iii) A penny sales equilibrium*, that is, the equilibrium features utility dispersion with a gap at the top, biased consumers stop searching only when offered the top offer $\overline{u}^*$, and otherwise, visit all firms and then return to the best offer, whereas standard consumers buy from the first firm they visit. Formally,*

$$Supp(\kappa^*) = [\underline{u}^*, \hat{u}^*] \cup \{\overline{u}^*\}, \quad U_0^* \leq \underline{u}^* < \hat{u}^* < \overline{u}^* = U_B^*. \tag{11}$$

Before we explain lemma, notice that in any equilibrium, $U_0^* < \overline{u}^*$. Intuitively, the best that a searching standard consumer can hope for is that the next firm offers $\overline{u}^*$ (with probability one) so that she is always willing to accept any current offer which supplies more than $\overline{u}^* - s$, as visiting another firm entails search costs $s > 0$. Formally, in any equilibrium, $U_0^* < \overline{u}^*$, because

$$U(\beta_0(\kappa)) \leq \overline{u} - s \quad \text{for all } \kappa \text{ and } s. \tag{12}$$

As a consequence, in any equilibrium, a standard consumer stops and buys if a firm offers $\overline{u}^*$.

To see the intuition behind Lemma 2, consider first the case that $\overline{u}^* > U_B^*$, that is, the best equilibrium offer also supplies enough utility so as to induce a biased consumer to stop and buy. In this case, an equilibrium arises as described in (i) with $v^* = 0$. In other words, the Diamond outcome obtains. The reason why all firms offer zero utility when $\overline{u}^* > U_B^*$ and $\overline{u}^* > U_0^*$ is that otherwise, when $\overline{u}^* > 0$, a firm which offers $\overline{u}^*$ could do better by offering marginally less utility. This would increase its share of the surplus and not affect its demand, because as $\overline{u}^* > U_B^*$ and $\overline{u}^* > U_0^*$, all consumers would still stop and buy.

In the case that $\overline{u}^* < U_B^*$, we have an equilibrium as described in (ii). Recall that $\overline{u}^* < U_B^*$ means that biased consumers never encounter an offer that induces them to stop searching. Hence, they visit all firms and return to the one which offers the highest utility, and it follows from standard arguments that κ^* has no mass points, and its support has no gaps. Because there is utility dispersion, part (i) of Lemma 1 together with $\overline{u}^* < U_B^*$ entails that $U_0^* \leq \underline{u}^*$ which means that standard consumers buy from the first firm they visit. Effectively, in equilibrium, the two consumer groups behave like the consumer groups in the benchmark due to Stahl (1989): biased consumers behave like shoppers, and standard consumers behave like consumers with positive search costs. Intuitively, both models feature therefore the same equilibrium structure with smooth utility dispersion. Accordingly, we refer to this type of equilibrium as a Stahl-type equilibrium.

Finally, consider the case that $\overline{u}^* = U_B^*$. In this case, an equilibrium with a mass point at $\overline{u}^*$ as described in (iii) may occur.[20] To highlight that in this equilibrium, the best offer $\overline{u}^*$ provides discontinuously more utility than any other equilibrium offer, we refer to $\overline{u}^*$ as a penny sales offer and to this equilibrium type as a penny sales equilibrium. To understand why a mass point can occur in equilibrium, note that when a firm makes a penny sales offer $\overline{u}^*$, a visiting biased consumer stops her search and accepts the firm's offer. Hence, on the equilibrium path, a biased consumer never visits two firms that both make penny sales offers. Hence, unlike in a Stahl-type equilibrium with shoppers, these firms do not compete with each other for returning consumers in the sense that that offering marginally more utility does not attract discontinuously more consumers, despite

[20] Also a Stahl-type equilibrium with $\overline{u}^* = U_B^*$ may occur, in which case there is no mass point at $\overline{u}^*$. There may also be a Diamond-type equilibrium with $\overline{u}^* = U_B^*$ in which firms offer $\overline{u}^* \geq 0$ with probability one.

the mass point. On the other hand, it is not profitable for a firm to offer marginally less utility, as otherwise a biased consumer would not stop anymore, but rather continue her search and not return if she encounters a firm with a penny sales offer (which occurs with positive probability).

To understand why the equilibrium utility distribution in (iii) also features an interval $[\underline{u}^*, \hat{u}^*]$, corresponding to smooth utility dispersion, note that with positive probability no firm in the market makes a penny sales offer. In this case, a biased consumer behaves like a shopper, visits all firms and returns to the one which offers the highest utility. Intuitively, apart from penny sales, the utility distribution can hence have no gaps and mass points, and features a region with smooth utility dispersion as described in (iii).

Lemma 2 makes clear that a completely novel equilibrium type may emerge when some consumers are biased. More specifically, a penny sales equilibrium, that is, an equilibrium with a mass point at $\overline{u}^*$, is a unique feature of our environment with biased beliefs in the following sense: In general, utility dispersion with a mass point at the top of the utility distribution may only occur if there is a consumer of type τ who satisfies $U^*_\tau = \overline{u}^*$.[21] Now, a consumer's reservation utility is bounded from above by $\overline{u}^*_\tau - s$ where $\overline{u}^*_\tau$ denotes the highest utility that she expects firms to offer, as we explain right after Lemma 2. Therefore, $U^*_\tau = \overline{u}^*$ can only occur if some consumers hold biased beliefs which attach positive probability to the event that there are golden eggs, that is, offers which are better than the truly best offer $\overline{u}^*$.

Equilibrium utility distribution

Next, we use Lemma 1 and 2 to identify the equilibrium utility distributions for each equilibrium type. To describe the distributions, we introduce a family of cdfs κ_ζ indexed by a parameter $\zeta \in [0,1]$ (which will correspond to the probability of a mass point at the upper support bound). Define the bounds of the support of κ_ζ as

$$\underline{u}_\zeta = \max\left\{0, \omega - \frac{s}{\mu(\zeta)}\right\}, \quad \hat{u}_\zeta = (1-\phi(\zeta))\,\omega + \phi(\zeta)\,\underline{u}_\zeta, \qquad \overline{u}_\zeta = (1-\rho(\zeta))\,\omega + \rho(\zeta)\,\underline{u}_\zeta, \tag{13}$$

where $\rho, \phi, \mu : [0,1] \to [0,1]$ are strictly increasing functions for which we provide closed form expressions in Appendix B. Let $\rho_0 = \rho(0)$, $\mu_0 = \mu(0)$, and define κ_ζ as the cdf that is continuous on $[\underline{u}_\zeta, \hat{u}_\zeta]$ with

$$\kappa_\zeta(u) \equiv \left[\frac{\rho_0}{1-\rho_0} \cdot \left(\frac{\omega - \underline{u}_\zeta}{\omega - u} - 1\right)\right]^{\frac{1}{N-1}} \quad \text{and has a mass point of mass } \zeta \in [0,1] \text{ at } \overline{u}_\zeta. \tag{14}$$

[21] To see this, if the reservation utility of all consumers is below $\overline{u}^*$, then a firm which offers $\overline{u}^*$ could offer marginally less utility so as to increase its mark-up without losing demand. On the other hand, if $U^*_\tau > \overline{u}^*$ for some consumers, then these consumers visit all firms and return to the best offer, meaning that there cannot be a mass point at the top as otherwise, a firm could offer marginally less utility and discontinuously increase its market share.

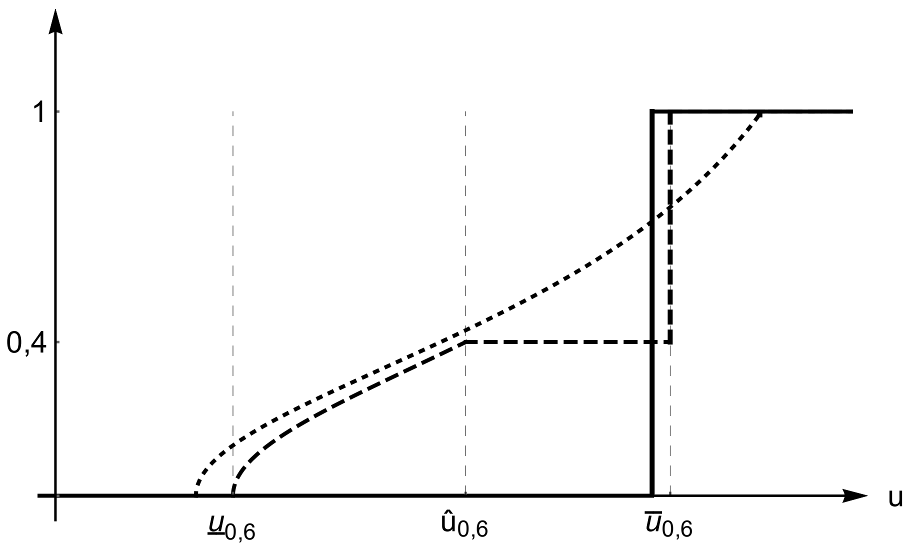

Figure 2: *The utility distributions κ_0 (dotted), $\kappa_{0,6}$ (dashed), and κ_1 (solid). The bounds of the support of $\kappa_{0,6}$ are indicated by $\underline{u}_{0,6}$, $\hat{u}_{0,6}$ and $\overline{u}_{0,6}$.*

Figure 8 illustrates κ_ζ for the values $\zeta \in \{0, 0.6, 1\}$. κ_0 is identical to the utility distribution in Stahl (1989). Moreover, κ_1 is the degenerate distribution with mass 1 on $\overline{u}_1$.

The next lemma establishes for each equilibrium type a necessary and sufficient condition for β_B for the equilibrium to exist. Here, δ_v denotes the (unit step) cdf which places all mass on the utility level v.

Lemma 3 *(i) There is a Diamond-type equilibrium with utility level v^* if and only if*

(a) $v^* = 0$ and $0 \geq U(\beta_B(\delta_0))$ and $\kappa^* = \delta_0$; or

(b) $v^* \in (0, \overline{u}_1]$ and $v^* = U(\beta_B(\delta_{v^*}))$ and $\kappa^* = \delta_{v^*}$.

(ii) There is a Stahl-type equilibrium if and only if $\overline{u}_0 \leq U(\beta_B(\kappa_0))$ and $\kappa^* = \kappa_0$.

(iii) There is a penny sales equilibrium if and only if there is $\zeta^* \in (0,1)$ so that $\overline{u}_{\zeta^*} = U(\beta_B(\kappa_{\zeta^*}))$ and $\kappa^* = \kappa_{\zeta^*}$.

To illustrate the logic behind Lemma 3, consider part (iii). By Lemma 2, in a penny sales equilibrium a firm is indifferent between all offers in $[\underline{u}^*, \hat{u}^*] \cup \{\overline{u}^*\}$ given the other firms adopt κ^* and consumers' search behaviour as described in the right part of (11). From this indifference condition, it can be deduced that the candidate equilibrium utility distribution κ^* takes the form (13) and (14) where the value of the lower support bound $\underline{u}^*$ is still a free variable (and not yet pinned down to be $\underline{u}_\zeta$). We then calculate the standard consumer's reservation utility $U_0^* = U(\kappa^*)$ as a function of $\underline{u}^*$ (remarkably, in closed form), and then determine $\underline{u}^*$ so that it satisfies the relationships between $U_0^* = U_0^*(\underline{u}^*)$ and $\underline{u}^*$ in the equilibrium characterization in Lemma 1, (i), and Lemma 2, (iii). This uniquely pins down the value of $\underline{u}^*$ as $\underline{u}_\zeta$ as defined in (13), and thus, $\kappa^* = \kappa_\zeta$, however, with ζ still a free variable. Finally, what pins down ζ^* is that in a penny sales equilibrium, given ζ^*, the biased consumer's equilibrium reservation utility $U_B^* = U(\beta_B(\kappa^*))$ must be equal to $\overline{u}^* = \overline{u}_{\zeta^*}$ by Lemma 2, and hence, a penny sales equilibrium exists if and only if there is such $\zeta^* \in (0,1)$.

The significance of Lemma 3 is that it explicitly constructs the equilibrium utility distribution κ^* and reduces the question whether a particular equilibrium type exists to a comparison between the biased consumer's reservation utility $U_B^* = U(\beta_B(\kappa^*))$ with the upper support bound $\overline{u}^*$ that is induced by κ^*.

4 Biased beliefs and equilibrium outcomes

In this section, we characterize the properties of β_B that determine which of the possible equilibrium outcomes obtains. Throughout, we assume that β_B satisfies the following regularity properties:

Definition 2 *Let* $\beta : CDF \to CDF$.
(A) β *is continuous if for all* κ *and* s *and all sequences* κ_n *and* s_n, $n = 1, 2, \ldots$*:*

$$U(\kappa_n, s_n) \to U(\kappa, s) \quad \Rightarrow \quad U(\beta(\kappa_n), s_n) \to U(\beta(\kappa), s). \tag{15}$$

(B) β *is monotone if for all* κ, κ'*:*

$$\kappa \ \text{ first order stochastically dominates (fosd) } \ \kappa' \quad \Rightarrow \quad \beta(\kappa) \ \text{ fosd } \ \beta(\kappa'), \tag{16}$$

with $U(\beta(\kappa)) > U(\beta(\kappa'))$ *whenever* $U(\kappa) > U(\kappa')$.

(C) For κ *and* $\Delta \in \mathbb{R}$*, let* κ_Δ *be the cdf that results from shifting* κ *by* Δ*:* $\kappa_\Delta(u) = \kappa(u - \Delta)$*. Then* β *is additive if for all* κ, Δ*:*

$$U(\beta(\kappa_\Delta)) = U(\beta(\kappa)) + \Delta. \tag{17}$$

Assumption 1 β_B *satisfies (A) to (C).*

Conditions (A) to (C) are natural regularity conditions that are satisfied for β_0 and which we assume to carry over to β_B.[22] We state conditions (A) to (C) in terms of reservation utilities, as

[22](A) and (B) hold trivially for β_0. To show (C), note that for U we have $R(U, \kappa_\Delta) = -s + \int \max\{u + \Delta, U\}\, d\kappa(u)$. Hence,

$$R(U(\kappa) + \Delta, \kappa_\Delta) = -s + \int \max\{u + \Delta, U(\kappa) + \Delta\}\, d\kappa(u) = U(\kappa) + \Delta, \tag{18}$$

where the last step follows from the definition of the reservation value $U(\kappa)$. Therefore, $U(\kappa_\Delta) = U(\kappa) + \Delta$.

this is how they are applied in the proofs for the results below. A sufficient condition for (A) is the joint requirement that (i) if κ_n converges to κ in distribution, then also $\beta_B(\kappa_n)$ converges to $\beta_B(\kappa)$ in distribution, and (ii) β_B is bounded in the sense that β_B maps sequences κ_n with $\max[Supp(\kappa_n)] < M$ for all n and some M into sequences $\beta_B(\kappa_n)$ with $\max[Supp(\beta_B(\kappa_n))] < \tilde{M}$ for all n and some $\tilde{M}$. A sufficient condition for (C) is that $\beta_B(\kappa_\Delta) = (\beta_B(\kappa))_\Delta$ (see footnote 22). That is, how a biased consumer forms her belief does not depend on the cardinal level of utility.

A class of belief functions that satisfies Assumption 1 is given by

$$\beta_B(\kappa)(u) = \int G(u - u')d\kappa(u'), \tag{19}$$

where G is an arbitrary cdf. If the mean of G is zero, then β_B is a mean-preserving spread of κ. This class includes the belief function induced by cursed beliefs given in (3), or any belief function that shifts the true distribution by some Δ to either higher or lower utility levels.[23] For example, a shift to higher utility levels occurs if a consumer systematically fails to account for price increases (e.g. inflation) and wrongly believes that he might find the product offered at the same price the last time he purchased it. A shift to lower utility levels occurs if a consumer has unrealistically low expectations, for example, because he fails to account for quality improvements through technological progress.

A further example that satisfies Assumption 1 is

$$\beta_B(\kappa)(u) = \kappa\Big(\alpha \cdot \big(u - \mathbf{E}(u)\big) + \mathbf{E}(u)\Big), \tag{20}$$

where $\alpha > 1$ is a parameter and $\mathbf{E}(u)$ is the mean of κ. β_B has the same mean as κ, but, because $\alpha > 1$, is more concentrated towards the mean. Thus, a biased consumer underestimates the degree to which true offers are dispersed. For example, as Grubb (2015) argues, consumers may underestimate the dispersion of prices, because their predictions about the next price quote suffer from overprecision or belief in the "law of small numbers".

An example ruled out by Assumption 1 is a constant belief function that assigns the same belief to any κ, because additivity is then violated.

Our next result establishes that the Diamond paradox is the unique equilibrium outcome if biased consumers are *pessimistic* in the sense that they (weakly) underestimate the value of the best offer in the market. Formally:

Definition 3 *$\beta : CDF \to CDF$ is pessimistic if for all κ:*

$$\max[Supp(\beta(\kappa))] \leq \max[Supp(\kappa)]. \tag{21}$$

[23] Cursed beliefs correspond to $G = \chi \cdot \delta_{-\upsilon} + (1 - 2\chi) \cdot \delta_0 + \chi \cdot \delta_\upsilon$. A belief function that shifts the true distribution by some Δ corresponds to $G = \delta_\Delta$.

Proposition 1 *If β_B is pessimistic, then the Diamond-type equilibrium (with utility level 0) is the unique symmetric equilibrium.*

The intuition behind this result is reminiscent of Diamond (1971). Because biased consumers are pessimistic, in any equilibrium, both biased and standard consumers believe that the best offer in the market supplies weakly less utility than $\overline{u}^*$. Hence, their reservation utilities are bounded from above by $\overline{u}^* - s$, as we explain after Lemma 2. From this and Lemma 1, we must have $\overline{u}^* = 0$ in any equilibrium, intuitively, as otherwise, a firm could offer slightly less utility than $\overline{u}^*$ without losing demand, because any visiting consumer is willing to stop and buy if it offers more than $\overline{u}^* - s$.

Notice that the rational consumer belief β_0 is pessimistic so that for $\beta_B = \beta_0$, Proposition 1 is consistent with the rational consumer benchmark. Moreover, any belief function given by (20) is pessimistic.

Our next result characterizes equilibrium outcomes if biased consumers' beliefs are not pessimistic (and a mild parameter restriction holds). Note that any belief function given by (19) which satisfies $G(0) < 1$, and especially the belief function (3) induced by cursed beliefs is non-pessimistic. We show that in this case, the three equilibrium types segment the search cost space, and almost everywhere the equilibrium is unique.

Proposition 2 *Suppose $\rho_0 < \mu_0$.*[24] *If β_B is non-pessimistic, then there are unique $0 < s_0 < s_1$ so that:*
(i) There is a Diamond-type equilibrium with utility level 0 if and only if $s \geq s_1$. Moreover, for all $v \in (0, \overline{u}_1]$, there is a Diamond-type equilibrium with utility level v if and only if $s = s_1$.
(ii) There is a Stahl-type equilibrium if and only if $s \leq s_0$.
(iii) There is a penny sales equilibrium, which is unique, if and only if $s \in (s_0, s_1)$.

Together, Propositions 1 and 2 provide an essentially complete picture of the possible market outcomes with biased beliefs. In particular, the propositions identify the class of pessimistic belief functions as the class of beliefs for which the Diamond outcome is the unique equilibrium for all levels of search costs. Moreover, a penny sales equilibrium obtains if and only if some consumers have non-pessimistic beliefs and search costs are in an intermediate range $s \in (s_0, s_1)$. Finally, for

[24] In Appendix B, we provide closed form expressions for μ_0 and ρ_0 which depend on N and γ only. Numerical calculations prove that $\rho_0 < \mu_0$ if and only if $\gamma > \gamma_{crit}(N)$ where the function γ_{crit} is decreasing in N, $\gamma_{crit}(2) \approx 0.51$, and $\gamma_{crit}(\infty) = 0$. Figure 8 in Appendix B plots $\gamma_{crit}(N)$. If $\rho_0 < \mu_0$ does not hold, part (i) and the "if-statement" of part (ii) of Proposition 2 remain true. As to part (iii), we can still show that the three equilibrium types segment the search cost space, each type exists and the equilibrium is almost everywhere unique. However, we cannot rule out that between s_0 and s_1 penny sales and Stahl equilibria alternate as s increases.

non-pessimistic beliefs and small search costs $s \leq s_0$, the equilibrium outcome is indistinguishable from the one in Stahl (1989), but the normative implications are rather different. While a shopper in Stahl (1989) loves shopping and has zero search costs, a biased consumer in our setting expends excessive search costs for visiting all firms due to her unrealistic expectations.

To understand the driving forces behind the proposition, observe that by Lemma 3, whether a particular equilibrium type exists depends only on the difference between the best offer in the market and the biased consumer's reservation utility. In fact, under Assumption 1, Lemma 3 can be re-phrased as follows:[25]

(i') There is a Diamond-type equilibrium if and only if $\overline{u}_1 - U(\beta_B(\kappa_1)) \geq 0$,
(i") There is a Stahl-type equilibrium if and only if $\overline{u}_0 - U(\beta_B(\kappa_0)) \leq 0$, and
(i"') There is a penny sales equilibrium if and only if there is $\zeta^* \in (0,1)$ with $\overline{u}_{\zeta^*} - U(\beta_B(\kappa_{\zeta^*})) = 0$.

Figure 3 illustrates the equilibrium outcome for the case when biased consumers hold cursed beliefs as in (3). Observe that the equilibrium values of the best offer $\overline{u}^*$ and the biased consumer's reservation value $U_B^* = U(\beta_B(\kappa^*))$ satisfy in each regime the respective equilibrium condition from above. In the Diamond- and Stahl regime the inequalities (i') and (i") hold strict. Indeed, to show that the three equilibrium types segment the search cost space, the proof of Proposition 2 establishes and utilizes single-properties for $\overline{u}_\zeta - U(\beta_B(\kappa_\zeta))$.[26] We next give an intuition.

For sufficiently large search costs, it becomes too expensive even for unrealistically optimistic biased consumers to engage in any search beyond the first firm. This grants firms full monopoly power, and the Diamond outcome obtains. The critical value s_1 is the level of search cost where a biased consumer is just indifferent between buying from the first firm and engaging in search, given the Diamond outcome.

Reversely, when search costs are very small, utility comparisons are very cheap for consumers. Indeed, in the limit as search costs vanish, only the Bertrand outcome, where firms offer the entire surplus to consumers, can obtain in equilibrium. Formally, it can be seen from (13) that $\kappa_\zeta \rightarrow \delta_\omega$ for all ζ as $s \rightarrow 0$. Now, as we show in the proof of Proposition 2, non-pessimism of β_B and Assumption 1 guarantee that a consumer overestimates the value of the best offer whenever she is confronted with a distribution involving a single atom, that is, $\max[Supp(\beta_B(\delta_\omega))] > \max[Supp(\delta_\omega)]$. This means that in any equilibrium, as search costs van-

[25] As we show in (55), additivity of β_B implies that $v - U(\beta_B(\delta_v))$ is independent of v. Part (i') follows then from the fact that for any v, we have that $v \geq U(\beta_B(\delta_v))$ is equivalent to $\overline{u}_1 \geq U(\beta_B(\kappa_1))$, because $\kappa_1 = \delta_{\overline{u}_1}$.

[26] Specifically, we show:

(i) $\overline{u}_0 - U(\beta_B(\kappa_0))$ is strictly increasing in s and is zero at $s = s_0$.
(ii) $\overline{u}_1 - U(\beta_B(\kappa_1))$ is strictly increasing in s and is zero at $s = s_1$.
(iii) $\overline{u}_\zeta - U(\beta_B(\kappa_\zeta))$ is strictly decreasing in ζ for all s.

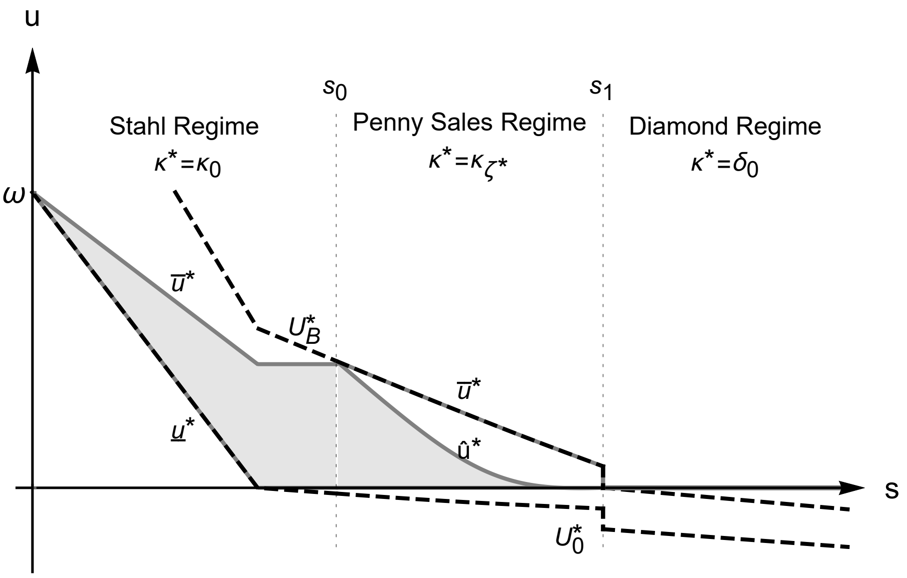

Figure 3: *The equilibrium reservation values U_0^* and U_B^* (dashed) as well as the support of the equilibrium utility distribution κ^* (gray area) as functions of s. In the penny sales regime, κ^* features a mass point of size ζ^* at $\overline{u}^*$ and its support is not connected.*

ish, a biased consumer believes that firms offer with positive probability golden eggs that actually supply more utility than the entire surplus. As a consequence, when search is essentially for free, a biased consumer is not willing to settle for less than a golden egg, and because in reality there is none, she ends up searching all firms in the market. As a result, biased consumers behave in any equilibrium like shoppers, and thus (only) a Stahl-type equilibrium obtains for sufficiently small search costs.

Indeed, as long as search costs stay in the range $s \in (0, s_0)$, biased consumers behave like shoppers and a Stahl-type equilibrium obtains, because it remains too costly for a firm to make an offer which induces a biased consumer to stop and buy, as this would require to offer more or close to the entire surplus. This changes at the critical level s_0. At this point, a biased consumer is indifferent between continuing search and accepting the best offer $\overline{u}_0$ that can obtain under the utility distribution κ_0 in the Stahl-type equilibrium. Beyond that point, search becomes too costly to support the Stahl-type equilibrium. For $s \in (s_0, s_1)$, firms make penny sales offers $\overline{u}^*$ with probability $\zeta^* \in (0,1)$ (notice the gap in the support in Figure 3) so as to keep a biased consumer indifferent between accepting an offer $\overline{u}^*$ and continuing their search. That is, ζ^* adjusts so that $\overline{u}^*_\zeta - U(\beta_B(\kappa^*_\zeta)) = 0$. In this regime, as search costs increase, all else equal, biased consumers become less selective and firms have to offer less utility so as to induce them to stop and buy. At $s = s_1$, a penny sale is offered with probability $\zeta^* = 1$, and all firms make the same single offer with probability 1, just as under the Diamond outcome. At this point, we have $\overline{u}_1 - U(\beta_B(\kappa_1)) = 0$, and the transition to the Diamond regime occurs.

5 Welfare Analysis

In this section, we return to specification (3) with cursed beliefs and study how changes in model parameters affect market outcomes and consumer welfare. In our model, biased consumer welfare is given as their expected utility given correct expectations. We show analytically that an intermediate level of cursedness, χ, maximizes both standard and biased consumers' welfare. We also provide numerical welfare comparative statics results with respect to the share of biased consumers, γ, and the number of firms, N. In particular, we show that biased consumer welfare might be non-monotonic in both γ and N.

Welfare effects of changes in the level of cursedness χ
As argued earlier, specification (3) is non-pessimistic so that from Proposition 2, when $\rho_0 < \mu_0$, the three equilibrium types segment the search cost space and almost everywhere the equilibrium is unique. The following proposition establishes a similar result for the cursedness parameter χ. To make the analysis tractable, assume in the following that $\upsilon > \omega$ in (3).[27]

Proposition 3 *There are unique $0 < \chi_0 < \chi_1$ so that:*
(i) There is a Diamond-type equilibrium with utility level 0 if and only if $\chi \leq \chi_0$. Moreover, for all $v \in (0, \overline{u}_1]$, there is a Diamond-type equilibrium with utility level v if and only if $\chi = \chi_0$.
(ii) There is a Stahl-type equilibrium if and only if $\chi \geq \chi_1$.
(iii) There is a penny sales equilibrium, which is unique, if and only if $\chi \in (\chi_0, \chi_1)$.

To derive the proposition, we show that under specification (3), irrespective of whether $\rho_0 < \mu_0$ or not, the equilibrium is as described in Proposition 2 with the cut-off values s_0 and s_1 (implicitly) given by

$$s_0 - \chi \left[\upsilon - (1 - \rho_0 - \mu_0) \cdot \min\left\{ \omega, \frac{s_0}{\mu_0} \right\} \right] = 0 \quad \text{and} \quad s_1 = \frac{\chi}{\upsilon}. \tag{22}$$

The proposition follows then from the fact that both s_0 and s_1 strictly increase in χ, as the left panel of Figure 4 illustrates.

Intuitively, as χ goes to zero, the difference in beliefs between biased and standard consumers disappears, because $\beta_B \to \beta_0$ as $\chi \to 0$. Consistent with this, for small χ the equilibrium outcome

[27] We impose $\upsilon > \omega$ only to simplify equation (87) in the appendix, using $\overline{u} - \underline{u} < \upsilon$ which follows from $\upsilon > \omega$. Intuitively, given (3), the latter means that a biased consumer believes that each firm offers with probability χ more and less utility than any firm truly does.

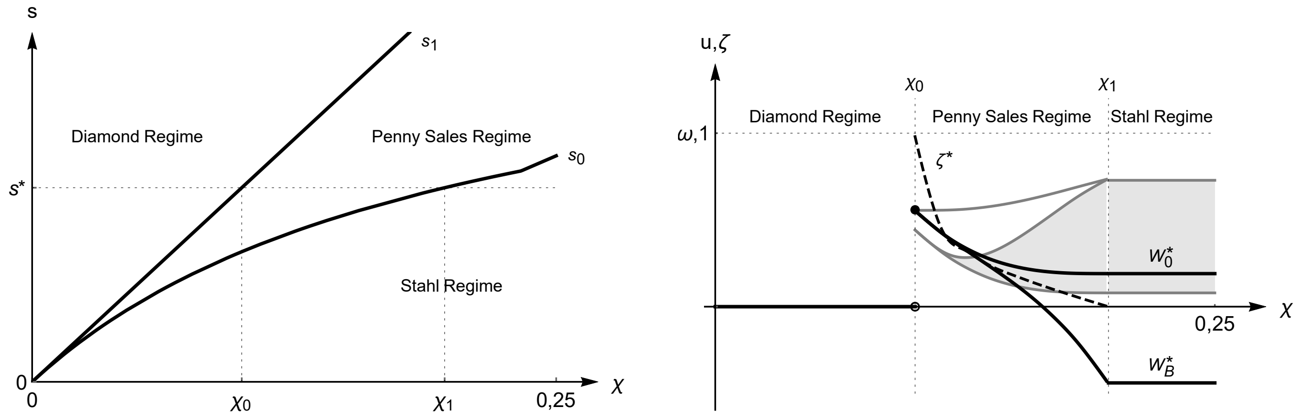

Figure 4: (Left) *The three equilibrium regimes as functions of s and* χ *with* s_0 *and* s_1 *(as functions of* χ*) given by (22).* (Right) *For fixed search costs* s^**, the equilibrium welfare of standard and biased consumers (black, solid), the support of the equilibrium utility distribution* κ^* *(gray area) as well as the equilibrium probability for a penny sale offer* ζ^* *(dashed) as functions of* χ*.*

coincides with the Diamond outcome. To the contrary, when χ is large (and search costs are not too large), biased consumers are strongly cursed and visit all firms due to their unrealistic expectations. Then the Stahl-type equilibrium obtains. For intermediate levels of χ, there is a penny sales equilibrium.

In what follows, we assume $\chi_1 < \frac{1}{4}$ so that any equilibrium type exists for some $\chi \in (0, \frac{1}{4})$. As the left panel of Figure 4 illustrates, $\chi_1 < \frac{1}{4}$ if s is sufficiently small.

The next proposition characterizes the welfare properties of the market as a function of the level of cursedness. Recall that the equilibrium is unique unless $\chi = \chi_0$. In what follows, we make the equilibrium selection that at $\chi = \chi_0$, the consumer optimal equilibrium is played, that is, firms offer the utility level $v^* = \overline{u}_1$ with probability 1. This simplifies, but is insubstantial for, our discussion.

Proposition 4 *(i) Standard consumer welfare is single-peaked in* χ *and maximal at* χ_0*.*
(ii) An intermediate value of cursedness $\chi \in [\chi_0, \chi_1]$ *maximizes biased consumer welfare. If*

$$\omega \geq \frac{s}{\mu_0}, \tag{23}$$

then biased consumer welfare is single-peaked in χ *and maximal at* χ_0*.*
(iii) A firm's profit is u-shaped in χ *and maximal in the Diamond regime for any* $\chi < \chi_0$*.*
(iv) Overall welfare decreases in χ *and is maximal in the Diamond regime for any* $\chi < \chi_0$ *and at* χ_0*.*

The main take-away from the proposition is that the welfare of both consumer types is maximized for intermediate levels of cursedness in the penny sales regime, as illustrated in the right panel of Figure 4, while firm profits and total welfare are maximized for low levels of cursedness in the Diamond regime.

To see what is behind consumer welfare, note that the more cursed biased consumers, the more optimistic they are about the benefits of search, as cursedness induces them to overestimate the degree to which offers are dispersed in the market.[28] Thus, an increase in cursedness increases their propensity to search. As χ moves from the Diamond regime into the penny sales regime at $\chi = \chi_0$, biased consumers begin to visit more than one firm. This creates competition and benefits both consumer types.

Interestingly however, for $\chi > \chi_0$, an increase in cursedness does not necessarily result in fiercer competition and better offers by firms. Even though more heavily cursed consumers visit more firms in expectation, they also become more demanding in terms of the utility that firms need to offer to induce them to stop and buy. This tilts firms' trade-off between maximizing market share and surplus share towards the latter. In other words, firms are less willing to compete for biased consumers and rather make low utility offers to extract more surplus from the less selective standard consumers, while speculating that some biased consumers may return after having visited all firms. As a result, the expected utility offered in the market decreases as cursedness increases when $\chi > \chi_0$, which implies that the welfare of standard consumers, who, in equilibrium, buy from the first firm they visit, decreases in χ for $\chi > \chi_0$.

The welfare of biased consumers is the sum of the expected consumption utility they derive from accepting an offer and the expected search expenditures. As to the former, as the right panel of Figure 4 shows, $\overline{u}^*$ increases with χ, since more heavily cursed consumers demand better offers to stop. Thus, conditional on receiving a penny sale offer, a biased consumer is better off as χ increases. However, as shown in the proof and the right panel of Figure 4, ζ^* decreases with χ, since, as explained above, firms increasingly target standard consumers. Thus, while penny sale offers become better as χ goes up, biased consumers receive such an offer less frequently which increases the likelihood that they end up with a relatively low offer after a futile search effort. As it turns out, under condition (23), these effects cancel each other and the overall effect of an increase in χ on their consumption utility is zero. However, biased consumers visit more firms in expectation so that their expected search expenditures go up, as they are less likely to encounter a penny sale offer that induces them to stop and buy. Thus, under condition (23), an increase in cursedness beyond χ_0 harms biased consumer welfare in the penny sales regime. Finally, in the Stahl regime, for $\chi > \chi_1$, biased consumers visit all firms irrespective of χ so that a change in χ does not affect the market outcome. Therefore, biased consumer welfare is independent of χ for

[28] Formally, $U(\beta_B(\kappa))$ increases in χ for any κ. Intuitively, from (3), for any κ, a biased consumer's belief given χ is a mean-preserving spread of her belief given χ' if $\chi > \chi'$ which entails that her reservation utility increases in χ.

$\chi > \chi_1$.

As to profits and total welfare, notice that in equilibrium, trade occurs with probability one, and consumers search only to find a good offer (rather than a good match). Hence, search is wasteful from the point of view of total welfare which is consequently maximal in the Diamond regime, where no search takes place, but also at χ_0 when firms make a penny sale offer with probability one so that no search occurs.[29] Likewise, profits are maximal in the Diamond regime since firms extract the entire surplus.

Finally, in any case, rational (weakly) exceeds biased consumer welfare, as both consumer types seek to maximize the same objective, however, biased consumers do so under incorrect beliefs. Nevertheless, all consumers may benefit from the presence of biased consumers in comparison to the rational consumer benchmark, including biased ones, because biased beliefs work as a commitment device to engage in search which can be seen as a public good that benefits all consumers.

Welfare effects of changes in the share of naive consumers γ

We next present comparative statics results with respect to γ and N. A formal treatment of the penny sales regime is analytically untractable due to the presence of various countervailing effects. Indeed, as is known from Stahl (1989), with respect to changes in N, the analysis is in general even untractable for the (less complicated) Stahl regime. Our results are therefore based on some partial analytical results which we present in the text below and the numerical examples depicted below in Figure 5 and 6.

As described in Proposition 2, for any γ, there is an interval in search costs (s_0, s_1) for which the penny sales equilibrium obtains. As the left panel of Figure 5 illustrates, the size of this interval vanishes as either $\gamma \to 0$ or $\gamma \to 1$ which is true for all parameter values.[30]

To understand this, observe first that given specification (3), when κ is degenerate, that is $\kappa = \delta_v$ for some v, then we have for a biased consumer's reservation utility as given in (4) and (5) that

$$\begin{aligned} U_B = v \quad &\Leftrightarrow \quad \chi \max\{U_B, v - \upsilon\} + (1-2\chi)\max\{U_B, v\} + \chi \max\{U_B, v+\upsilon\} - s \Big|_{U_B = v} = v \qquad (24) \\ &\Leftrightarrow \quad s = \chi \upsilon. \end{aligned}$$

[29]The fact that welfare is maximal in the Diamond regime, when the monopoly outcome obtains, is partially a result of the unit good assumption which implies that there are no monopoly distortions.

[30]Formally, given (22), $\lim_{\gamma\to 0} s_0 = \lim_{\gamma\to 0} s_1$ follows from $\lim_{\gamma\to 0} \rho_0 = 1$ and $\lim_{\gamma\to 0} \mu_0 = 0$ by Lemma 4, and $\lim_{\gamma\to 1} s_0 = \lim_{\gamma\to 1} s_1$ follows from $\lim_{\gamma\to 1} \rho_0 = 0$ and $\lim_{\gamma\to 1} \mu_0 = 1$ by Lemma 4. On a related note, there is a discontinuity at $\gamma = 0$ and $\gamma = 1$, because when there is only one consumer group in the market, then the Diamond-type equilibrium with utility level zero is the unique equilibrium.

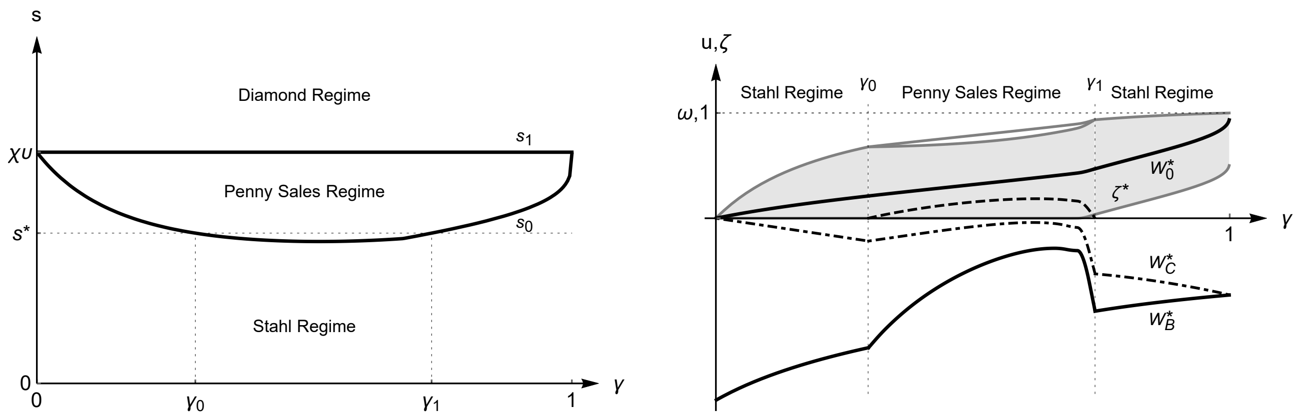

Figure 5: (Left) *The three equilibrium regimes as functions of* s *and* γ *with* s_0 *and* s_1 *(as functions of* γ*) given by (22).* (Right) *For fixed search costs* s^**, the equilibrium welfare of standard and biased consumers (black, solid), total consumers welfare (black, dot-dashed), the support of the equilibrium utility distribution* κ^* *(gray area) as well as the equilibrium probability for a penny sale offer* ζ^* *(dashed) as functions of* γ*.*

Because reservation utilities decrease in s, given $\kappa = \delta_v$, we have $U_B \le v$ if and only if $s \ge \chi\upsilon$, and otherwise, $U_B > v$. From this and Lemma 3, it follows immediately that a Diamond-type equilibrium with $\kappa^* = \delta_{\overline{u}^*}$ and $U_B^* \le \overline{u}^*$ exists if and only if $s \ge \chi\upsilon$ which is equivalent to $s \ge s_1$ by (22).

When $s < \chi\upsilon$ and $\gamma \to 0$ or $\gamma \to 1$, then, similar to Stahl (1989), all firms target the predominant consumer group, that is, when $\gamma \to 0$, they target the selective biased consumers, resulting in good offers, and when $\gamma \to 1$, they target the less search prone standard consumers resulting in bad offers. Indeed, in any Stahl-type or penny sales equilibrium, if it exists, we have that the utility distribution becomes degenerate: $\kappa^* \to \delta_0$ and $\overline{u}^* \to 0$ as $\gamma \to 0$ as well as $\kappa^* \to \delta_\omega$ and $\overline{u}^* \to \omega$ as $\gamma \to 1$.[31] From this, it follows that in any equilibrium, when $s < \chi\upsilon$ and $\gamma \to 0$ or $\gamma \to 1$, we must have $U_B^* < \overline{u}^*$, as argued in the previous paragraph. Thus, from Lemma 3, only a Stahl-type equilibrium with $U_B^* < \overline{u}^*$ exists when $s < \chi\upsilon$ and $\gamma \to 0$ or $\gamma \to 1$, but not a penny sales equilibrium which would require that $U_B^* = \overline{u}^*$.

The comparative statics results with respect to γ in the Stahl regime are identical to the comparative statics with respect to the share of shoppers in Stahl (1989):[32] As argued above and as

[31] To see this, recall that in any Stahl-type ($\zeta^* = 0$) or penny sales equilibrium ($\zeta^* \in (0,1)$) the upper and lower bound of the support are given by (13) as $\overline{u}^* = (1-\rho(\zeta^*))\,\omega + \rho(\zeta^*)\,\underline{u}^*$ and $\underline{u}^* = \max\left\{0, \omega - \frac{s}{\mu(\zeta^*)}\right\}$. From this and $\lim_{\gamma\to 0}\rho(\zeta) = 1$, $\lim_{\gamma\to 1}\rho(\zeta) = 0$ as well as $\lim_{\gamma\to 0}\mu(\zeta) = 0$ for all $\zeta \in [0,1]$ by part (v) of Lemma 4, it follows that $\kappa^* \to \delta_0$ and $\overline{u}^* \to 0$ as $\gamma \to 0$ as well as $\overline{u}^* \to \omega$ as $\gamma \to 1$. Moreover, $\kappa^* \to \delta_\omega$ as $\gamma \to 1$ follows from $\mathbf{E}(u) = \mu(\zeta)\,\omega + (1-\mu(\zeta))\,\underline{u}^*$ by (45) and $\lim_{\gamma\to 1}\mu(\zeta) = 1$ for all $\zeta \in [0,1]$ by part (v) of Lemma 4.

[32] The only difference in the welfare analysis is that in our Stahl-type equilibrium, biased consumers incur the total

the right panel of Figure 5 illustrates, when $\gamma \to 0$, then $\kappa^* \to \delta_0$, and when $\gamma \to 1$, then $\kappa^* \to \delta_\omega$. To see the latter in the plot, recall that standard consumer welfare is $W_0^* = \mathbf{E}(u)$ and notice that $W_0^* \to \omega$ as $\gamma \to 1$. Moreover, in the Stahl regime a marginal increase in γ increases κ^* in the sense of first order stochastic dominance which improves standard and biased consumer welfare.

In contrast, in the penny sales regime, biased consumer welfare may be non-monotonic. This may be the case even if an increase in γ increases the average utility being offered, and hence standard consumer welfare, as the example depicted in the right panel of Figure 5 shows. To understand this, recall that a biased consumer strongly benefits from the presence of penny sales offers both because they supply much utility and, importantly, because they induce her to stop her search, thus preventing her from sampling all firms and incurring large total search costs. Intuitively, the non-monotonicity in biased consumer welfare thus reflects that the probability for a penny sale offer, ζ^*, is non-monotonic in the penny sales regime, as we (must) have $\zeta \to 0$ when $\gamma \to \gamma_0$ or $\gamma \to \gamma_1$, as the plot shows.

Total consumer welfare is given as the weighted sum of standard and biased consumer welfare: $W_C^* = \gamma \cdot W_B^* + (1-\gamma) \cdot W_0^*$. As γ increases, there is a direct effect on total consumer welfare, because some standard consumers turn into biased consumers. As a consequence, the welfare of these consumers declines, as $W_B \leq W_0$. Hence, an increase in γ never induces a Pareto improvement from the consumers' point of view. Figures 5 shows that total consumer welfare may be non-monotonic in γ.

Welfare effects of changes in the number of firms N

The left panel of Figure 6 illustrates the different equilibrium regimes as a function of s and N, and shows that a penny sales equilibrium is more likely to exist if the numbers of firms is large. To understand this, consider the case of intermediate search costs and small N. As the plot shows, a Stahl-type equilibrium obtains in this case. Now recall that in a Stahl-type equilibrium a biased consumer visits all firms. Thus, as the number of firms increases, each firm is still visited by the same mass of biased consumers but the share of standard consumers that visit it first, and buy its product in equilibrium, becomes smaller as N increases. Thus, intuitively, at some point, it pays off to make a penny sale offer that induces all visiting biased consumers to stop and buy. As a consequence, a penny sales equilibrium obtains for sufficiently large N for intermediate search costs.[33]

search costs $N \cdot s$ from visiting all firms while shoppers in Stahl (1989) incur no costs from doing so. However, as total search costs are independent of γ, this difference is irrelevant for the comparative statics wrt γ.

[33] This argument does not hold for very small search costs, when $s < \chi \cdot (\upsilon - \omega)$, as in this case, a firm must offer more than the entire surplus from trade to induce a biased consumer to stop and buy. Formally, from Lemma 3, when $s < \chi \cdot (\upsilon - \omega)$, only a Stahl-type equilibrium exists, because $U_B^* > \overline{u}^*$ in any equilibrium, as $U_B^* > \omega$ in any

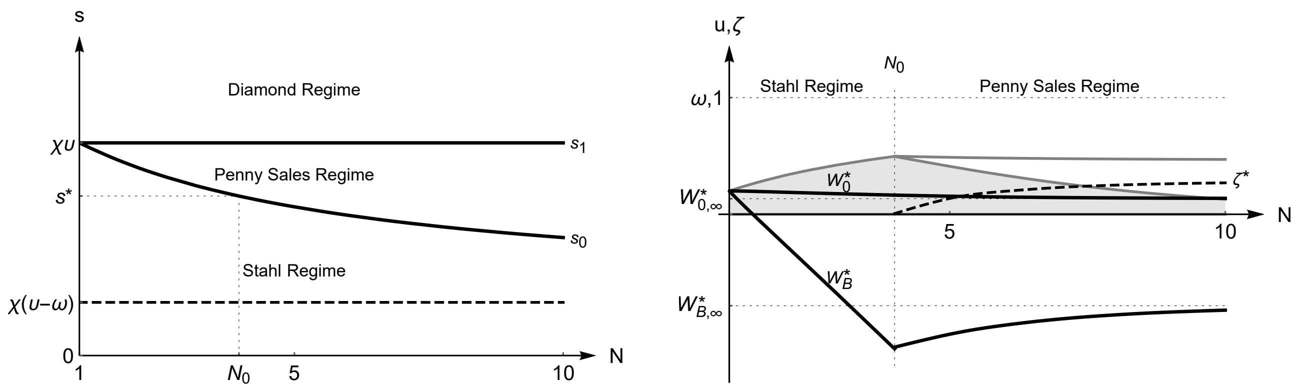

Figure 6: (Left) *The three equilibrium regimes as functions of* s *and* N *with* s_0 *and* s_1 *(as functions of* N*) given by (22).* (Right) *For fixed search costs* s^**, the equilibrium welfare of standard and biased consumers (black, solid), the support of the equilibrium utility distribution* κ^* *(gray area) as well as the equilibrium probability for a penny sale offer* ζ^* *(dashed) as functions of* N*.*

An increase in the number of firms may reduce standard consumer welfare both in the Stahl as well as the penny sales regime as the right panel of Figure 6 shows. This is similar to Stahl (1989) where the same effect can occur. Intuitively, in the Stahl regime, as N increases, firms must offer more and more utility to make the best offer in the market and obtain the entire demand from biased consumers. This increases a firm's incentive to evade competition for biased consumers and target the less demanding standard consumers with low utility offers. As a consequence, the average utility offered (and thus standard consumer welfare) may drop as N increases. Indeed, when $s < \chi \cdot (\upsilon - \omega)$ so that there is a Stahl-type equilibrium for large N, this has the perverse consequence that the average utility offered, and thus standard consumer welfare, goes to zero, as $N \to \infty$.

On the other hand, in the penny sales regime, primarily two effects affect standard consumers' welfare as N increases. On the plus side, the probability for a penny sale offer may increase. On the minus side, firms who target standard consumers offer less utility as N increases. As the plot illustrates, the latter effect may dominate. This reflects that as $N \to \infty$, from (14), κ^* becomes bi-modal: intuitively, firms who offer less utility than a penny sale offer are not competing for the demand from returning biased consumers any more, since the probability that the latter encounter a penny sale elsewhere (and do not return) goes to one, as $N \to \infty$. In contrast to the Stahl regime, standard consumer welfare does not vanish in the penny sales regime as $N \to \infty$. This is because firms make a strictly positive penny sale offer with strictly positive probability, both

equilibrium. To see the latter, given (3), $U_B = \omega$ when $s = \chi \cdot (\upsilon - \omega)$ and $\kappa = \delta_0$ which implies that $U_B^* > \omega$ when $s < \chi \cdot (\upsilon - \omega)$, because κ^* fosd δ_0 in any equilibrium.

bounded away from zero, as $N \to \infty$.[34] The underlying reason for this difference is that in a penny sales equilibrium, a firm does not have to make the best offer in the market to derive demand from biased consumers. Intuitively, this prevents that all firms target standard consumers, as $N \to \infty$.

Biased consumers' welfare may be non-monotonic in N. They may suffer from an increase in the number of firms in the Stahl regime: while they may derive a higher expected consumption utility, this may be outweighed by the search costs they incur from visiting all firms increase, which increases as the number of firms increases. In the penny sales regime, an increase in N may result in an increase of ζ^*, as the right panel of Figure 6 shows. As argued before, biased consumers strongly benefit from the presence of penny sale offers. As a result, biased consumer welfare may increase as N increases in the penny sales regime.

6 Discussion

Our analysis builds on the seminal search model due Diamond (1971) which considers the important benchmark of a market for a homogeneous product. This basic framework allows us to isolate the role of biased consumer beliefs for equilibrium outcomes. In what follows, we discuss to what extent our insights carry over to richer models.

Heterogeneous products: While firms offer homogeneous products in our model, our notion of a belief function can also be applied to search models where firms offer heterogeneous products and consumer have heterogeneous tastes as captured by idiosyncratic match values for each firm (Wolinsky (1986)). For example, biased consumers may have systematically wrong beliefs about the distribution of match values. Our key qualitative insights, that biased beliefs can have pro-competitive effects, and that this depends on whether biased consumer overestimate the upper tails of the distribution (of match values), is likely to be robust in such a model. The reason is that overestimating the upper tails of the match-value distribution results in stronger incentives to search in the form of higher reservation match-values often resulting in lower prices as is argued in Anderson and Renault (1999).

Heterogeneous search costs We assume that standard and biased consumers have the same search costs. The fact that the search costs are the same across the two consumer groups is insubstantial for our analysis. If search costs within each consumer group were characterized by a distribution of search costs (with support bounded away from zero), then Proposition 1 remains true, and

[34] While the right panel of Figure 6 illustrates this for some parameter values, it can be shown analytically that, in general, firms make a strictly positive penny sale offer with strictly positive probability, both bounded away from zero, as $N \to \infty$. We omit the details.

whether the Diamond Paradox breaks down under non-pessimism depends only on the search behaviour of the biased consumers with the lowest search costs. However, for general search cost distributions the model is not tractable enough to characterize the equilibria as this would give rise to a non-trivial distribution of reservation values.[35]

Different knowledge: We make the simplifying assumption that all biased consumers entertain the same belief function. In practise, biased consumers are likely to have different knowledge about the true market condition which would result in different belief functions. Equilibria other than the Diamond equilibrium would then exist (for some search costs) if and only if there is a positive mass of consumers with non-pessimistic beliefs.

No recall: We assume that consumers' search is with recall which means that a consumer can return to any past offer at not cost. This significantly enhances tractability, because it allows us to work with stationary reservation values. At a general level, a consumer's search incentives with recall are driven only by her beliefs about the upper tail of the utility distribution, because receiving a low utility offer tomorrow has no payoff consequences, as the consumer can return to her best past offer. In contrast, with no recall, when a consumer does not stop today, she may become stuck with a low utility offer, and hence the lower tail of the distribution matters, too. This suggests that non-pessimism alone, which only implies that a consumer overestimates the best offer in the market, might therefore not be sufficient to foster search with no recall.

Against this backdrop it is perhaps surprising that even under no recall there are other equilibria than the Diamond equilibrium if and only the belief function is non-pessimistic. One reason for this is that, regardless of the belief function, in our setting a strategy profile is a Diamond-type equilibrium with utility level 0 under no recall if and only if it is so under recall. Intuitively, when all firms offer the outside option with probability 1, a consumer, even without recall, can guarantee himself the best past offer by choosing the outside option, and so search incentives are the same with and without recall. In particular, this implies that whenever the Diamond paradox breaks down with recall, it does so without recall. Similarly, Proposition 1 remains valid without recall. A general analysis of how equilibria look like under no recall when the Diamond paradox breaks down, involves non-stationary reservation values and is beyond the scope of the current article.

Learning We assume that biased consumers (stubbornly) think that their belief about the utility distribution in the market is correct with probability 1. While this is somewhat similar to recent work on misspecified models such as Heidhues et al. (2021), it implies that biased consumers

[35]Formally, our characterization of equilibria builds on Lemma 2 which characterizes, in general, the set of equilibria that may exist given there are two consumer groups with two (exogenous) reservation utilities. The assumption of two consumer groups is common in the search literature with homogeneous products to enhance tractability.

do not learn about the underlying distribution from observed offers. In fact, without further restrictions on the belief function, it cannot be ruled out that a biased consumer deems offers impossible that are actually in the support of the true utility distribution but still does not change her beliefs when encountering such an offer. It is thus noteworthy that such an extreme form of denial cannot occur within our specification of cursed beliefs where the support of equilibrium offers is contained in the support of the biased consumer's beliefs. The fact that a biased consumer does not learn also means that, in a Stahl-type equilibrium, where she searches all firms, she incurs arbitrary large total search costs as the number of firms grows.[36] This feature is reminiscent of the "searching forever after" logic in Antler and Bachi (2021) and would not be present in a model in which the agent's beliefs would gradually become more realistic. An analysis of a model which incorporates learning would entail non-stationary reservation values and is beyond the scope of this article.

7 Conclusion

This articles illustrates how biased consumer beliefs shape search market outcomes. More specifically, it shows that when the consumers' beliefs are non-pessimistic, then novel equilibria with penny sales offers may arise in which firms respond to the consumers' inflated expectations by actually making excellent offers. As a result, biased consumer beliefs may actually improve the welfare of all consumers.

Consumer protection policies that aim at educating consumers such as publishing information about the true offer distribution may move biased consumers' beliefs closer to the truth or even debias some consumers entirely. In our model, the former corresponds to reducing the size of the bias whereas the latter corresponds to reducing the share of biased consumers.

Our analysis shows that reducing the size of the bias may benefit all consumers when the size of the bias is large. However, at intermediate levels, eliminating the bias entirely may actually harm all consumers. By contrast, reducing the share of biased consumers may not constitute a Pareto improvement from the consumers' point of view: While it always benefits consumers who hold rational beliefs after but not before the intervention, the equilibrium effect may harm all other consumers in the market.

Our article considers a setting with consumers who are characterized by a general belief function, instead of focussing on a specific behavioural consumer bias. This allows us to identify

[36] More precisely, as the left panel of Figure 6 illustrates, there is a Stahl-type equilibrium for all N, when search costs are very small, that is $s < \chi \cdot (\upsilon - \omega)$. Then, biased consumer welfare goes to $-\infty$, as $N \to \infty$, because biased consumers incur the search costs $N \cdot s$ from visiting all firms.

pessimism and non-pessimism of the belief function as the driving forces behind the search market outcomes in our setting. To explain observed consumer search behavior in a specific market, it might be useful to be more specific about the consumers' behavioural biases. In this article, we provide examples of biases that induce pessimistic respectively non-pessimistic beliefs and may help explain why in some markets some consumers appear to search too little respectively too much. For example, our introductory example suggests that when consumers suffer from correlation neglect and have cursed beliefs, then they may search too much in vertically differentiated markets. Likewise, markets in which consumer beliefs are pessimistic due to an overprecision bias or a belief in the law of small numbers (see Grubb (2015)).

In our setting, consumers are endowed with an exogenous biased belief function. In practice, firms can influence beliefs through advertising and marketing messages. An interesting avenue for future research is to extend the framework in this article to allow for the belief function to depend on strategic actions by firms.

Appendix

Proof of Lemma 1 We first show that $\underline{u}^* \geq \max\{0, \min\{U_0^*, U_B^*\}\}$. To the contrary, suppose a firm offers $u < \max\{0, \min\{U_0^*, U_B^*\}\}$. Then, from (8), no consumer stops and buys, and, as we argue below, no consumer returns, contradicting that firms make positive profits in equilibrium (see Footnote 19). Note for consumers to return to a firm that offers $\underline{u}^*$, there must be a mass point at $\underline{u}^*$ because otherwise the probability that $\underline{u}^*$ is the best offer has probability 0. This can, however, not be the case in equilibrium, as otherwise, the firm could offer slightly more utility than $\underline{u}^*$ and discontinuously increase its demand (the latter is possible because $\underline{u}^* < \omega$ as firms make strictly positive profits by Footnote 19).

Next, we show $\underline{u}^* \leq \max\{0, \min\{U_0^*, U_B^*\}\}$. From part (ii) below, $\overline{u}^* \leq \max\{0, \max\{U_0^*, U_B^*\}\}$, and hence, $\underline{u}^* < \max\{0, \max\{U_0^*, U_B^*\}\}$, because $\underline{u}^* < \overline{u}^*$ by assumption. Moreover, $\underline{u}^* \notin (\max\{0, \min\{U_0^*, U_B^*\}\}, \max\{0, \max\{U_0^*, U_B^*\}\})$, because otherwise, a firm which offers $\underline{u}^*$ only derives demand from those consumers with the lower reservation utility who visit it first, because, as argued in the preceding paragraph, no consumer returns when it offers the least utility. Thus, as $\underline{u}^* > \max\{0, \min\{U_0^*, U_B^*\}\}$, it could offer slightly less utility without losing any demand from these consumers.

The argument for (ii) is given in the text. ■

Proof of Lemma 2 The arguments for the cases $\overline{u}^* > U_B^*$ and $\overline{u}^* < U_B^*$ are given in the main text. For $\overline{u}^* = U_B^*$, let ζ denote the mass of a (potential) mass point at $\overline{u}^*$. If $\zeta = 0$, then a biased consumer never encounters an offer that stops her. She visits all firms and returns to the best offer. From standard arguments, κ^* has thus no mass points and its support has no gaps, as described in part (ii) of Lemma 2. Next, if $\zeta = 1$, then all firms offer the same utility and $\underline{u}^* = \overline{u}^* = v^*$ with $U_0^* \leq U_B^* = v^*$, as described in part (i) of Lemma 2. Finally, if $\zeta \in (0, 1)$, then with probability $(1-\zeta)^N$ no firm offers $\overline{u}^*$ so that a biased consumer visits all firms and returns to best offer. Then, standard arguments imply that on $Supp(\kappa^*) \cap [\underline{u}^*, \overline{u}^*)$,[37] κ^* has no mass points and its support no gaps in the sense that $Supp(\kappa^*) \cap [\underline{u}^*, \overline{u}^*)$ is connected, as described in part (iii) of Lemma 2. To complete the proof, we show that the support has a "gap between" $Supp(\kappa^*) \cap [\underline{u}^*, \overline{u}^*)$ and $\overline{u}^*$, that is, $\{Supp(\kappa^*) \cap [\underline{u}^*, \overline{u}^*)\} \cup \overline{u}^*$ is not connected. The reason for this is that it cannot be profit maximizing for a firm to offer marginally less utility than $\overline{u}^*$ when there is a mass point at $\overline{u}^*$ and $\overline{u}^* = U_B^*$. This follows from the fact that offering marginally less utility than $\overline{u}^*$ increases the firm's share of the surplus only marginally whereas its loss in demand increases discontinuously, because a biased consumer does not stop anymore and only returns if she does not encounter a

[37] The restriction to $Supp(\kappa^*) \cap [\underline{u}^*, \overline{u}^*)$ reflects that there are only consumers that return if no firm offers $\overline{u}^*$.

firm that offers $\overline{u}^*$ (which occurs with positive probability because there is a mass point at $\overline{u}^*$ by assumption). ■

Proof of Lemma 3 As to (i). From Lemma 2, (U_0^*, U_B^*, κ^*) is a Diamond-type equilibrium if and only if

(α) $\kappa^* = \delta_{v^*}$ is profit-maximizing given U_0^*, U_B^*, and all other firms play $\kappa^* = \delta_{v^*}$,

(β) $\max\{U_0^*, U_B^*\} \leq v^*$ with $U_0^* = U(\kappa^*)$ and $U_B^* = U(\beta_B(\kappa^*))$.

Thus, it suffices to show that (α) and (β) is met if and only if (a) and (b) from part (i) of Lemma 3 are met.

We begin with the "only if"-part: First, for $v^* = 0$, (α) and (β) trivially imply (a). Second, let $v^* > 0$. We have to show that (α) and (β) imply (b), that is, $U(\beta_B(\delta_{v^*})) = v^*$ and $v^* \leq \overline{u}_1$. Since $\max\{U_0^*, U_B^*\} \leq v^*$ by (β), all consumers buy from the first firm they visit. Moreover, $U_0^* = U(\delta_{v^*}) = v^* - s < v^*$ by (4). Hence, if, by contradiction, $U_B^* = U(\beta_B(\delta_{v^*})) < v^*$, then a firm could offer slightly less (non-negative) utility than v^* so as to increase its mark-up without losing any demand, as all consumers would still buy from the first firm they visit. Thus, $U(\beta_B(\delta_{v^*})) = v^*$.

To see that $v^* \leq \overline{u}_1$, note that (α) implies in particular that offering v^* is (weakly) more profitable than offering $u^{dev} = \max\{0, v^* - s\}$, that is,

$$\pi(v^*) \geq \pi(\max\{0, v^* - s\}) \quad \Leftrightarrow \quad \frac{1}{N}(\omega - v^*) \geq \frac{1-\gamma}{N}(\omega - \max\{0, v^* - s\}). \tag{25}$$

To understand (25), recall that $\omega - u$ is the firm's mark-up if it offers u. Its market share when it offers v^* is $1/N$, because in this case, the demand of all consumers is split equally among all firms. On the other hand, $U_0^* = v^* - s$ and $U_B^* = v^*$ as argued in the previous paragraph. Thus, from (8), u^{dev} is the minimal utility that induces a standard consumer to stop and buy, but does not offer enough utility for a biased consumer to stop and buy, as $u^{dev} < U_B^*$, leading to a market share of $(1-\gamma)/N$, because no biased consumers returns, as all other firms offer $v^* = U_B^*$. We now show that (25) is equivalent to $v^* \leq \overline{u}_1$: First, if $v^* \geq s$, then (25) is true if, and only if,

$$v^* \leq \omega - \frac{1-\gamma}{\gamma}s. \tag{26}$$

Second, if $v^* \in [0, s)$, then (25) holds if, and only if,

$$v^* \leq \gamma\omega. \tag{27}$$

Now, observe that for $v^* = s$, (26) is equivalent to (27). Indeed, (27) implies (26) for $v^* \geq s$ and (26) implies (27) for $v^* \in [0, s)$. Thus, for any $v^* \geq 0$, (25) is true if and only if the less restrictive condition among the two is met. Therefore, (25) is true if and only if

$$v^* \in \left[0, \max\left\{\gamma\omega, \omega - \frac{1-\gamma}{\gamma}s\right\}\right]. \tag{28}$$

This is equivalent to $v^* \in [0, \overline{u}_1]$, because

$$\overline{u}_1 = (1-\rho(1))\omega + \rho(1)\max\left\{0, \omega - \frac{s}{\mu(1)}\right\} = \gamma\omega + \max\left\{0, (1-\gamma)\omega - \frac{1-\gamma}{\gamma}s\right\}, \tag{29}$$

where we used $\rho(1) = 1-\gamma$ and $\mu(1) = \gamma$ from Lemma 4 in the appendix. This completes the proof of the "only if"-part.

As to the "if"-part, let (a) and (b) be given. Since $U(\beta_B(\delta_{v^*})) \leq v^*$ by assumption, and since (4) entails that $U_0^* = U(\delta_{v^*}) = v^* - s < v^*$, part ($\beta$) follows. To see ($\alpha$), we argue that offering v^* is profit-maximizing. Again, offering more utility than v^* only reduces the firm's mark-up without generating additional demand. On the other hand, if a firm offers less utility than v^*, then this is clearly not optimal if $v^* = 0$ (because it derives no demand). If $v^* > 0$, then offering less than v^* implies that the firm loses the entire demand from biased consumers, since $U(\beta_B(\delta_{v^*})) = v^*$ by assumption. Therefore, conditional on offering less than $v^* > 0$, offering $u^{dev} = \max\{0, U_0^*\}$ is optimal, because this is the lowest offer that induces a standard consumer to stop and buy. Hence, (25) and the ensuing calculations imply that offering v^* is profit-maximizing for $v^* \leq \overline{u}_1$. This implies (α) and completes the proof.

As to (ii). The proof of part (ii) is identical to the proof of part (iii) for the special case that $\zeta^* = 0$ and where instead of the condition $\overline{u}^* = U_B^*$ in (31), we impose the condition $\overline{u}^* \leq U_B^*$.

Part (iii) follows from the following four claims. Claim 1 reformulates the equilibrium conditions for a penny sales equilibrium in terms of a system of equations. Claims 2 and 3 show that this system pins down the utility distribution up to the probability ζ^* with which a penny sale is offered in equilibrium. Claim 4 then completes the proof by establishing the desired condition on ζ^* for equilibrium existence.

Claim 1: (U_0^*, U_B^*, κ^*) is a penny sales equilibrium if and only if it satisfies (30), (31), (32) as well as (33) and (34).

To see Claim 1, note that by Lemma 1 and 2, in a penny sales equilibrium

$$\underline{u}^* = \max\{0, U_0^*\}, \quad \text{with} \quad U_0^* = U(\kappa^*), \tag{30}$$

$$\overline{u}^* = U_B^*, \quad \text{with} \quad U_B^* = U(\beta_B(\kappa^*)), \tag{31}$$

$$Supp(\kappa^*) = [\underline{u}^*, \hat{u}^*] \cup \{\overline{u}^*\}, \quad \text{with a mass point of size } \zeta^* \text{ at } \overline{u}^*. \tag{32}$$

For (U_0^*, U_B^*, κ^*) to be an actual equilibrium, in addition, κ^* needs to be profit maximizing, given all other firms adopt κ^* and (30), (31) and (32). Note first that, given (30), (31) and (32), $u \notin [\underline{u}^*, \hat{u}^*] \cup \{\overline{u}^*\}$ is never profit maximizing. Offering less than $\underline{u}^*$ generates no demand, as no consumer stops or returns, and offering more than $\overline{u}^*$ does not generate any additional demand in comparison to offering $\overline{u}^*$ but reduces the firm's margin (in either case any visiting consumer buys). Similarly, offering $u \in (\hat{u}^*, \overline{u}^*)$ does not generate any additional demand in comparison to offering $\hat{u}^*$ but reduces the firm's margin. It induces no biased consumer to stop, as $u < \overline{u}^* = U_B^*$. Moreover, no biased consumer returns who would not do so if the firm offered just $\hat{u}^*$. Therefore, κ^* is profit maximizing if and only if any $u \in Supp(\kappa^*)$ generates the same profit, that is,

$$\pi^* = (\omega - u) \cdot \left[\frac{1-\gamma}{N} + \gamma \kappa^*(u)^{N-1} \right] \quad \text{for all } u \in [\underline{u}^*, \hat{u}^*], \quad \text{and} \tag{33}$$

$$\pi^* = \left[\omega - \overline{u}^* \right] \cdot \left[\frac{(1-\gamma)}{N} + \gamma \cdot \eta(\zeta^*) \right], \tag{34}$$

where the function $\eta : [0,1] \to [0,1]$ is defined by[38]

$$\eta(\zeta) \equiv \sum_{t=1}^{N} \binom{N-1}{t-1} \frac{\zeta^{t-1}(1-\zeta)^{N-t}}{t}. \tag{35}$$

To see that Condition (33) describes a firm's profit when it supplies $u \in [\underline{u}^*, \hat{u}^*]$, note that given (30) and (32), a standard consumer buys from any firm with probability $1/N$, as with equal probability each firm is the first one she visits and buys from. Moreover, given (31) and (32), a biased consumer never stops at the firm since $u < \overline{u}^*$ but buys from it if it supplies the highest utility in the market (in which case the biased consumer visits all firms and returns to the firm) which occurs with probability $\kappa^*(u)^{N-1}$.

To understand (34), if a firm offers $\overline{u}^*$, then its mark-up is $\omega - \overline{u}^*$ and it receives its share $\frac{1-\gamma}{N}$ of standard consumers. In addition, if there are $t-1$ many other firms that offer $\overline{u}^*$, then a biased consumer buys from this firm with probability $1/t$, as this is the probability that she visits it first among the t many firms that offer $\overline{u}^*$. Furthermore, the probability that among its $n-1$ many competitors exactly $t-1$ offer $\overline{u}^*$ is $\binom{N-1}{t-1}(\zeta^*)^{t-1}(1-\zeta^*)^{N-t}$. Therefore, the probability that a biased consumer buys from a firm that offers $\overline{u}^*$ is $\eta(\zeta^*)$, and (34) follows.

Claim 2: Let $\overline{u}_{\zeta,\underline{u}}$ and $\hat{u}_{\zeta,\underline{u}}$ be defined as the support bounds $\overline{u}_\zeta$ and $\hat{u}_\zeta$ defined in (13), and let $\kappa_{\zeta,\underline{u}}$ be defined as the function κ_ζ defined in (14) with the interpretation that ζ and $\underline{u}$ are independent variables. Then we have that, given $\underline{u}^*$ and ζ^*, κ^* satisfies (32), (33), and (34) if and only if $\kappa^* = \kappa_{\zeta^*,\underline{u}^*}$.

Indeed, observe first, because $\kappa^*(\underline{u}^*) = 0$, (33) implies that

$$\pi^* = \pi(\underline{u}^*) = (\omega - \underline{u}^*) \cdot \frac{1-\gamma}{N}. \tag{36}$$

From this, it follows straightforwardly that (33) holds if and only if κ^* satisfies (14).

[38] We set $\eta(0) = 1$ and $\eta(1) = \frac{1}{N}$ so that $\eta(0) = \lim_{\zeta \to 0} \eta(\zeta)$ and $\eta(1) = \lim_{\zeta \to 1} \eta(\zeta)$.

Second, (32) holds if and only if $\hat{u}^*$ satisfies $\kappa^*(\hat{u}^*) = 1-\zeta^*$. Because κ^* satisfies (14),

$$\kappa^*(\hat{u}^*) = 1-\zeta^* \quad \Leftrightarrow \quad \frac{\rho_0}{1-\rho_0}^{\frac{1}{N-1}}\left(\frac{\omega-\underline{u}^*}{\omega-\hat{u}^*}-1\right)^{\frac{1}{N-1}} = 1-\zeta^* \tag{37}$$

$$\Leftrightarrow \quad \frac{\omega-\underline{u}^*}{\omega-\hat{u}^*}-1 = (1-\zeta^*)^{N-1}\cdot\frac{1-\rho_0}{\rho_0}. \tag{38}$$

Re-arranging terms and using the definition of $\phi(\cdot)$ in (121), yields $\hat{u}^* = \hat{u}_{\zeta^*,\underline{u}^*}$.

Third, given (33), condition (34) holds if and only if

$$\pi(\underline{u}^*) = \pi(\overline{u}^*) \quad \Leftrightarrow \quad (\omega-\underline{u}^*)\cdot\frac{\rho_0}{1-\rho_0}\gamma = (\omega-\overline{u}^*)\cdot\left(\frac{\rho_0}{1-\rho_0}\gamma+\gamma\eta(\zeta^*)\right) \tag{39}$$

$$\Leftrightarrow \quad \frac{\omega-\underline{u}^*}{\omega-\overline{u}^*}-1 = \eta(\zeta^*)\frac{1-\rho_0}{\rho_0}. \tag{40}$$

Re-arranging terms and using the definition of $\rho(\cdot)$ in (120), yields $\overline{u}^* = \overline{u}_{\zeta^*,\underline{u}^*}$. Because κ^* satisfies (14) with $\hat{u}^* = \hat{u}_{\zeta^*,\underline{u}^*}$ and $\overline{u}^* = \overline{u}_{\zeta^*,\underline{u}^*}$, we have $\kappa^* = \kappa_{\zeta^*,\underline{u}^*}$.

Finally, we verify that $\hat{u}^* < \overline{u}^*$, as required by (32). With $\hat{u}^* = \hat{u}_{\zeta^*,\underline{u}^*}$ and $\overline{u}^* = \overline{u}_{\zeta^*,\underline{u}^*}$ and because $\underline{u}^* < \omega$, we have

$$\hat{u}^* < \overline{u}^* \quad \Leftrightarrow \quad \phi(\zeta^*) > \rho(\zeta^*). \tag{41}$$

We establish the right hand side in part (iii) of Lemma 4.

Claim 3: Let ζ^* be given. Then $\underline{u}^*$, κ^* and U_0^* satisfy (30), (32), (33), and (34) if and only if $\kappa^* = \kappa_{\zeta^*,\underline{u}^*}$ and $\underline{u}^* = \underline{u}_{\zeta^*}$ as given by (13).

Indeed, from Claim 2, κ^* satisfies (32), (33), and (34) if and only if $\kappa^* = \kappa_{\zeta^*,\underline{u}^*}$, and therefore, condition (30) is, in addition, satisfied if and only if $\underline{u}^*$ is a solution to the fixed point equation

$$\underline{u} \quad = \quad \max\{0, U(\kappa_{\zeta^*,\underline{u}})\}. \tag{42}$$

Now, for any solution $\underline{u}^*$ to (42), we have that $\underline{u}^* \geq U(\kappa_{\zeta^*,\underline{u}^*})$ and from (4),

$$U(\kappa_{\zeta^*,\underline{u}^*}) = -s + \int \max\{u, U(\kappa_{\zeta^*,\underline{u}^*})\}\, d\kappa_{\zeta^*,\underline{u}^*}(u) \tag{43}$$

$$= -s + \int u\, d\kappa_{\zeta^*,\underline{u}^*}(u), \tag{44}$$

Below, we show that for all $\underline{u}$ and ζ, we have

$$\int u\, d\kappa_{\zeta,\underline{u}}(u) = \mu(\zeta)\omega + (1-\mu(\zeta))\,\underline{u}. \tag{45}$$

Hence,

$$U(\kappa_{\zeta^*,\underline{u}^*}) = -s + \mu(\zeta^*)\omega + (1-\mu(\zeta^*))\,\underline{u}^*, \tag{46}$$

and we infer that $\underline{u}^*$ is indeed a solution to (42) if and only if

$$\underline{u}^* \;=\; \max\{0, -s + \mu(\zeta^*)\omega + (1-\mu(\zeta^*))\,\underline{u}^*\}. \tag{47}$$

Notice that (a) $\underline{u}^* = 0$ is a solution to (47) if and only if $\omega - \frac{s}{\mu(\zeta^*)} \leq 0$, and (b) $\underline{u}^* > 0$ is a solution to (47) if and only if $\underline{u}^* = \omega - \frac{s}{\mu(\zeta^*)} > 0$. Clearly, there is no solution to (47) with $\underline{u}^* < 0$. Therefore, $\underline{u}^* = \max\left\{0, \omega - \frac{s}{1-\mu(\zeta^*)}\right\} = \underline{u}_{\zeta^*}$ is the unique solution to (42).

Claim 4: There is a penny sales equilibrium if and only if there is $\zeta^* \in (0,1)$ such that $\overline{u}_{\zeta^*} = U(\beta_B(\kappa_{\zeta^*}))$.

From Claims 1 and 3, (U_0^*, U_B^*, κ^*) is a penny sales equilibrium if and only if $\kappa^* = \kappa_{\zeta^*,\underline{u}^*}$, $\underline{u}^* = \underline{u}_{\zeta^*}$ and (31). Note that $\kappa^* = \kappa_{\zeta^*,\underline{u}^*}$ and $\underline{u}^* = \underline{u}_{\zeta^*}$ is equivalent to $\kappa^* = \kappa_{\zeta^*}$. Hence, (31) holds if and only if there is $\zeta^* \in (0,1)$ such that $\overline{u}_{\zeta^*} = U(\beta_B(\kappa_{\zeta^*}))$.

This completes the proof of Lemma 3. It only remains to show (45). To see this, note that

$$\int_{\underline{u}}^{\overline{u}} (\omega - u)\, d\kappa_{\zeta,\underline{u}}(u) = \int_{\underline{u}}^{\hat{u}} (\omega - u)(\kappa_{\zeta,\underline{u}})'(u)\, du + \zeta(\omega - \overline{u}), \tag{48}$$

and inserting $\kappa_{\zeta,\underline{u}}$ from (14) and some manipulations yield

$$\int_{\underline{u}}^{\overline{u}} (\omega - u)\, d\kappa_{\zeta,\underline{u}}(u) = \frac{1}{N-1} \cdot \left(\frac{\rho_0}{1-\rho_0}\right)^{\frac{1}{N-1}} \int_{\underline{u}}^{\hat{u}} \left(\frac{\omega - \underline{u}}{\omega - u} - 1\right)^{-\frac{N-2}{N-1}} \frac{\omega - \underline{u}}{\omega - u}\, du + \zeta(\omega - \overline{u}). \tag{49}$$

A change of variables and inserting $(\omega - \overline{u}) = \rho(\zeta)(\omega - \underline{u})$ from (13) yield[39]

$$\int_{\underline{u}}^{\overline{u}} (\omega - u)\, d\kappa_{\zeta,\underline{u}}(u) = \left[\frac{1}{N-1} \cdot \left(\frac{\rho_0}{1-\rho_0}\right)^{\frac{1}{N-1}} \int_{1}^{\frac{1}{\phi(\zeta)}} (v-1)^{-\frac{N-2}{N-1}} \cdot \frac{1}{v}\, dv + \zeta\rho(\zeta)\right](\omega - \underline{u}) \tag{50}$$

$$= (1-\mu(\zeta)) \cdot (\omega - \underline{u}), \tag{51}$$

where the last equality follows from the definition of $\mu(\cdot)$ in (122). Hence, (45) follows from the fact that $\mathbf{E}(u) = \omega - \mathbf{E}(\omega - u)$. ■

Proof of Proposition 1 To see the claim, notice that for any cdf κ with $\overline{u} = \max(Supp(\kappa))$, we have for β_τ, $\tau \in \{0, B\}$:

$$U_\tau = U(\beta_\tau(\kappa)) \leq U(\beta_\tau(\delta_{\overline{u}})) \leq U(\delta_{\overline{u}}) = \overline{u} - s, \tag{52}$$

[39]Let $v(u) = \frac{\omega-u}{\omega-u}$, then $du = \frac{(\omega-u)^2}{\omega-\underline{u}} dv = \frac{\omega-u}{v}\, dv$, $v(\underline{u}) = 1$, and $v(\hat{u}) = 1/\phi(\zeta)$ from (37).

where the first inequality follows from two observations: first, that β_τ is monotone which entails that $\beta_\tau(\delta_{\overline{u}})$ fosd $\beta_\tau(\kappa)$ and monotonicity of $U(\cdot)$ with respect to fosd.[40] Similarly, the second inequality follows from monotonicity of $U(\cdot)$ and the fact that $\delta_{\overline{u}^*}$ fosd $\beta_\tau(\delta_{\overline{u}^*})$ because β_τ is pessimistic. Finally, the last equality follows from the definition of $U(\cdot)$ as explained in (12).

Now, (52) implies existence because for $\overline{u} = 0$, the conditions in part (i) of Lemma (2) are met, and it implies uniqueness because the only equilibrium type consistent with $U_B^* < \overline{u}^*$ is $\kappa^* = \delta_0$.■

Proof of Proposition 2 Recall the definitions of $\overline{u}_\zeta$ and κ_ζ. For the proof, it is useful to make the dependency on s explicit and to write $\overline{u}(\zeta,s)$ and $\kappa(\zeta,s)$ instead. Define

$$\Gamma(\zeta,s) = \overline{u}(\zeta,s) - U(\beta_B(\kappa(\zeta,s)),s). \tag{53}$$

Below, we show:

(a) Γ is continuous.

(b) $\Gamma(0,s)$ is strictly increasing in s, and $\lim_{s\to 0}\Gamma(0,s) < 0$.

(c) $\Gamma(1,s)$ is strictly increasing in s, and $\lim_{s\to\infty}\Gamma(1,s) > 0$.

(d) $\Gamma(\zeta,s)$ is strictly decreasing in ζ for all s.

Taken together, these properties imply Proposition 2. Indeed, notice first that (a)–(d) imply that there are unique $0 < s_0 < s_1$ so that

$$\Gamma(0,s_0) = 0, \quad \text{and} \quad \Gamma(1,s_1) = 0. \tag{54}$$

Too show part (i) of Proposition 2, by part (i) of Lemma 3, a Diamond-type equilibrium with utility level 0 exists if and only if $\Gamma(1,s) \geq 0$, which by (c) and (54) is equivalent to $s \geq s_1$. Moreover, a Diamond-type equilibrium with utility level $v \in (0,\overline{u}_1]$ exists if and only if $v - U(\beta_B(\delta_v)) = 0$. From additivity of β_B,

$$\begin{aligned} v - U(\beta_B(\delta_v)) = 0 \quad &\Leftrightarrow \quad v + \Delta - U(\beta_B(\delta_{v+\Delta})) = 0 \quad \text{for all } \Delta, \\ &\Leftrightarrow \quad \overline{u}_1 - U(\beta_B(\delta_{\overline{u}_1})) = 0. \end{aligned} \tag{55}$$

The latter is equivalent to $\Gamma(1,s_1) = 0$, which by (c) and (54) is equivalent to $s = s_1$.

To show part (ii) of Proposition 2, by part (ii) of Lemma 3, a Stahl-type equilibrium exists if and only if $\Gamma(0,s) \leq 0$, which by (b) and (54) is equivalent to $s \leq s_0$.

To show part (iii) of Proposition 2, by part (iii) of Lemma 3, a (unique) penny sales equilibrium exists if and only if there is a (unique) $\zeta^* \in (0,1)$ with $\Gamma(\zeta^*,s) = 0$, which by (a)–(d) and (54) is equivalent to $s \in (s_0,s_1)$.

---

[40] As is well-known in the search literature, the definition of $U(\cdot)$ as given in (4) entails that $U(\beta) \geq U(\beta')$ if β fosd β'. Moreover, for every β, $U(\beta)$ strictly decreases in s with $\lim_{s\to 0} U(\beta) = \max(Supp(\beta))$ and $\lim_{s\to\infty} U(\beta) = -\infty$. We omit the details.

To complete the proof, it remains to show (a)–(d). Because $\overline{u}(\zeta,s)$ is clearly continuous, property (a) follows if $U(\beta_B(\kappa(\zeta,s)),s)$ is continuous in (ζ,s), which, by continuity of β_B, is the case if $U(\kappa(\zeta,s),s)$ is continuous in (ζ,s). From (46) and (13),

$$U(\kappa(\zeta,s),s) = -s + \mu(\zeta)\omega + (1-\mu(\zeta)) \max\left\{0, \omega - \frac{s}{\mu(\zeta)}\right\}, \tag{56}$$

which is continuous in (ζ,s), because $\mu(\cdot)$ is continuous with $\mu(\zeta) \geq \mu_0 > 0$.

As to (b). To see the first part of the claim, recall from (13),

$$\overline{u}(0,s) = (1-\rho_0)\omega + \rho_0 \max\left\{0, \omega - \frac{s}{\mu_0}\right\}. \tag{57}$$

Therefore, for $s' < s$, we have that

$$\overline{u}(0,s) - \overline{u}(0,s') \geq -\frac{\rho_0}{\mu_0}(s-s'). \tag{58}$$

Below we show that

$$U(\beta_B(\kappa(0,s)),s) - U(\beta_B(\kappa(0,s')),s') \quad \leq \quad -(s-s'). \tag{59}$$

Together with (53), (58) and (59) yield that

$$\Gamma(0,s) - \Gamma(0,s') \quad \geq \quad -\frac{\rho_0}{\mu_0}(s-s') + s - s', \tag{60}$$

which is positive since $\mu_0 > \rho_0$ by assumption. This establishes that $\Gamma(0,s)$ strictly increases in s.

To show (59), notice that $\kappa(0,s')$ first order stochastically dominates $\kappa(0,s)$ (compare Proposition 7 Stahl (1989)). Together with monotonicity of β_B, we conclude that

$$\beta_B(\kappa(0,s')) \quad fosd \quad \beta_B(\kappa(0,s)), \tag{61}$$

which entails that

$$U(\beta_B(\kappa(0,s)),s) < U(\beta_B(\kappa(0,s')),s'), \tag{62}$$

because $U(\beta,s)$ increases in β if β increases in the fosd sense and decreases in s by Footnote 40. From definition (4), we obtain therefore that

$$U(\beta_B(\kappa(0,s)),s) = -s + \int \max\{u, U(\beta_B(\kappa(0,s),s)\} \, d\beta_B(\kappa(0,s))(u) \tag{63}$$

$$\leq -(s-s') - s' + \int \max\{u, U(\beta_B(\kappa(0,s'),s')\} \, d\beta_B(\kappa(0,s'))(u), \tag{64}$$

where the inequality follows from the fact that $\max\{u, U\}$ increases in u and U in combination with (61) and (62). This establishes (59), as (64) is equivalent to (59) by (4).

To see that $\lim_{s\to 0}\Gamma(0,s) < 0$, recall from (13) that $\underline{u}(0,s) = \max\left\{0, \omega - \frac{s}{\mu_0}\right\}$ and hence $\lim_{s\to 0}\underline{u}(0,s) = \omega$. Since, in addition, $\overline{u}(0,s) \leq \omega$ for all s, we conclude that $\lim_{s\to 0}\kappa(0,s)(u) = \delta_\omega(u)$ for all u and $\lim_{s\to 0}\overline{u}(0,s) = \omega$. Since $\Gamma(0,s)$ is continuous in s by (a):

$$\lim_{s\to 0} U(\beta_B(\kappa(0,s)),s) = U(\beta_B(\kappa(0,0)),0) = \max(Supp(\beta_B(\delta_\omega))), \tag{65}$$

where the second equality follows by assumption.[41] Thus,

$$\lim_{s\to 0}\Gamma(0,s) = \lim_{s\to 0}[\overline{u}(0,s) - U(\beta_B(\kappa(0,s)),s)] = \omega - \max(Supp(\beta_B(\delta_\omega))). \tag{66}$$

To complete the proof, it suffices to show that $\omega - \max(Supp(\beta_B(\delta_\omega))) < 0$.

Indeed, because β_B is non-pessimistic, there is $\hat{\kappa}$ with

$$\hat{\overline{u}} = \max(Supp(\hat{\kappa})) < \max(Supp(\beta_B(\hat{\kappa}))). \tag{67}$$

Because β_B is monotone and $\delta_{\hat{\overline{u}}}$ fosd $\hat{\kappa}$, we have that $\beta_B(\delta_{\hat{\overline{u}}})$ fosd $\beta_B(\hat{\kappa})$, and thus $\max(Supp(\beta_B(\hat{\kappa}))) \leq \max(Supp(\beta_B(\delta_{\hat{\overline{u}}})))$. Thus, by (67),

$$\max(Supp(\delta_{\hat{\overline{u}}})) < \max(Supp(\beta_B(\delta_{\hat{\overline{u}}}))). \tag{68}$$

Let $\Delta = \omega - \hat{\overline{u}}$ and notice that $\max(Supp(\beta_B(\delta_{\hat{\overline{u}}}))) = U(\beta_B(\delta_{\hat{\overline{u}}}),0)$ by Footnote 40. Hence, from (68),

$$\omega < U(\beta_B(\delta_{\hat{\overline{u}}}),0) + \Delta = U(\beta_B(\delta_\omega),0) = \max(Supp(\beta_B(\delta_\omega))), \tag{69}$$

where the second equality follows from additivity of β_B, hence the desired inequality.

As to (c). To show that $\Gamma(1,s)$ is strictly increasing in s, let $s' < s$ and recall that $\kappa(1,s) = \delta_{\overline{u}(1,s)}$. Because $U(\beta,s)$ strictly decreases in s for every β by Footnote 40:

$$\Gamma(1,s) = \overline{u}(1,s) - U(\beta_B(\delta_{\overline{u}(1,s)}),s) > \overline{u}(1,s) - U(\beta_B(\delta_{\overline{u}(1,s)}),s'). \tag{70}$$

Let $\Delta = \overline{u}(1,s) - \overline{u}(1,s')$. By additivity of β_B, the right hand side is equal to

$$\overline{u}(1,s) - U(\beta_B(\delta_{\overline{u}(1,s')+\Delta}),s') = u(1,s) - \left\{U(\beta_B(\delta_{\overline{u}(1,s')}),s') + \Delta\right\} \tag{71}$$

$$= \overline{u}(1,s') - U(\beta_B(\delta_{\overline{u}(1,s')}),s') = \Gamma(1,s'), \tag{72}$$

as desired.

To see that $\lim_{s\to\infty}\Gamma(1,s) > 0$, observe that for $\Delta = \overline{u}(1,s)$, additivity of β_B delivers

$$\Gamma(1,s) = \overline{u}(1,s) - U(\beta_B(\delta_\Delta),s) = \overline{u}(1,s) - \left\{U(\beta_B(\delta_0,s) + \Delta\right\} = -U(\beta_B(\delta_0),s). \tag{73}$$

[41] For $s = 0$, let $U(\beta) \equiv \max(Supp(\beta))$ consistent with $\lim_{s\to 0} U(\beta) = \max(Supp(\beta))$.

From Footnote 40, $U(\beta(\delta_0),s)$ goes to $-\infty$ as $s \to \infty$. Hence, $\lim_{s\to\infty}\Gamma(1,s) > 0$, as desired.

As to (d). Let $\zeta' < \zeta$. We have to show that $\Gamma(\zeta',s) > \Gamma(\zeta,s)$ for all s. Let α denote the level of utility that satisfies $\kappa(\zeta',s)(\alpha) = 1-\zeta$. Denote by ψ the cdf which coincides with $\kappa(\zeta',s)$ up to α and has a mass point of size ζ at $\overline{u}(\zeta',s)$:

$$\psi(u) = \begin{cases} \kappa(\zeta',s)(u) & if \quad u < \alpha \\ 1-\zeta & if \quad u \in [\alpha, \overline{u}(\zeta',s)) \\ 1 & if \quad u \geq \overline{u}(\zeta',s). \end{cases} \tag{74}$$

Figure 7 illustrates how ψ relates to $\kappa(\zeta,s)$ and $\kappa(\zeta',s)$. In particular, it shows that by construction, ψ first order stochastically dominates $\kappa(\zeta',s)$. Because ψ fosd $\kappa(\zeta',s)$, monotonicity of β_B

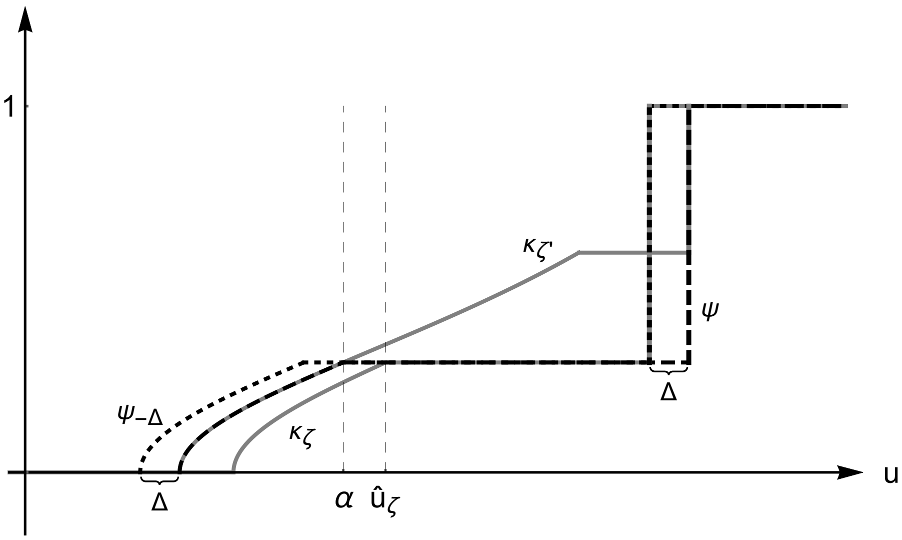

Figure 7: *The utility distributions $\kappa_\zeta = \kappa(\zeta,s)$ and $\kappa_{\zeta'} = \kappa(\zeta',s)$ (solid, gray) as well as the auxiliary functions ψ (dashed) and $\psi_{-\Delta}$ (dotted). Notice that $\psi_{-\Delta}$ is identical to ψ when shiftet by Δ to the "left". Moreover, the critical values α and $\hat{u}_\zeta = \hat{u}(\zeta,s)$.*

implies that[42]

$$\Gamma(\zeta',s) = \overline{u}(\zeta',s) - U(\beta_B(\kappa(\zeta',s))) > u(\zeta',s) - U(\beta_B(\psi)). \tag{77}$$

[42]To see that the inequality is strict, we will show that

$$U(\psi,s') > U(\kappa(\zeta',s),s') \quad \text{for all } s' > 0. \tag{75}$$

Together with monotonicity of β_B, this implies (77).

To see (75), note that by (12), we have for every $s' > 0$ both $U(\psi,s') < \overline{u}(\zeta',s)$ and $U(\kappa(\zeta',s),s') < \overline{u}(\zeta',s)$. Moreover,

$$\int \max\{u,U\}\, d\psi > \int \max\{u,U\}\, d\kappa(\zeta',s) \quad \text{for every } U < \overline{u}(\zeta',s), \tag{76}$$

Further, for $\Delta = \overline{u}(\zeta', s) - \overline{u}(\zeta, s) > 0$, let $\psi_{-\Delta}(u) = \psi(u + \Delta)$ be the cdf that results from shifting ψ by Δ to the left. Hence, additivity of β_B entails that

$$\overline{u}(\zeta', s) - U(\beta_B(\psi)) = \overline{u}(\zeta, s) - U(\beta_B(\psi_{-\Delta})). \tag{78}$$

Below, we will argue that $\kappa(\zeta, s)$ fosd $\psi_{-\Delta}$ which together with monotonicity of β_B and the monotonicity of $U(\cdot)$ with respect to fosd implies that

$$U(\beta_B(\psi_{-\Delta})) \le U(\beta_B(\kappa(\zeta, s))). \tag{79}$$

Putting all inequalities together delivers:

$$\Gamma(\zeta', s) > \overline{u}(\zeta, s) - U(\beta_B(\kappa(\zeta, s))) = \Gamma(\zeta, s), \tag{80}$$

as desired.

To see that $\kappa(\zeta, s)$ fosd $\psi_{-\Delta}$, note that by construction (compare Figure 7)

$$\kappa(\zeta, s)(u) \le \psi_{-\Delta}(u) \quad \text{if } u < \underline{u}(\zeta, s), \tag{81}$$
$$\kappa(\zeta, s)(u) < \psi_{-\Delta}(u) \quad \text{if } u \in [\underline{u}(\zeta, s), \hat{u}(\zeta, s)), \tag{82}$$
$$\kappa(\zeta, s)(u) = \psi_{-\Delta}(u) \quad \text{if } u \ge \hat{u}(\zeta, s), \tag{83}$$

where (81) follows from $\kappa(\zeta, s)(u) = 0$ for $u < \underline{u}(\zeta, s)$ and (83) holds by construction. To see (82), because $\kappa(\zeta, s)(u) < \min\{\kappa(\zeta', s)(u), 1 - \zeta\}$ for all $u \in [\underline{u}(\zeta, s), \hat{u}(\zeta, s))$, we have $\kappa(\zeta, s)(u) < \psi(u)$ for all $u \in [\underline{u}(\zeta, s), \hat{u}(\zeta, s))$. From this (82) follows, because $\psi_{-\Delta}(u) \ge \psi(u)$ for all u by construction. This completes the proof. ■

Proof of Proposition 3 As argued in the text, it suffices to show that the equilibrium is as described in Proposition 2 and the cut-off values s_0 and s_1 are given by (22) and strictly increasing in χ. Define

$$\Gamma_B(\zeta, s) \equiv s + \chi \left[(1 - \rho(\zeta) - \mu(\zeta)) \min\left\{ \omega, \frac{s}{\mu(\zeta)} \right\} - \upsilon \right]. \tag{84}$$

As we show below, for all $\zeta \in [0, 1]$, we have

$$\Gamma_B(\zeta, s) \gtreqless 0 \quad \Leftrightarrow \quad \Gamma(\zeta, s) \gtreqless 0. \tag{85}$$

Following along the same lines as the first part of the proof of Proposition 2, it follows from (85) and the fact that (3) satisfies Assumption (2) that the equilibrium is as described in Proposition 2 if:

because ψ fosd $\kappa(\zeta', s)$ and the mass point at $\overline{u}(\zeta', s)$ of ψ contains more mass than the mass point at $\overline{u}(\zeta', s)$ of $\kappa(\zeta', s)$ as $\zeta > \zeta'$. Together, this implies (75) by definition (4).

(a') Γ_B is continuous.

(b') $\Gamma_B(0,s)$ is strictly increasing in s, and $\lim_{s\to 0}\Gamma_B(0,s) < 0$.

(c') $\Gamma_B(1,s)$ is strictly increasing in s, and $\lim_{s\to\infty}\Gamma_B(1,s) > 0$.

(d') $\Gamma_B(\zeta,s)$ is strictly decreasing in ζ for all s.

Notice that (a') - (d') follow trivially from the fact that $\Gamma_B(\zeta,s)$ is strictly increasing in s for all $\zeta\in[0,1]$, because $1-\rho(\zeta)-\mu(\zeta)\geq 0$ for all $\zeta\in[0,1]$ from Lemma 4, and the fact that $\Gamma_B(\zeta,s)$ is strictly decreasing in ζ for all s, as $\mu(\cdot)$ and $\rho(\cdot)$ strictly increase in ζ from Lemma 4.

Next, from (54), the cut-off values s_0 and s_1 are implicitly given by $\Gamma(0,s_0)=0$ and $\Gamma(1,s_1)=0$. From (85), this is equivalent to $\Gamma_B(0,s_0)=0$ and $\Gamma_B(1,s_1)=0$ which implies that s_0 and s_1 are indeed given by (22), because $(1-\rho(1)-\mu(1))=0$ from Lemma 4. Finally, s_0 and s_1 are strictly increasing in χ, because $\upsilon>\omega$ by assumption and $1-\rho_0-\mu_0\in(0,1)$ from Lemma 4.

To complete the proof, it remains to show (85). To do so, we establish first the following auxiliary result:

$$\Gamma_B(\zeta,s)=\overline{u}(\zeta,s)-R\big(\overline{u}(\zeta,s),\beta_B(\kappa(\zeta,s))\big). \tag{86}$$

To see this, by the definition of $R(U,\beta)$ from (5) and β_B from (3), we have

$$R\big(\overline{u}(\zeta,s),\beta_B(\kappa(\zeta,s))\big)=-s+\int\max\{u,\overline{u}(\zeta,s)\}\,d\beta_B(\kappa(\zeta,s))(u) \tag{87}$$

$$=-s+[1-2\chi]\int\max\{u,\overline{u}(\zeta,s)\}\,d\kappa(\zeta,s)(u) \tag{88}$$

$$+\chi\int\max\{\upsilon+u,\overline{u}(\zeta,s)\}\,d\kappa(\zeta,s)(u)$$

$$+\chi\int\max\{-\upsilon+u,\overline{u}(\zeta,s)\}\,d\kappa(\zeta,s)(u).$$

From (13), we have $\overline{u}(\zeta,s)-\underline{u}(\zeta,s)<\omega$. Together with $\omega<\upsilon$ by assumption, this implies $\overline{u}(\zeta,s)-\underline{u}(\zeta,s)<\upsilon$. From this, it follows that $\upsilon+u>\overline{u}(\zeta,s)$ and $-\upsilon+u<\overline{u}(\zeta,s)$ for all u in the support of $\kappa(\zeta,s)$. Hence, (88) is equivalent to

$$R\big(\overline{u}(\zeta,s),\beta_B(\kappa(\zeta,s))\big)=-s+[1-2\chi]\,\overline{u}(\zeta,s)+\chi\left(\upsilon+\int u\,d\kappa(\zeta,s)(u)\right)+\chi\overline{u}(\zeta,s) \tag{89}$$

$$=-s+\overline{u}(\zeta,s)+\chi\left(-\overline{u}(\zeta,s)+\upsilon+\mu(\zeta)\omega+(1-\mu(\zeta))\,\underline{u}(\zeta,s)\right), \tag{90}$$

where the last equality follows from (45) and re-arranging terms. Inserting (90) in (86) yields

$$\Gamma_B(\zeta,s)=s+\chi\left[\overline{u}(\zeta,s)-\upsilon-\mu(\zeta)\omega-(1-\mu(\zeta))\,\underline{u}(\zeta,s)\right]. \tag{91}$$

Inserting $\overline{u}(\zeta,s) = (1-\rho(\zeta))\,\omega + \rho(\zeta)\,\underline{u}(\zeta,s)$ from (13) yields

$$\Gamma_B(\zeta,s) = s + (1-\rho(\zeta)-\mu(\zeta))\omega - \upsilon - (1-\rho(\zeta)-\mu(\zeta))\,\underline{u}(\zeta,s) \tag{92}$$

which is equivalent to (84), because $\underline{u}(\zeta,s) = \max\left\{0, \omega - \frac{s}{\mu(\zeta)}\right\}$ from (13).

Finally, to show the equivalence (85), from (86), we have

$$\begin{aligned} \Gamma_B(\zeta,s) \gtreqless 0 \quad &\Leftrightarrow \quad \overline{u}(\zeta,s) \gtreqless R\big(\overline{u}(\zeta,s), \beta_B(\kappa(\zeta,s))\big) \\ &\Leftrightarrow \quad \overline{u}(\zeta,s) \gtreqless U(\beta_B(\kappa(\zeta,s)),s) \\ &\Leftrightarrow \quad \Gamma(\zeta,s) \gtreqless 0, \end{aligned}$$

where the second equivalence follows from the fact that $U(\beta_B(\kappa(\zeta,s)),s)$ is a solution to

$$U = R\big(U, \beta_B(\kappa(\zeta,s))\big)$$

and $(\partial R(U,\beta))/(\partial U) \leq 1$ for all β. ■

Proof of Proposition 4 Denote by W_0, W_B, π^*, W_T, respectively, the equilibrium values of standard and biased consumer welfare, profits and total welfare. We first argue that in the Diamond and Stahl regime, for $\chi \in [0,\chi_0) \cup (\chi_1, 1/4]$, these values are independent of χ.

In the Diamond regime, for $\chi \in [0,\chi_0)$, both standard and biased consumers buy from the first firm they visit which supplies zero utility. Because the first search is free, both consumer types obtain zero utility while firms share the entire surplus from trade. Total welfare equals ω, as trade occurs with probability one and no consumer engages in costly search:

$$W_0 = W_B = 0, \quad \pi^* = \frac{\omega}{N} \quad \text{and} \quad W_T = \omega, \quad \text{when } \chi \in [0,\chi_0). \tag{93}$$

In the Stahl regime, for $\chi \in [\chi_1, \frac{1}{4})$, W_0, W_B, π^*, W_T are independent of χ, because first, $\kappa^* = \kappa_0$, and κ_0 is independent of χ, and second, irrespective of χ, standard consumers buy from the first firm while biased consumers visit all firms.

Next, we consider the transition points χ_1 and χ_0 and show first that W_0, W_B, π^*, W_T are continuous at χ_1. As we show below,

$$\frac{d\zeta^*(\chi)}{d\chi} < 0 \quad \text{for } \chi \in (\chi_0,\chi_1) \quad \text{and} \quad \lim_{\chi \downarrow \chi_0} \zeta^* = 1 \quad \text{and} \quad \lim_{\chi \uparrow \chi_1} \zeta^* = 0. \tag{94}$$

The facts that $\lim_{\chi \uparrow \chi_1} \zeta^* = 0$, and that $\kappa^* = \kappa_{\zeta^*}$ in the penny sales regime imply that $\lim_{\chi \uparrow \chi_1} \kappa^* = \kappa_0$. Moreover, as $\chi \uparrow \chi_1$, biased consumers visit all firms with probability one, just as in the Stahl regime, because from (94), the probability ζ^* for a penny sale that induces them to stop and buy goes to zero. Finally, in both regimes standard consumers buy from the first firm they visit in equilibrium. Therefore, W_0, W_B, π^*, W_T are continuous at χ_1.

Next, we show that W_0, W_B, π^* are discontinuous from below at χ_0, as $\chi \uparrow \chi_0$. Recall that, by assumption, for $\chi = \chi_0$ we focus on the consumer optimal equilibrium which is a Diamond-type equilibrium with $\kappa^* = \delta_{\overline{u}_1^*}$. From the same arguments used to derive (93), we obtain for $\chi = \chi_0$:

$$W_0 = W_B = \overline{u}_1, \quad \pi^* = \frac{\omega - \overline{u}_1}{N} \quad \text{and} \quad W_T = \omega, \quad \text{when } \chi = \chi_0. \tag{95}$$

Therefore, from (93) and (95), W_0 and W_B discontinuously increase at χ_0, because $\overline{u}_1 > 0$, while π^* discontinuously decreases. Notice also that total welfare remains constant.

In contrast, W_0, W_B, π^*, W_T are continuous from above at χ_0, as $\chi \downarrow \chi_0$. This follows immediately from the fact that $\lim_{\chi \downarrow \chi_0} \zeta^* = 1$ by (94) which, in particular, entails that $\lim_{\chi \downarrow \chi_0} \kappa^* = \delta_{\overline{u}_1}$.

These considerations imply that the proposition follows from the following comparative statics results for the penny sales regime. For the rest of the proof, let $\chi \in (\chi_0, \chi_1)$. We show: As χ increases,

(a) W_0 decreases,
(b) W_B decreases when $\omega \geq \frac{s}{\mu_0}$,
(c) π^* increases, and
(d) W_T decreases.

To show (a), because standard consumers buy from the first firm they visit, and as the first search is for free, we have $W_0 = \mathbf{E}(u)$ and with (45), we obtain:

$$W_0 = \mu(\zeta^*)\omega + (1 - \mu(\zeta^*))\, \underline{u}^* = \omega - (1 - \mu(\zeta^*)) \cdot \min\left\{\omega, \frac{s}{\mu(\zeta^*)}\right\}, \tag{96}$$

where the second equality follows from inserting (13) for $\underline{u}^*$ and straightforward algebra. Because $\mu(\cdot)$ is increasing by Lemma 4, and ζ^* decreases in χ by (94), W_0 decreases in χ.

To show (b), it is useful to decompose biased consumer welfare into the expected consumption value W_B^U which the consumer attains from accepting an offer, and the expected total search costs W_B^S which the consumer incurs in equilibrium:

$$W_B = W_B^U - W_B^S. \tag{97}$$

To show part (b), it is therefore sufficient to show:

$$\frac{dW_B^U}{d\chi} = 0 \text{ if } \omega \geq \frac{s}{\mu_0}, \quad \text{and} \quad \frac{dW_B^S}{d\chi} > 0. \tag{98}$$

To show the first part of (98), in equilibrium trade occurs with probability one and the surplus from trade ω is split among consumers and firms. Hence,

$$\gamma \cdot W_B^U + (1-\gamma) \cdot W_0 + N \cdot \pi^* = \omega. \tag{99}$$

From (96) and (112), it follows that

$$\begin{aligned} W_B^U &= \frac{1}{\gamma}\left[\omega - (1-\gamma) \cdot W_0 - N \cdot \pi^*\right] \\ &= \frac{1}{\gamma}\left[\omega - (1-\gamma)\omega + (1-\gamma) \cdot (1-\mu(\zeta^*)) \cdot \min\left\{\omega, \frac{s}{\mu(\zeta^*)}\right\} - (1-\gamma) \cdot \min\left\{\omega, \frac{s}{\mu(\zeta^*)}\right\}\right] \\ &= \omega - \frac{1-\gamma}{\gamma} \cdot \mu(\zeta^*) \cdot \min\left\{\omega, \frac{s}{\mu(\zeta^*)}\right\}. \end{aligned} \tag{100}$$

Observe that this expression depends on χ only indirectly through ζ^*. Differentiating (100) yields

$$\frac{\partial W_B^U(\zeta^*)}{\partial \zeta} = \begin{cases} -\frac{1-\gamma}{\gamma}\mu'(\zeta^*)\omega < 0 & \text{if} \quad \omega < \frac{s}{\mu(\zeta^*)} \\ 0 & \text{otherwise,} \end{cases} \tag{101}$$

where the inequality follows since $\mu(\cdot)$ increasing by Lemma 4. Therefore, since $\partial W_B^U/\partial \chi = \partial W_B^U/\partial \zeta \cdot \partial \zeta^*/\partial \chi$ and $\partial \zeta^*/\partial \chi < 0$ from (94), the first part of (98) follows from the fact that $\omega \geq \frac{s}{\mu_0}$ implies $\omega \geq \frac{s}{\mu(\zeta^*)}$, as $\mu(\cdot)$ is increasing with $\mu(0) = \mu_0$ by Lemma 4.

To show the second part of (98), let i denote the number of searches that a biased consumer conducts until she encounters a firm which makes a penny sales offer. The probability that she visits exactly i firms to find one that makes a penny sales offer is

$$\zeta^* \cdot (1-\zeta^*)^{i-1}.$$

In this case, she incurs the total search costs $(i-1)s$, as her first search is free. On the contrary, with probability $(1-\zeta^*)^N$ she does not encounter any firm that makes a penny sales offer (as no firm offers one). Then, she visits all firms and incurs the total search costs $(N-1)s$. Together, it follows that

$$\frac{W_B^S}{s} = \left[\sum_{i=1}^{N}(i-1) \cdot \zeta^* \cdot (1-\zeta^*)^{i-1}\right] + (N-1) \cdot (1-\zeta^*)^N \tag{102}$$

$$= \zeta^* \cdot \sum_{i=1}^{N-1} i \cdot (1-\zeta^*)^i + (N-1) \cdot (1-\zeta^*)^N \tag{103}$$

$$= \frac{1-\zeta^*}{\zeta^*} \cdot [1-(1-\zeta^*)^{N-1}], \tag{104}$$

where the final equality follows from a straightforward calculation.[43]

[43] Recall that

$$\sum_{i=1}^{N-1} i \cdot (1-\zeta^*)^i = \frac{\zeta^*(1-\zeta^*)^N - \zeta^* N(1-\zeta^*)^N - (1-\zeta^*)^N + (1-\zeta^*)}{(\zeta^*)^2}, \tag{105}$$

Again, this expression depends on χ only through ζ^*, and by (94) the second part of (98) is shown if $\partial W_B^S(\zeta^*)/\partial\zeta < 0$ for all $\zeta^* \in (0,1)$. In fact,

$$\frac{\partial W_B^S(\zeta^*)}{\partial\zeta} = -\frac{s}{(\zeta^*)^2}\cdot[1-(1-\zeta^*)^{N-1}] + s\frac{1-\zeta^*}{\zeta^*}(N-1)(1-\zeta^*)^{N-2} \tag{108}$$

$$= -s\cdot\frac{1-(1-\zeta^*)^{N-1}-\zeta^*(N-1)(1-\zeta^*)^{N-1}}{(\zeta^*)^2}. \tag{109}$$

This is negative if the denominator

$$h(\zeta^*) \equiv 1-(1-\zeta^*)^{N-1}-\zeta^*(N-1)(1-\zeta^*)^{N-1} \tag{110}$$

is positive. To see this, observe that $h(0)=0$, and

$$h'(\zeta^*) = (N-1)\cdot(1-\zeta^*)^{N-2}\cdot[1-(1-\zeta^*)+\zeta^*(N-1)]$$

$$= (N-1)\cdot(1-\zeta^*)^{N-2}\cdot N\zeta^* > 0. \tag{111}$$

This completes the proof of part (b).

To show (c), a firm's profit in a penny sales equilibrium is equal to the profit of a firm that offers the lowest utility which is given as the product of the demand it derives from standard consumers $(1-\gamma)/N$ and its margin $\omega-\underline{u}^*$. Inserting from (13) for $\underline{u}^*$ yields

$$\pi^* = \frac{1-\gamma}{N}\cdot\min\left\{\omega, \frac{s}{\mu(\zeta^*)}\right\}. \tag{112}$$

Because $\mu(\cdot)$ is increasing from Lemma 4 and ζ^* decreases in χ by (94), π^* increases in χ.

To show (d), because trade occurs with probability one and no standard consumer engages in costly search, we have

$$W_T = \omega - \gamma W_B^S, \tag{113}$$

and (d) follows from the fact W_B^S increases in χ as argued above.

To complete the proof, it remains to show (94). For fixed s, let

$$\tilde{\Gamma}_B(\zeta,\chi) \equiv \frac{\Gamma_B(\zeta,s)}{\chi} = \frac{s}{\chi} + (1-\rho(\zeta)-\mu(\zeta))\min\left\{\omega, \frac{s}{\mu(\zeta)}\right\} - v \tag{114}$$

---

and thus,

$$\frac{W_B^S}{s} = \zeta^*\cdot\frac{\zeta^*(1-\zeta^*)^N-\zeta^*N(1-\zeta^*)^N-(1-\zeta^*)^N+(1-\zeta^*)}{(\zeta^*)^2} + (N-1)\cdot(1-\zeta^*)^N \tag{106}$$

$$= \zeta^*\cdot\frac{-(1-\zeta^*)^N+(1-\zeta^*)}{(\zeta^*)^2}, \tag{107}$$

which is equivalent to (104).

with $\Gamma_B(\zeta,s)$ as given in (84). Recall that ζ^* is defined by the equilibrium condition $\Gamma(\zeta^*,s)=0$. Hence, the equivalence (85) implies that $\tilde{\Gamma}_B(\zeta^*,\chi)=0$ for all χ in equilibrium, and it follows that

$$\frac{d\zeta^*}{d\chi}=-\frac{\partial\tilde{\Gamma}_B(\zeta^*,\chi)}{\partial\chi}\Big/\frac{\partial\tilde{\Gamma}_B(\zeta^*,\chi)}{\partial\zeta}. \tag{115}$$

Differentiating (114) yields

$$\frac{\partial\tilde{\Gamma}_B(\zeta^*,\chi)}{\partial\chi}=-\frac{s}{\chi^2}<0. \tag{116}$$

Moreover,

$$\frac{\partial\tilde{\Gamma}_B(\zeta^*,\chi)}{\partial\zeta}=\begin{cases}-[\rho'(\zeta^*)+\mu'(\zeta^*)]\,\omega<0 & \text{if}\quad \omega<\frac{s}{\mu(\zeta^*)}\\ -\left[\frac{\rho'(\zeta^*)}{\mu(\zeta^*)}+\frac{1-\rho(\zeta^*)}{\mu^2(\zeta^*)}\mu'(\zeta^*)\right]s<0 & \text{otherwise,}\end{cases} \tag{117}$$

where the inequalities follow from $\mu'(\zeta^*)>0$ and $\rho'(\zeta^*)>0$ by Lemma 4. Hence, $d\zeta^*\big/d\chi<0$.

Finally, $\lim_{\chi\downarrow\chi_0}\zeta^*=1$ and $\lim_{\chi\uparrow\chi_1}\zeta^*=0$, because ζ^* is unique and implicitly given by $\tilde{\Gamma}_B(\zeta^*,\chi)=0$, and $\tilde{\Gamma}_B$ is continuous in ζ and χ with $\tilde{\Gamma}_B(1,\chi_0)=0$ and $\tilde{\Gamma}_B(0,\chi_1)=0$. This completes the proof. ■

8 Appendix B

Recall the definition of $\eta(\cdot)$ in (35). Define

$$\rho_0 \equiv \frac{1-\gamma}{N\gamma+1-\gamma}, \tag{118}$$

$$\mu_0 \equiv 1-\frac{1}{N-1}\cdot\left(\frac{\rho_0}{1-\rho_0}\right)^{\frac{1}{N-1}}\int_1^{\frac{1}{\rho_0}}\frac{(\nu-1)^{-\frac{N-2}{N-1}}}{\nu}\,d\nu, \tag{119}$$

$$\rho(\zeta) \equiv \frac{\rho_0}{\rho_0+(1-\rho_0)\,\eta(\zeta)}, \tag{120}$$

$$\phi(\zeta) \equiv \frac{\rho_0}{\rho_0+(1-\rho_0)(1-\zeta)^{N-1}}, \tag{121}$$

$$\mu(\zeta) \equiv 1-\frac{1}{N-1}\cdot\left(\frac{\rho_0}{1-\rho_0}\right)^{\frac{1}{N-1}}\int_1^{\frac{1}{\phi(\zeta)}}\frac{(\nu-1)^{-\frac{N-2}{N-1}}}{\nu}\,d\nu-\zeta\rho(\zeta). \tag{122}$$

Lemma 4 *We have:*
(i) $\eta(\cdot)$ is strictly decreasing with $\eta(0)=1$ and $\eta(1)=\frac{1}{N}$.
(ii) $\rho(\cdot)$ is strictly increasing with $\rho(0)=\rho_0$ and $\rho(1)=1-\gamma$.
(iii) $\phi(\cdot)$ is strictly increasing with $\phi(0)=\rho_0$ and $\phi(1)=1$. Moreover, $\phi(\zeta)>\rho(\zeta)$ for all $\zeta\in(0,1]$.
(iv) $\mu(\cdot)$ is strictly increasing with $\mu(0)=\mu_0>0$ and $\mu(1)=\gamma$. Moreover, $1-\rho(\zeta)-\mu(\zeta)>0$ for all $\zeta\in(0,1)$ and $1-\rho(1)-\mu(1)=0$.
(v) For all $\zeta\in[0,1]$, $\lim_{\gamma\to 0}\rho(\zeta)=1$, $\lim_{\gamma\to 1}\rho(\zeta)=0$, $\lim_{\gamma\to 0}\mu(\zeta)=0$ and $\lim_{\gamma\to 1}\mu(\zeta)=1$.

Proof of Lemma 4 (i) Recall from definition (35) that

$$\eta(\zeta)=\sum_{t=1}^{N}\binom{N-1}{t-1}\frac{\zeta^{t-1}(1-\zeta)^{N-t}}{t}, \tag{123}$$

with $\eta(0)=1$ and $\eta(1)=\frac{1}{N}$. Differentiating (123) yields

$$\eta'(\zeta)=\sum_{t=2}^{N}\binom{N-1}{t-1}\frac{t-1}{t}\zeta^{t-2}(1-\zeta)^{N-t}-\sum_{t=1}^{N-1}\binom{N-1}{t-1}\frac{N-t}{t}\zeta^{t-1}(1-\zeta)^{N-t-1}. \tag{124}$$

By an index change, the second term on the right hand side can be written as[44]

$$\sum_{t=2}^{N}\binom{N-1}{t-2}\frac{N-t+1}{t-1}\zeta^{t-2}(1-\zeta)^{N-t}=\sum_{t=2}^{N}\binom{N-1}{t-1}\zeta^{t-2}(1-\zeta)^{N-t}, \tag{126}$$

[44]Note that

$$\binom{N-1}{t-2}=\frac{(N-1)!}{(t-2)!(N-t+1)!}=\frac{(N-1)!}{(t-1)!(N-t)!}\frac{t-1}{N-t+1}=\binom{N-1}{t-1}\frac{t-1}{N-t+1}. \tag{125}$$

Therefore,

$$\eta'(\zeta) = \sum_{t=2}^{N} \binom{N-1}{t-1} \left[\frac{t-1}{t} - 1\right] \zeta^{t-2}(1-\zeta)^{N-t} \tag{127}$$

$$= -\frac{1}{\zeta} \sum_{t=2}^{N} \binom{N-1}{t-1} \cdot \frac{1}{t} \cdot \zeta^{t-1}(1-\zeta)^{N-t} \tag{128}$$

$$= -\frac{1}{\zeta}\left[\eta(\zeta) - (1-\zeta)^{N-1}\right] \tag{129}$$

$$= -\frac{1}{\zeta}\frac{\rho_0}{1-\rho_0}\left[\frac{1}{\rho(\zeta)} - \frac{1}{\phi(\zeta)}\right] < 0, \tag{130}$$

where the third line follows from the definition of $\eta(\cdot)$ in (123) and the last line from the definition of $\rho(\cdot)$ in (120) and $\phi(\cdot)$ in (121). Finally, $\eta(\cdot)$ strictly decreasing follows from (128).

(ii) Because $\eta(0) = 1$ and $\eta(1) = \frac{1}{N}$ from part (i), we have

$$\rho(0) = \rho_0 \quad \text{and} \quad \rho(1) = 1-\gamma, \tag{131}$$

where the last equality follows straightforwardly from the definition of ρ_0 in (118) and rearranging terms. Moreover, differentiating (120) yields

$$\rho'(\zeta) = -\frac{\rho_0}{(\rho_0 + (1-\rho_0)\,\eta(\zeta))^2} \cdot (1-\rho_0)\,\eta'(\zeta) \tag{132}$$

$$= -\frac{1-\rho_0}{\rho_0} \cdot \rho(\zeta)^2 \cdot \eta'(\zeta) \tag{133}$$

$$= \frac{\rho(\zeta)^2}{\zeta} \cdot \left[\frac{1}{\rho(\zeta)} - \frac{1}{\phi(\zeta)}\right], \tag{134}$$

where the last equality follows from inserting η' from (130). Moreover, $\rho(\cdot)$ strictly increasing follows from (133) and $\eta(\cdot)$ strictly decreasing by part (i).

(iii) Observe that $\phi(0) = \rho_0$ and $\phi(1) = 1$. Moreover, differentiating (121) yields

$$\phi'(\zeta) = \frac{1-\rho_0}{\rho_0} \cdot \phi(\zeta)^2 \cdot (N-1)(1-\zeta)^{N-2} > 0, \tag{135}$$

as desired. Finally, $\phi(\zeta) > \rho(\zeta)$ for all $\zeta \in (0,1]$ follows immediately from equation (130) in the proof of part (i).

(iv) Observe that $\mu(1) = 1-\rho(1) = \gamma$ and $\mu(0) = \mu_0$, as $\phi(0) = \rho_0$ and $\phi(1) = 1$ from part (iii). Differentiating (122) yields

$$\mu'(\zeta) = -\frac{1}{N-1}\cdot\left(\frac{\rho_0}{1-\rho_0}\right)^{\frac{1}{N-1}}\cdot\left(-\frac{\phi'(\zeta)}{\phi(\zeta)^2}\right)\cdot\phi(\zeta)\left(\frac{1}{\phi(\zeta)}-1\right)^{\frac{1}{N-1}-1} - \rho(\zeta) - \zeta\rho'(\zeta)$$

$$= \left\{\left(\frac{\rho_0}{1-\rho_0}\right)^{\frac{1}{N-1}-1}(1-\zeta)^{N-2}\left(\frac{(1-\rho_0)(1-\zeta)^{N-1}}{\rho_0}\right)^{\frac{1}{N-1}-1}\right\}\cdot\phi(\zeta) \tag{136}$$

$$-\rho(\zeta)-\rho(\zeta)^2\cdot\left[\frac{1}{\rho(\zeta)}-\frac{1}{\phi(\zeta)}\right]$$

$$= \frac{(\phi(\zeta)-\rho(\zeta))^2}{\phi(\zeta)} \tag{137}$$

where the second line follows from inserting ρ' from (134) and ϕ' from (135), and the third line from the fact that the term in curly brackets in the second line is equal to one. Now, the term in (137) is positive for all $\zeta \in (0,1]$ by part (iii), which implies that $\mu(\cdot)$ is strictly increasing. Next, $\mu_0 > 0$ follows from

$$\mu_0 = 1 - \frac{1}{N-1}\cdot\left(\frac{\rho_0}{1-\rho_0}\right)^{\frac{1}{N-1}}\int_1^{\frac{1}{\rho_0}}\frac{(v-1)^{-\frac{N-2}{N-1}}}{v}\,dv \tag{138}$$

$$> 1 - \frac{1}{N-1}\cdot\left(\frac{\rho_0}{1-\rho_0}\right)^{\frac{1}{N-1}}\int_1^{\frac{1}{\rho_0}}\frac{(v-1)^{-\frac{N-2}{N-1}}}{1}\,dv = 0, \tag{139}$$

where the second line follows from $1/v < 1$ for $v \in [1, \frac{1}{\rho_0}]$, and the final equality from a straightforward calculation. Finally, $1-\rho(\zeta)-\mu(\zeta) > 0$ for all $\zeta \in (0,1)$ follows from the fact that $\rho(\cdot)$ and $\mu(\cdot)$ are strictly increasing and $1-\rho(1)-\mu(1) = 0$, since $\rho(1) = 1-\gamma$ and $\mu(1) = \gamma$.

(v) To show $\lim_{\gamma\to 0}\rho(\zeta) = 1$, from definition (118), it follows that $\lim_{\gamma\to 0}\rho_0 = 1$ which entails that for all $\zeta \in [0,1]$, we have $\lim_{\gamma\to 0}\rho(\zeta) = 1$, because $\rho(\cdot)$ is strictly increasing in ζ with $\rho(0) = \rho_0$ and $\rho(1) \le 1$ from part (ii).

Analogously, to show $\lim_{\gamma\to 1}\mu(\zeta) = 1$, it suffices to show $\lim_{\gamma\to 1}\mu_0 = 1$. To show this, as $\mu(0) = \mu_0$ by part (iv), inserting $\underline{u} = 0$ and $\zeta = 0$ in (45) yields

$$\int u \, d\kappa_{0,0}(u) = \mu_0 \cdot \omega \tag{140}$$

where $\kappa_{\zeta,\underline{u}}$ denotes, as before, function κ_ζ defined in (14) with the interpretation that ζ and $\underline{u}$ are independent variables. From (13), (14) and $\rho(0) = \rho_0$ by part (ii), we have

$$\kappa_{0,0}(u) = \left[\frac{1-\gamma}{N\gamma}\cdot\frac{u}{\omega-u}\right]^{\frac{1}{N-1}} \tag{141}$$

with support $[0,(1-\rho_0)\omega]$. Observe that from (141), as $\gamma \to 1$, $\kappa_{0,0}$ converges weakly to δ_ω.[45] From this, it follows that as $\gamma \to 1$, the LHS of (140) converges to ω, and thus, $\lim_{\gamma\to 1}\mu_0 = 1$.

[45] Consistent with this, as $\gamma \to 1$, we have $\bar{u}_{0,0} \to \omega$, because $\lim_{\gamma\to 1} 1-\rho_0 = 1$, as $\lim_{\gamma\to 1}\rho_0 = 0$ as argued before.

Finally, for all $\gamma \in (0,1)$, we have $\rho(\zeta) > 0$, because $\rho(\cdot)$ is increasing with $\rho(0) = \rho_0$ by part (ii) and $\rho_0 > 0$ by definition (118). Analogously, $\mu(\zeta) > 0$ for all $\gamma \in (0,1)$ follows from part (iv). Therefore, as $1-\rho(\zeta)-\mu(\zeta) \geq 0$ for all $\zeta \in [0,1]$ by part (iv), $\lim_{\gamma\to 0}\mu(\zeta) = 0$ follows from $\lim_{\gamma\to 0}\rho(\zeta) = 1$ and $\lim_{\gamma\to 1}\rho(\zeta) = 0$ follows from $\lim_{\gamma\to 1}\mu(\zeta) = 1$. ■

Numerical results for $\rho_0 < \mu_0$

By definition (118) and (119), μ_0 and ρ_0 depend only on the parameters N and γ. Our numerical results show that the difference $\mu_0 - \rho_0$ increases in γ. Hence, for each N, there exists $\gamma_{crit}(N)$ such that $\rho_0 < \mu_0$ holds if and only if γ exceeds $\gamma_{crit}(N)$. The following plot shows $\gamma_{crit}(N)$.

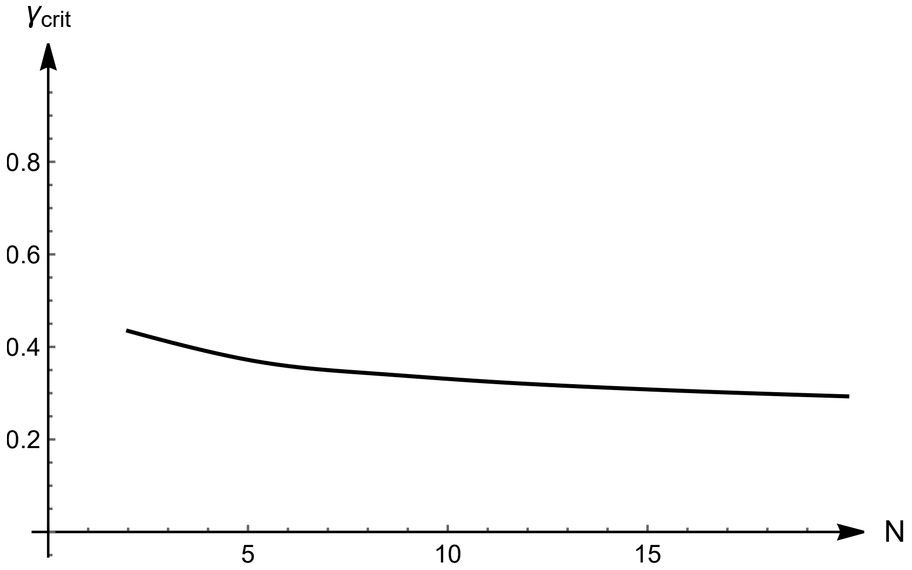

Figure 8: $\gamma_{crit}(N)$. *The condition $\rho_0 < \mu_0$ holds if and only if γ exceeds $\gamma_{crit}(N)$.*

References

Anderson, S. P. and R. Renault (1999). Pricing, Product Diversity, and Search Costs: a Bertrand-Chamberlin-Diamond Model. *The RAND Journal of Economics 30*(4), 719–735.

Antler, Y. and B. Bachi (2021). Searching Forever After. *American Economic Journal: Microeconomics*, (forthcoming).

Armstrong, M. and Y. Chen (2009). Inattentive Consumers and Product Quality. *Journal of the European Economic Association 7*(2-3), 411–422.

Armstrong, M., J. Vickers, and J. Zhou (2009). Prominence and Consumer Search. *Rand Journal of Economics 40*(2), 209–233.

Benabou, R. and R. Gertner (1993). Search with learning from prices does increased inflationary uncertainty lead to higher markups. *Review of Economic Studies 60*(1), 69–94.

Burdett, K. and K. L. Judd (1983). Equilibrium Price Dispersion. *Econometrica 51*(4), 955–969.

Dana, J. D. (1994). Learning in an Equilibrium Search Model. *International Economic Review 35*(3), 745–771.

Diamond, P. A. (1971). A model of price adjustment. *Journal of Economic Theory 3*(2), 156–168.

Enke, B. and F. Zimmermann (2019). Correlation Neglect in Belief Formation. *The Review of Economic Studies 86*, 313–332.

Esponda, I. and D. Pouzo (2016). Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models. *Econometrica 84*(3), 1093–1130.

Eyster, E. and M. Rabin (2005). Cursed equilibrium. *Econometrica 73*(5), 1623–1672.

Gamp, T. and D. Krähmer (2021). Competition in Search Markets with Naive Consumers. *RAND Journal of Economics*, (forthcoming).

Garcia, D. and S. Shelegia (2018). Consumer search with observational learning. *RAND Journal of Economics 49*(1), 224–253.

Grubb, M. D. (2015). Failing to Choose the Best Price : Theory, Evidence, and Policy. *Review of Industrial Organization 47*(3), 303–340.

Heidhues, P., J. Johnen, and B. Köszegi (2021). Browsing versus Studying Offers. *Review of Economic Studies 88*(2), 708–729.

Heidhues, P., B. Koszegi, and P. Strack (2021). Convergence in models of misspecified learning. *Theoretical Economics 16*(1), 73–99.

Heidhues, P. and B. Kőszegi (2018). Behavioral Industrial Organization. *Handbook of Behavioral Economics 1*, 517–612.

Heller, Y. and E. Winter (2020). Biased-Belief Equilibrium. *American Economic Journal: Microeconomics 12*(2), 1–40.

Ispano, A. and P. Schwardmann (2020). Competition over cursed consumers. *Working Paper*.

Janssen, M. and S. Shelegia (2020). Beliefs and Consumer Search in a Vertical Industry. *Journal of the European Economic Association 18*(5), 2359–2393.

Janssen, M. C., A. Parakhonyak, and A. Parakhonyak (2017). Non-reservation price equilibria and consumer search. *Journal of Economic Theory 172*, 120–162.

Jehiel, P. (2005). Analogy-based expectation equilibrium. *Journal of Economic Theory 123*(2), 81–104.

Karle, H., H. Schumacher, and R. Volund (2020). Consumer Search and the Uncertainty Effect. *Working Paper*.

Kohn, M. G. and S. Shavell (1974). The Theory of Search. *Journal of Economic Theory 123*, 93–123.

Lauermann, S., W. Merzyn, and G. Virág (2018). Learning and price discovery in a search market. *Review of Economic Studies 85*(2), 1159–1192.

Matsumoto, B. and F. Spence (2016). Price beliefs and experience: Do consumers' beliefs converge to empirical distributions with repeated purchases? *Journal of Economic Behavior and Organization 126*, 243–254.

Mauring, E. (2020). Partially Directed Search for Prices. *Working Paper*.

Parakhonyak, A. and A. Sobolev (2015). Non-Reservation Price Equilibrium and Search without Priors. *Economic Journal 125*(584), 887–909.

Price, C. W. and M. Zhu (2016). Empirical evidence of consumer response in regulated markets. *Journal of Competition Law and Economics 12*(1), 113–149.

Salop, S. and J. E. Stiglitz (1977). Bargains and Ripoffs: A Model of Competitive Monopolistically Price Dispersion. *The Review of Economic Studies 44*(3), 493–510.

Stahl, D. O. I. (1989). Oligopolistic Pricing with Sequential Consumer Search. *The American Economic Review 79*(4), 700–712.

Stiglitz, J. E. (1979). Equilibrium in Product Markets with Imperfect Information. *The American Economic Review 69*(2).

Varian, H. R. (1980). A Model of Sales. *The American Economic Review 70*(4), 651–659.

Wolinsky, A. (1986). True Monopolistic Competition as a Result of Imperfect Information. *The Quarterly Journal of Economics 101*(3), 493–512.