DISCUSSION PAPER SERIES

DP13842

NONLINEAR PRICING WITH AVERAGE-PRICE BIAS

David Martimort and Lars Stole

INDUSTRIAL ORGANIZATION

ISSN 0265-8003

NONLINEAR PRICING WITH AVERAGE-PRICE BIAS

David Martimort and Lars Stole

Discussion Paper DP13842
Published 06 July 2019
Submitted 17 June 2019

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

NONLINEAR PRICING WITH AVERAGE-PRICE BIAS

Abstract

Empirical evidence suggests that consumers facing complex nonlinear pricing often make choices based on average (not marginal) prices. Given such behavior, we characterize a monopolist's optimal nonlinear price schedule. In contrast to the textbook setting, nonlinear prices designed for ``average-price bias'' distort consumption downward for consumers at the top, may produce efficient consumption for consumers at the bottom, and typically feature quantity premia rather than quantity discounts. These properties arise because the bias replaces consumer information rents with curvature rents. Whether or not a monopolist prefers consumers with average-price bias depends upon underlying preferences and costs.

JEL Classification: D82

Keywords: Nonlinear Pricing, average-price bias, curvature rents, price discrimination

David Martimort - david.martimort@psemail.eu
Paris School of Economics and CEPR

Lars Stole - lars.stole@chicagobooth.edu
Chicago Booth School of Business

NONLINEAR PRICING WITH AVERAGE-PRICE BIAS*

DAVID MARTIMORT† LARS STOLE‡

June 15, 2019

Abstract

Empirical evidence suggests that consumers facing complex nonlinear pricing often make choices based on average (not marginal) prices. Given such behavior, we characterize a monopolist's optimal nonlinear price schedule. In contrast to the textbook setting, nonlinear prices designed for "average-price bias" distort consumption downward for consumers at the top, may produce efficient consumption for consumers at the bottom, and typically feature quantity premia rather than quantity discounts. These properties arise because the bias replaces consumer *information rents* with *curvature rents*. Whether or not a monopolist prefers consumers with average-price bias depends upon underlying preferences and costs.

KEYWORDS. Nonlinear pricing, average-price bias, curvature rents, price discrimination.

JEL CODES. D82.

*We thank Austan Goolsbee, Bruno Jullien, Emir Kamenica, Amin Sayedi and seminar participants at University of Chicago, Toulouse School of Economics and UC Berkeley for comments and useful suggestions on earlier drafts. The usual disclaimer applies.

†Paris School of Economics-EHESS. E-mail: martimort.david@gmail.com.

‡Chicago Booth School of Business. E-mail: lars.stole@chicagobooth.edu.

1 Introduction

The theory of optimal screening with privately-informed agents has been fruitfully applied to many settings, including optimal taxation (e.g., Mirrlees (1971)), second-degree price discrimination (e.g., Mussa and Rosen (1978), Maskin and Riley (1984)), and regulation (e.g., Baron and Myerson (1982)). In each case, a principal (e.g., tax authority, monopolist, regulator) designs an optimal menu of choices (e.g., a nonlinear schedule or tariff) from which the privately-informed agent selects. Importantly, it is assumed that the agent rationally equates marginal benefits with marginal costs.

Suppose instead the agent has difficulty discerning the relevant margin from the offered nonlinear schedule and wrongly focuses on average price. What is the optimal contract in light of this bias? We specialize this question to the setting of monopoly nonlinear pricing, supposing that buyers wrongly consume where marginal benefit equals average price – an error we call *average-price bias*. This paper's purpose is to understand how a strategic monopolist would exploit such a bias, and to explore the welfare implications of such behavior.

Several papers document the presence of average-price bias in decision making. For example, in Koichiro Ito (2014), spatial discontinuities in electricity service areas in California from 1999-2007 are used to identify the perceived consumer price that determines demand; the paper finds that average price explains short-run demand variations far better than the marginal price predicted by neoclassical theory. Blake Shaffer (2018) finds additional supporting evidence in the British Columbia electricity market.[1] Laboratory evidence also documents a bias toward average prices. Using a sample of MBA students instructed to choose between two investment decisions (one tax free and one taxed), de Bartolome (1995) finds that most students fail to compute the correct marginal tax rate from the provided tables, relying instead on the average rate.

The evidence to date of average-price bias is provocative, if not persuasive, but we

[1] Shaffer (2018) includes a richer model of consumer heterogeneity, allowing for multiple behavioral types, and estimates 85% of behavioral types exhibit average-price bias.

do not seek to rationalize average-price bias in this paper. Our purpose is simply to take average-price bias as an exogenous feature of consumer decision making and draw out the theoretical implications.[2] To this end, we combine the textbook model of monopoly nonlinear pricing with the assumption of average-price bias. We find the mechanics of optimal nonlinear pricing with average-price biased consumers is remarkably simple and an interesting contrast to the standard model. First, in a world with biased consumers, we show that the monopolist no longer faces a tradeoff between information-rent extraction and efficient consumption. Instead, because consumers effectively perceive the nonlinear price offering of the monopolist as a linear function, private information is freely (and indirectly) revealed. That said, average-price biased consumers still earn consumer surplus through what we call *curvature* rents. We document how the monopolist's desire to extract curvature rents rather than information rents leads to significant differences in the shape of optimal nonlinear prices. In the simplest quadratic preference model, the quantity discounts arising in Maskin and Riley (1984) are replaced with quantity premia. Second, because the average-price bias represents both a cost and a benefit to the monopolist (i.e., curvature rents have replaced information rents), it is unclear that a monopolist will prefer to face rational consumers or those with bias. Indeed, with constant unit costs, quadratic utility, and uniformly-distributed preferences, we show that expected monopoly profits and consumer surplus are equal across the regimes, although the quantity allocations are distinct. We then introduce cost nonlinearities to identify sources of variation in profit and surplus across the two settings, illuminating the consequences of average-price bias on cost pass through.

At least a few theoretical contributions are related to this paper. First, many papers have incorporated models of decision-making from behavioral economics and psychology into the standard contracting paradigms.[3] These are directly related to

[2]The literature on inattention and salience provides evidence that consumers are inattentive to the correct prices: e.g., ignoring taxes that are not salient (Chetty, et al. (2009)), treating shipping costs on eBay differently from product prices (Hossain and Morgan (2006)), and ignoring fixed-price offers that are dominated by eBay auctions (Malmendier and Lee (2011)). Della Vigna (2009) surveys the evidence of limited attention, salience and suboptimal heuristics from field experiments.

[3]Köszegi (2014) provides a survey.

the nonlinear pricing framework of the present paper. Carbajal and Ely (2016) and Hahn, et al (2018), for example, introduce loss aversion and reference-dependent preferences (as in Köszegi and Rabin (2006)) into a nonlinear pricing model. Our paper is in the same spirit but focuses on average-price bias. A separate literature considers "unawareness" in contracts where the agent may be unaware of contingencies and the principal may leave contracts incomplete, lest attention be drawn to the unconsidered dimensions.[4] In our setting, the consumer is unaware that the price schedule is nonlinear, but we do not consider the possibility that changes in the principal's contract could alter this belief.

The working paper of Liebman and Zeckhauser (2004) is perhaps closest to ours.[5] Liebman and Zeckhauser's focus on two forms of agent bias, one of which is average-price bias (what LZ refer to as "ironing"), which they consider in a model of monopoly pricing with two consumer types. In their nonlinear pricing application, the monopolist prefers biased consumers to rational consumers – in contrast to our results – because they restrict attention to price schedules that are piecewise linear, convex, and *originate from the origin*. Such price schedules leave positive rents to low-demand consumers which is suboptimal with unbiased consumers. LZ also do not identify the tradeoff between information rents and *curvature rents*. Illuminating this tradeoff is a primary contribution of the present paper.

[4] See, for example, Filiz-Obay (2012), in the context of insurance.

[5] Farhi and Gabaix (2018) develop a general framework of optimal taxation with behavioral agents which is sufficiently general that it can embed the biases in Liebman and Zeckhauser (2004), as well as a long list of other well-studied heuristics and biases. We emphaszie the specific implications of average-price bias, for which the economics analysis in LZ is more directly related.

2 Nonlinear pricing without bias

2.1 Preferences

We begin with the textbook[6] model of nonlinear pricing by a monopolist, and suppose that a consumer's payoff is

$$u(q, \theta) - P(q),$$

where $q \in \mathcal{Q} \equiv [0, \overline{q}]$ is the quantity (or quality) consumed, $\theta \in \Theta \equiv [\theta_0, \theta_1]$ is the consumer's type which is distributed according to a continuous, log-concave distribution, F (density, f), and $P(q)$ is the total payment made to the firm for q units. We assume that u is twice continuously differentiable, strictly concave in q, $u(0, \theta) = 0$, $u_\theta(q, \theta)$ is nonnegative and bounded, and $u_{q\theta}(q, \theta) > 0$. The familiar single-crossing property implies higher type consumers have a higher marginal value of consumption. The consumer's outside option is normalized to zero.

The firm's per-consumer cost is $C(q)$ – convex and continuously differentiable. The firm offers the consumer a nonlinear price schedule, $P : \mathcal{Q} \to \Re_+$, in order to maximize expected profits,

$$\mathrm{E}_\theta \left[P(q(\theta)) - C(q(\theta)) \right],$$

where $q(\theta)$ denotes the incentive-compatible consumption of the type-θ consumer when offered $P(\cdot)$.

2.2 Optimal nonlinear pricing without bias

We initially assume that consumers are unbiased in order to reproduce the standard nonlinear pricing results as a benchmark. Denote the optimal price schedule as $P^n(\cdot)$ and the corresponding allocation for type θ as $q^n(\cdot)$. Throughout we use the superscipt "n" to indicate the unbiased, neoclassical consumers from the textbook model. In contrast, for the average-price-bias setting, we will denote the optimal price schedule and allocation as $\overline{P}(\cdot)$ and $\overline{q}(\theta)$, respectively.

[6] See, for example, Tirole (1988), chapter 3, and Maskin and Riley (1984).

Given a price schedule, define the consumer surplus for the unbiased setting as

$$U^n(\theta) \equiv \max_{q \in \mathcal{Q}} \; u(q,\theta) - P(q).$$

As is well known, an allocation $q(\theta)$ is implementable by some nonlinear price schedule, P, if and only if q is nondecreasing and the derivative of U^n is equal to $u_\theta(q(\theta),\theta)$ almost everywhere. The latter envelope condition implies that the consumer's expected surplus in any incentive-compatible mechanism is

$$\mathrm{E}_\theta\left[U^n(\theta)\right] = \mathrm{E}_\theta\left[U^n(\theta_0) + \int_{\theta_0}^{\theta} u_\theta(q(s),s))ds\right] = U^n(\theta_0) + \mathrm{E}_\theta\left[u_\theta(q(\theta),\theta)\frac{1-F(\theta)}{f(\theta)}\right].$$

Because profits equal the difference between the social and consumer surpluses,

$$\begin{aligned}
\mathrm{E}_\theta\left[\Pi^m(\theta)\right] &= \mathrm{E}_\theta\left[u(q(\theta),\theta) - C(q(\theta)) - U^n(\theta)\right] \\
&= \mathrm{E}_\theta\left[u(q(\theta),\theta) - C(q(\theta)) - \frac{1-F(\theta)}{f(\theta)}u_\theta(q(\theta),\theta) - U^n(\theta_0)\right].
\end{aligned}$$

We make the standard regularity assumption that this integrand exhibits increasing differences in (q,θ), ensuring that the pointwise maximum of q over θ is nondecreasing.[7] Thus, the optimal allocation for the monopolist is

$$q^n(\theta) = \arg\max_{q \in \mathcal{Q}} \; u(q,\theta) - C(q) - \frac{1-F(\theta)}{f(\theta)}u_\theta(q,\theta), \tag{1}$$

and the corresponding price schedule satisfies

$$dP^n(q)/dq = u_q(q, \vartheta^n(q)),$$

where $\vartheta^n(q)$ is the inverse of $q^n(\theta)$, and

$$P^n(q^n(\theta_0)) = u(q^n(\theta_0), \theta_0).$$

[7] Given that F is log concave, increasing differences arise if $u(q,\theta)$ and $C(q)$ are quadratic functions.

The optimal nonlinear pricing solution exhibits three properties.

1. *No surplus at the bottom.* The lowest-type consumer gets zero surplus, regardless of whether $q^n(\theta_0)$ is positive.

2. *No distortion at the top.* The highest-type consumes efficiently, $u_q(q^n(\theta_1), \theta_1) = C'(q^n(\theta_1))$. The monopolist distorts consumption only to extract information rents, but the latter (captured by $u_\theta(q,\theta)(1 - F(\theta))/f(\theta)$) is zero for the highest type.

3. *Quantity discounts are optimal if $C(q) = cq$ and u is quadratic.* As noted by Maskin and Riley (1984), under mild assumptions the price schedule exhibits quantity discounts.[8]

These three properties are absent when consumers exhibit average-price bias.

3 Average-price bias

3.1 Behavioral Assumption of Average-price Bias

The contribution of this paper is to consider the possibility that consumers incorrectly use *average price* when choosing consumption, q. A simple behavioral rule which encapsulates this bias requires that for a given price schedule, $P(q)$, a consumer chooses output to satisfy

$$q \in \arg\max_{\tilde{q} \in Q} \; u(\tilde{q}, \theta) - \frac{P(q)}{q}\tilde{q}. \tag{2}$$

Notice that this optimization problem is self-referential: we require that the q that is used to determine the average unit price is also the solution to the consumer's choice problem. For simplicity, we assume that if there are multiple solutions to (2), then the monopolist's preferred solution is selected; if no solution exists then the consumer

[8]More generally, quantity discounts arise if $u_{q\theta\theta}(q,\theta) \leq 0$ and $u_{qq\theta}(q,\theta) \leq 0$.

chooses $q = 0$.[9] To this end, define the correspondence

$$\overline{Q}(\theta, P) \equiv \{q \mid u_q(q, \theta)q = P(q)\} \cup \{0\}.$$

The monopolist chooses the schedule $P(\cdot)$ and a selection, $q(\theta) \in \overline{Q}(\theta, P)$ to maximize expected profits.

Although our primary goal is to explore the consequences of (2) for optimal nonlinear pricing, we briefly posit two foundations for our behavioral rule. Both arise from a misspecification error by an otherwise rational consumer – the consumer incorrectly believes the firm's price schedule is linear.

Model 1: Learning with misspecification bias

The behavior in (2) can be rationalized with a dynamic-learning model similar to that in Bengt Holmström (1999). Assume the consumer incorrectly believes that the data generating process for prices is governed by $P_t = (p_t + \varepsilon_t)q_t$, where p_t is the perceived constant linear price of the monopolist in period t and ε_t is i.i.d. normal measurement error. The consumer also believes the monopolist changes prices from period to period according to $p_t = p_{t-1} + \eta_t$, where η_t is another i.i.d. normal process with zero mean; this assumption conveniently ensures the consumer never stops learning. Finally, assume an initial, normally-distributed prior estimate of the monopolist's price at the start of the game, p_1. For this normal-learning model with misspecification bias, an otherwise statistically-capable consumer would estimate the unit price of the monopolist in period t $(t \geq 2)$ to be

$$\overline{p}_t = \lambda_t \overline{p}_{t-1} + (1 - \lambda_t) \frac{P(q_{t-1})}{q_{t-1}},$$

where λ_t gives the variance-minimizing weights on the past-period's belief and recent price information. λ_t is a deterministic function of the variances of ε, η and the prior

[9]This assumption is inconsequential, as we will see the optimal price schedule will never exhibit multiple, positive solutions.

belief, p_1.[10] The consumer will choose q_t each period to equate marginal utility to the expected (perceived) marginal price, yielding a first-order dynamic relationship

$$u_q(q_t, \theta) = \lambda_t u_q(q_{t-1}, \theta) + (1 - \lambda_t)\frac{P(q_{t-1})}{q_{t-1}},$$

where λ_t converges to $\overline{\lambda} \in (0, 1)$ and the output satisfying (2) is a steady-state solution.[11]

Model 2: A communication mechanism leveraging misspecification bias

Our second foundation for (2) is also a heuristic to illustrate that the firm is able to learn a biased consumer's information for free.

Consider the following direct-revelation recommendation mechanism:

- Stage 1: the firm commits to a menu of prices and quantities, $P(\cdot)$:
- Stage 2: the consumer reports type, $\hat{\theta}$;
- Stage 3: the firm recommends $q = r(\hat{\theta})$ and announces the associated "unit" price $p(\hat{\theta}) = P(r(\hat{\theta}))/r(\hat{\theta})$ from the set price schedule;
- Stage 4: the consumer either follows the recommendation, choosing $q = r(\hat{\theta})$ and paying $p(\hat{\theta})r(\hat{\theta})$, or chooses an outside option.

The consumer values this recommendation mechanism because it reduces computational costs, though it is unrealistic to imagine the consumer can

[10] If $\varepsilon \sim \mathcal{N}(0, h_\varepsilon)$, $\eta \sim \mathcal{N}(0, h_\eta)$ and $p_1 \sim \mathcal{N}(0, h_1)$, where the h_i's are precisions (inverse variances) of the random variables, then following Holmström (1999),

$$\lambda_t = \left(2 + \frac{h_\varepsilon}{h_\eta} - \lambda_{t-1}\right)^{-1}.$$

[11] For λ_t sufficiently close to 1 (which arises if η has sufficiently small variance and sufficient time has passed), consumption around $\overline{q}$ is locally stable if $u_q(q, \theta)q - P(q)$ is locally concave. Furthermore, if utility is quadratic and unit costs are constant, for the static-optimal price schedule P (derived in Proposition 1) , the difference equation reduces to

$$(q_t - \overline{q}) = (2\lambda_{t-1} - 1)(q_{t-1} - \overline{q}),$$

and consumption converges globally, $q_t \to \overline{q} = \frac{\theta - c}{2\gamma}$. Formally, we have derived (2) as a "restricted perceptions" learning equilibrium, in the words of Evans and Honkapohja (2001).

communicate a complicated object such as θ without noticing that $P(q)$ is nonlinear. Nonetheless, this mechanism illustrates why the consumer will reveal θ to the firm without any information rent. The consumer believes $P(q)$ is linear and, therefore, the announced unit price $p(\hat{\theta})$ is a sufficient statistic for the entire schedule. Under this misperception, at stage 4, the consumer will follow the firm's recommendation if either $r(\hat{\theta}) = 0$ or

$$u_q(r(\hat{\theta}), \theta) = \frac{P(r(\hat{\theta}))}{r(\hat{\theta})}.$$

At stage 3, the firm will make recommendations to satisfy these conditions (doing otherwise would lead the consumer to exit the market). At stage 2, because the firm has committed to $P(\cdot)$ and cannot introduce off-menu recommendations, the consumer, believing the menu is linear, believes the firm has no incentive to give an improper recommendation. At the reporting stage, the consumer misperceives that the firm and consumer are playing a common-interest game. As such, the consumer finds it optimal to report truthfully, $\hat{\theta} = \theta$. In stage 1, the firm generally does not find it optimal to offer a linear price schedule, so the stage 3 recommendation serves only to direct to consumer to a choice $q \in \overline{Q}(\theta, P)$, without allowing the consumer to learn of the misspecification bias. The result of consumer's misspecification is that a price schedule with a recommendation mechanism can elicit the consumer's type information for free in stage 2, though the firm is compelled to direct the consumer to a point satisfying (2) at stage 3. Consistent with this intuition, we will find that nonlinear pricing does not generate information rents to the biased consumer satisfying (2), though consumer surplus will be positive and driven by the curvature of the consumer's utility function.

As a specific example of a recommendation mechanism, suppose that $u(q, \theta) = \theta q - \frac{1}{2}q^2$, $C(q) = cq$ and the firm offers a Stage 1 price schedule of $P(q) = qc + q^2$. (We will show below that this price schedule is profit maximizing.) As illustrated in Figure 1, an agent who reports type θ can be induced to choose $\overline{q}(\theta) = (\theta - c)/2$ at total price of $P(q(\theta)) = (\theta + c)(\theta - c)/4$ – a point where the convex price schedule *intersects* his indifference curve – because he incorrectly perceives a linear tariff through that point (dotted line) which is tangent to his indifference curve with slope $(\theta + c)/2$.

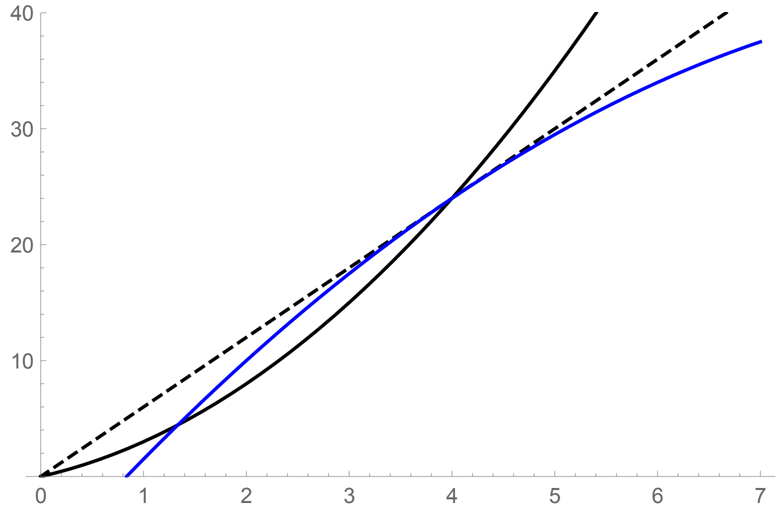

Figure 1: Here, $c = 2$ and $\theta = 10$. The principal offers the convex menu $P(q) = 2q + q^2$ (solid curve), recommends the agent to choose the menu item $\{\overline{p}(\theta) = 24, \overline{q}(\theta) = 4\}$, and induces this by letting the consumer believe he is facing the tariff $\tilde{P}(q) = 6q$ (dashed line). The concave indifference curve of the agent is tangent to $\tilde{P}$ at $q = 4$, but of course the agent could do better by selecting $q = 8/3$ if they were aware of their mis specification.

3.2 Implementability and curvature rents

Recall that the monopolist chooses a price schedule, $P(\cdot)$, and a selection from the bias-choice correspondence, $\overline{Q}(\theta, P)$. We now determine which quantity allocations, $q : \Theta \to \mathcal{Q}$, are implementable and the price schedules that implement them. Compared to the setting with unbiased consumers, the answer is relatively straightforward. The key is to note that for any type θ, $q(\theta) > 0$ can be implemented by offering a price $P(q(\theta))$ such that

$$P(q(\theta)) = u_q(q(\theta), \theta) q(\theta).$$

Any type θ can be induced to choose any $q > 0$, in isolation, by choosing a price schedule whose average price $P(q)/q$ is equal to the marginal utility of the type-θ consumer, $u_q(q, \theta)$ at q. Similarly, the monopolist can implement $q(\theta) = 0$ for any type θ by fiat. Not every allocation mapping, $q(\cdot) : \Theta \to \mathcal{Q}$ is implementable, however, because $P(q)$ must be a function giving a single price for each q. Consequently, it is impossible to implement the same $q > 0$ for two different types of consumers: e.g., $P(q)/q = u_q(q,\theta)$ implies $P(q)/q \neq u_q(q,\theta')$ for $\theta \neq \theta'$. Except for the infeasibility of implementing pooling allocations for some $q > 0$, there are no other restrictions. We conclude that an allocation $\overline{q}(\cdot)$ is implementable under average-price bias iff it is fully separating over positive quantities. In such a case, the implementing price schedule is

$$\overline{P}(q) = \begin{cases} u_q(q,\overline{\vartheta}(q))q & \text{for all } q \in \overline{q}([\theta_0,\theta_1]) \text{ and } q > 0 \\ \infty & \text{otherwise} \end{cases} \tag{3}$$

where $\overline{\vartheta}(q) = \overline{q}^{-1}(q)$ for $q > 0$.

Because $u(q,\theta)$ is strictly concave, a consumer equating marginal utility equal to average price necessarily obtains surplus. To be clear, define the difference between a consumer's average utility and marginal utility by

$$\Delta(q,\theta) \equiv \frac{u(q,\theta)}{q} - u_q(q,\theta) \geq 0.$$

Consumer surplus, given (3), is therefore

$$CS(q,\theta) = u(q,\theta) - P(q) = u(q,\theta) - u_q(q,\theta)q = \Delta(q,\theta)q.$$

A consumer who chooses positive consumption will obtain surplus that is proportional to the curvature of the utility function. Consequently, we call $\Delta(q,\theta)$ the marginal *curvature rent* of type θ from consuming q.[12] Because of this curvature rent, a consumer's participation constraint is trivially satisfied. Indeed, if $\overline{q}(\theta_0) > 0$, then the lowest-type agent will earn a positive rent, unlike the standard nonlinear pricing model.

[12] Δ may be independent of type (e.g., if u is linear in θ).

3.3 Optimal nonlinear pricing with average-price bias

Given (3), finding the optimal nonlinear price under average-price bias is straightforward. The firm maximizes expected profit, subject only to (3) and a separation requirement. We proceed by considering the relaxed program (ignoring separation) and verify that the solution does not induce pooling on positive outputs.

In the relaxed program, the firm chooses the allocation q pointwise in θ to maximize

$$\mathrm{E}_\theta\left[P(q(\theta)) - C(q(\theta))\right] = \mathrm{E}_\theta\left[u_q(q(\theta),\theta)q(\theta) - C(q(\theta))\right].$$

For simplicity, we maintain the following assumption:

Assumption 1. *For all θ, $u_q(q,\theta)q - C(q)$ is concave in q.*

Given A.1, the solution to the relaxed program satisfies

$$u_q(\overline{q}(\theta),\theta) - C'(\overline{q}(\theta)) = -u_{qq}(\overline{q}(\theta),\theta)\overline{q}(\theta)$$

for $\overline{q}(\theta) > 0$ and $\overline{q}(\theta) = 0$ otherwise. Because $u_{q\theta} > 0$, the relaxed solution is strictly increasing in θ for for positive consumption and, hence, is a solution to the unrelaxed program.

Proposition 1. *The optimal allocation under average-price bias satisfies*

$$u_q(\overline{q}(\theta),\theta) - C'(\overline{q}(\theta)) = -u_{qq}(\overline{q}(\theta),\theta)\overline{q}(\theta); \tag{4}$$

the optimal price schedule is

$$\overline{P}(q) = u_q(q,\overline{\vartheta}(q))q, \tag{5}$$

where $\overline{\vartheta}(q)$ is the inverse of $\overline{\theta}(q)$.

Several comparisons with the standard model are now available. First, the optimality allocation is independent of F. As suggested above, average-price bias eliminates the costs of incentive compatibility; information rents (captured by the inverse hazard rate of the type distribution) are no longer earned by consumers.

Second, although there are no information rents, there are curvature rents. Indeed, one can rewrite the firm's expected profits as social surplus minus curvature rents,

$$\mathrm{E}_{\theta}\left[u(q(\theta),\theta) - C(q(\theta)) - \Delta(q(\theta),\theta)q(\theta)\right],$$

and thus interpret the first-order condition above as a marginal tradeoff between efficient consumption and reducing curvature rents.

Third, for all types that consume, consumption is inefficient – this is true even for the highest type. This everywhere-downward distortion reflects the tradeoff between efficiency and curvature rents. This sharply contrasts with the standard nonlinear pricing model, as in Maskin and Riley (1984).

Fourth, the lowest consumer type with average-price bias has (weakly) higher surplus than the same type in the unbiased setting. In the average-price bias setting, the lowest-type consumer earns $\overline{U}(\theta_0) = \Delta(\overline{q}(\theta_0),\theta_0)\overline{q}(\theta_0)$, which is positive if $\overline{q}(\theta_0) > 0$. The monopolist leaves rents to a type-θ_0 consumer because it can't extract the agent's surplus without raising the agent's perceived average price, leading to reductions in purchases.

Fifth, quantity premia are typically optimal. Recall that if u is quadratic and unit cost is constant, then the optimal price schedule in the unbiased model exhibits quantity discounts. In contrast, consider the optimal price schedule when consumers have average-price bias. Letting $\overline{\vartheta}(q)$ represent the inverse allocation, substituting (4) into (5) reveals

$$\overline{P}(q) = C'(q)q - u_{qq}(q,\overline{\vartheta}(q))q^2.$$

If utility is quadratic and $C'''(q) \geq 0$, then $\overline{P}''(q) > 0$ and there are quantity premia rather than discounts. Intuitively, an increase in the true marginal price raises the perceived marginal price less for high-consumption agents than low-consumption agents, and so (ignoring terms above second-order) the firm will want to exploit this misperception by increasing marginal prices with consumption. For average-price bias consumers, quantity premia rather than quantity discounts are the norm. This basic intuition is beneath the welfare results in Section 4.

We have exceptionally clear empirical predictions across the two settings. In the standard setting with rational consumers, nonlinear prices depend upon the underlying distribution of types and typically (e.g., with preferences well approximated by a quadratic function) exhibit quantity discounts with full efficiency for the highest-value consumer and no consumer surplus for the lowest types. In contrast, when consumers exhibit average-price bias, the optimal nonlinear price schedule does not depend upon the details of the type distribution, the schedule exhibits increasing marginal prices, consumption is distorted downward even for the highest-type consumer, and the lowest-type purchasing consumer earns a positive (curvature) rent.

4 Welfare comparisons

We are interested in how profits, consumer surplus and welfare differ across the standard model and the average-price bias setting. For this purpose it is useful to observe that the monopolist's program when facing biased consumers is mathematically equivalent to perfect third-degree price discrimination with unbiased consumers. To illustrate, suppose that the monopolist is facing a population of fully rational consumers, observes each consumer's type, but is constrained to offer a linear price to each type: i.e., $P(q \,|\, \theta) = p(\theta)q$. In this case, optimal prices are determined using standard monopoly pricing: For any θ, the firm faces a demand curve $p = u_q(q, \theta)$ and chooses price to maximize profit conditional on θ. Notice, however, that for any price $p(\theta)$, the revenue generated is $p(\theta)q(\theta) = u_q(q(\theta), \theta)q(\theta)$. This, of course, is the same revenue earned in the case of average-price biased consumers, (3). Hence, the allocation, profit and consumer surplus corresponding to the case of average-price bias are equivalent to those arising from perfect third-degree price discrimination. Welfare comparisons between settings of average-price bias and rational consumers are equivalent to comparisons between third-degree and second-degree price discrimination.

Returning to the task at hand, we have noted that the lowest-type consumer may earn a strictly positive rent in the average-price bias game, implying that such consumers do better compared to the standard model. This preference may not hold for all consumers, however, as the slope of rent with respect to type may be flatter in the bias setting. There is also ambiguity in terms of profits. In the bias game, the monopolist eliminates all information rents, but at the cost of leaving curvature rents. It is not *a priori* obvious which effect will dominate. In order to explore these different effects, we investigate a model with quadratic preferences and uniformly-distributed heterogeneity. We then illustrate how perturbations of this setting can impact welfare properties.

4.1 Results for Uniform-Quadratic Preferences

We begin with a stylized but common model of monopoly pricing in which (i) unit costs are constant, (ii) consumers have linear demand curves with identical slopes but heterogeneity in the intercepts, and (iii) consumer heterogeneity is distributed uniformly. Formally, we assume $C(q) = cq$,

$$u(q, \theta) = \theta q - \frac{\gamma}{2} q^2,$$

$\gamma > 0$, and θ is uniformly distributed on $[\theta_0, \theta_1]$. Consequently, the θ-type consumer has a neoclassical demand curve given by

$$p = \theta - \gamma q.$$

We also assume $\theta_1 > c > \theta_0$ so that it is inefficient to serve all types.

We first calculate the consumer surplus and profits associated with the standard model. Using (1),

$$q^n(\theta) = \arg\max_{q\geq 0} u(q,\theta) - cq - (\theta_1 - \theta)q = \frac{1}{\gamma}\max\{0, 2\theta - \theta_1 - c\},$$

and the corresponding optimal nonlinear price is

$$P^n(q) = \frac{q}{2}\left(\theta_1 + c - \frac{\gamma}{2}q\right).$$

The corresponding consumer surplus for type θ is

$$CS^n(\theta) \equiv \max_{q\geq 0} u(q,\theta) - P^n(q) = \frac{1}{4\gamma}\left(\max\{0, 2\theta - \theta_1 - c\}\right)^2,$$

and therefore ex ante surplus is

$$CS^n = E[CS^n(\theta)] = \frac{(\theta_1 - c)^3}{24\gamma(\theta_1 - \theta_0)}.$$

Similarly, we compute the firm's profit from a type-θ consumer and its ex ante expected profit:

$$\Pi^n(\theta) = P^n(q^n(\theta)) - cq^n(\theta),$$

$$\Pi^n = E[\Pi^n(\theta)] = \frac{(\theta_1 - c)^3}{12\gamma(\theta_1 - \theta_0)}.$$

Note that $\Pi^n = 2CS^n$. A similar relationship will hold in the case of average-price bias.

We next compute profits and consumer surplus with optimal pricing in the average-price bias game. This setting is mathematically equivalent to third-degree price discrimination conditional on the consumer's type, θ. The consumption allocation under average-price bias is given by (4), which specializes to

$$\bar{q}(\theta) = \arg\max_{\tilde{q}\geq 0} (\theta - \gamma\tilde{q} - c)\tilde{q} = \max\left\{0, \frac{\theta - c}{2\gamma}\right\}.$$

The corresponding consumer surpluses and profits are

$$\overline{CS}(\theta) = \frac{1}{8\gamma}\left(\max\{0, \theta - c\}\right)^2, \quad \overline{CS} = E[\overline{CS}(\theta)] = \frac{(\theta_1 - c)^3}{24\gamma(\theta_1 - \theta_0)},$$

$$\overline{\Pi}(\theta) = \frac{1}{4\gamma}(\max\{0, \theta - c\})^2, \quad \overline{\Pi} = E[\overline{\Pi}(\theta)] = \frac{(\theta_1 - c)^3}{12\gamma(\theta_1 - \theta_0)}.$$

As in the case of unbiased consumers, expected profits are twice the expected consumer surplus. More remarkably, these ex ante values are equal across the regimes. The expected benefit of eliminating information rents equals the expected cost curvature rents.

Proposition 2. *Suppose that $u(q, \theta) = \theta q - \frac{\gamma}{2}q^2$, $C(q) = cq$ and θ is uniformly distributed. Then*

$$\overline{\Pi} = \Pi^n, \quad \overline{CS} = CS^n.$$

Although profits and consumer surpluses are equal across the regimes, it is worth emphasizing that the optimal allocations in the two regimes, $\overline{q}$ and q^n, are almost everywhere different. Furthermore, the marginal non-participating consumer type under average-price bias, c, is different from the marginal consumer under the standard second-degree price discrimination setting, $\frac{1}{2}(c + \theta_1)$. Figure 2 illustrates these differences.

4.2 Welfare for nonlinear costs

The result in Proposition 2 is knife edge and relies on constant unit costs, affording an opportunity to understand how departures from constant costs will favor one regime over the other, either in terms of profit or consumer surplus. In particular, we consider simple variations with a more general cost function $C(q) = c_0 + cq + \frac{c_2}{2}q^2$ and explore the impact on relative changes to surplus and profits.

Strictly convex costs: Suppose that $c_0 = 0$ but $c_2 > 0$; i.e., the cost function is strictly convex with $C(0) = 0$. Assume also $c_2 < \gamma$, which implies that the unbiased

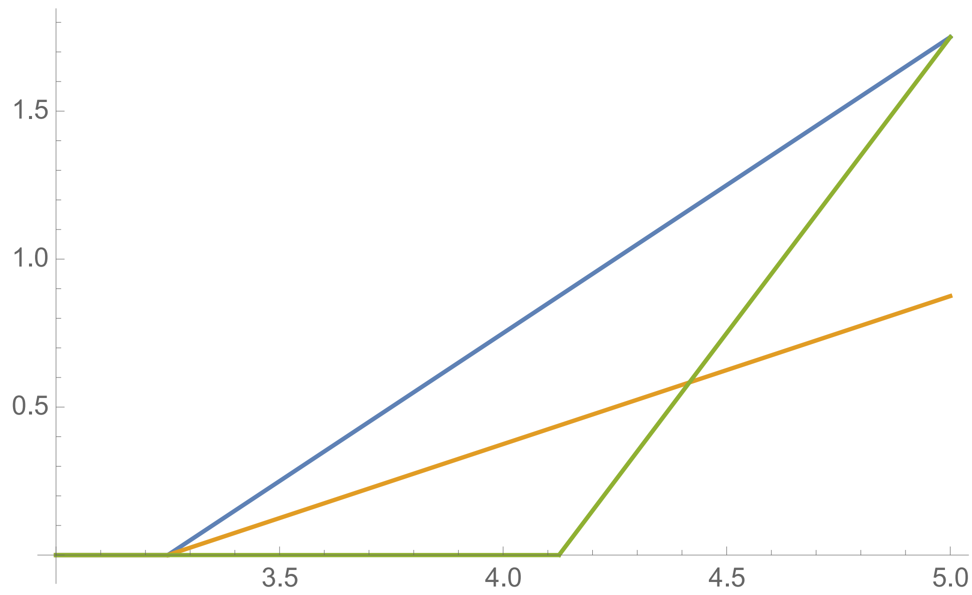

Figure 2: First-best allocation, $q_{fb}(\theta) = \max\{0, \theta - 3.25\}$; "average-price bias" allocation, $\overline{q}(\theta) = \frac{1}{2}\max\{0, \theta - 3.25\}$; second-degree price discrimination allocation, $q_{2nd}(\theta) = \max\{0, 2\theta - 8.25\}$. In all cases, $\theta_0 = 3$, $\theta_1 = 5$, $c = 3.25$, $= \gamma = 1$, $\alpha = 0$.

nonlinear price schedule exhibits quantity discounts; this is sufficient for our results below. The previous analysis easily extends to this setting, albeit with slightly more complex expressions. (The algebraic computations are available in an online appendix.) We summarize these in the following proposition.

Proposition 3. *In the quadratic-uniform setting with* $C(q) = cq + \frac{c_2}{2}q^2$, *for all* $c_2 \in (0, \gamma)$,

$$\frac{d}{dc_2}\left(\overline{\Pi} - \Pi^n\right) > 0,$$

$$\frac{d}{dc_2}\left(\overline{CS} - CS^n\right) < 0,$$

$$\frac{d}{dc_2}\left(\overline{\Pi} + \overline{CS} - \Pi^n - CS^n\right) > 0.$$

To be clear, an increase in c_2 lowers profit and consumer surplus, but such an increase adversely impacts profits less (and consumer surplus more) under average-price bias than in the unbiased setting. To understand the economics of the result, note that an increase in c_2 of dc_2 raises marginal cost by $q\,dc_2$. The firm in either setting increases the marginal price for higher outputs in response. In the standard setting,

$$P^n(q) = \frac{q}{2}\left(\theta_1 + c + \frac{c_2 - \gamma}{2} q\right),$$

and thus

$$\frac{d}{dc_2} P^{m\prime}(q) = \frac{q}{2};$$

half of the increase in marginal cost is passed along to unbiased consumers in higher marginal prices. In the case of average-price bias, however, increases in the price margin are less salient and an excessive amount of the marginal cost increase is passed along

$$\overline{P}(q) = cq + (\gamma + c_2)q^2,$$

and thus

$$\frac{d}{dc_2} \overline{P}'(q) = 2q.$$

Although marginal pass through is excessive, the consumer perceives only half as much pass through of the marginal cost increase:

$$\frac{d}{dc_2}\left(\frac{\overline{P}(q)}{q}\right) = q.$$

It follows that a slight convexity in costs will tilt a firm toward a preference for consumers with average-price bias. The reverse is the case for changes in fixed per-consumer costs.

Fixed costs: Suppose that marginal costs are constant, but there is a positive per-consumer fixed cost: $c_2 = 0$, $c_0 > 0$. In the general case in which not all consumers are served under either regime, $q^n(\theta)$ and $\overline{q}(\theta)$ exhibit positive jumps for the marginal consumer types so the algebra to derive Π^n and $\overline{\Pi}$ is decidedly more tedious. Nonetheless, computing a general comparative static result for c_0 is straightforward if c_0 is not too large relative to the gross profit earned for the highest type.[13]

[13] The firm will only participate in the average-price bias regime if it makes profit for the highest type: $\theta_1 - c > 2\sqrt{c_0\gamma}$. The requirement of Proposition 4 is slightly stronger. The algebraic computations are available in an online appendix.

Proposition 4. *In the quadratic-uniform setting with $C(q) = cq + \frac{c_2}{2}q^2$, In the quadratic-uniform setting with $C(q) = c_0 + cq$, for c_0 sufficiently small ($\sqrt{c_0\gamma}(2+\sqrt{2}) < \theta_1 - c$),*

$$\frac{d}{dc_0}\left(\overline{\Pi} - \Pi^n\right) < 0,$$

$$\frac{d}{dc_0}\left(\overline{CS} - CS^n\right) > 0,$$

$$\frac{d}{dc_0}\left(\overline{\Pi} + \overline{CS} - \Pi^n - CS^n\right) < 0.$$

The introduction of a small fixed cost reduces profitability and consumer surplus, but such a cost increase adversely impacts profits more (and consumer surplus less) under average-price bias than in the textbook setting. Fixed costs can be easily passed onto neoclassical consumers with increases in the fixed component of a price schedule, but introducing a fixed increase in $\overline{P}$ will be misunderstood as an increase in marginal price for consumer with average-price bias. Consequently, less of the fixed fee will be passed along to biased consumers which is evident in the price functions:

$$P^n(q) = \frac{1}{2}c_0 + \frac{q}{2}(\theta_1 + c) - \frac{\gamma}{4}q^2,$$

$$\overline{P}(q) = cq + \gamma q^2.$$

Half the fixed cost is passed along to neoclassical consumers, yet it is entirely absorbed by firms serving average-price biased consumers (though a measure $2\sqrt{c_0\gamma}$ of biased consumers will be excluded from the increase in fixed cost). Firms fare worse following the introduction of a fixed cost when facing a population of biased consumers.

5 Extensions and concluding comments

A few extensions suggest themselves. First, it is plausible that the population of consumers contains some who are neoclassical and others that suffer from average-price bias. This is certainly consistent with the experimental work in de Bartolome (1996) and in recent work by Shaffer (2018) which indicates distinct behavior types. How should the firm alter its nonlinear price schedule if it must offer the same schedule to all consumer types?

A second extension is to take seriously our model of adaptive consumer learning with a misspecified model, and ask how the firm would want to manipulate the consumer *on the path to a steady state*. In a fully dynamic model, the firm may find it optimal to depart from the steady-state contact in Proposition 1 to better extract curvature rents and improve profits by manipulating the consumer's learning dynamics along some path.

A third direction for future research is a deeper analysis of recommendation mechanism design problem when facing a population of average-price-bias consumers – an analysis that would allow a firm to inform consumers of their bias as part of the communication mechanism. A general investigation of such mechanism design with behavioral agents seems potentially valuable.

References

BARTOLOME, CHARLES (1995). Which tax rate do people use: average or marginal? *Journal of Public Economics*, **56**, 79-96.

CARBAJAL, JUAN CARLOS AND JEFFREY ELY (2016). A model of price discrimination under loss aversion and state-contingent reference points. *Theoretical Economics*, **11**, 455-485.

CHETTY, RAJ, ADAM LOONEY AND KORY KROFT (2009). Salience and taxation: Theory and evidence. *American Economic Review*, **99**: 1145-1177.

DELLA VIGNA, STEFANO (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, **47(2)**: 315-72.

EVANS, GEORGE AND SEPPO HONKAPOHJA (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press.

FARHI, EMMANUEL AND XAVIER GABAIX (2018). "Optimal Taxation with Behavioral Agents," June 7, 2018.

FILIZ-OBAY, EMEL (2012). Incorporating unawareness into contract theory. *Games and Econmic Behavior*, **76**: 181-194.

HAHN, JONG-HEE, JINWOO KIM, SANG-HYUN KIM ABD JOHONG LEE (2018). Price discrimination with loss averse consumers. *Economic Theory*, **65**: 681-728.

HEIDHUES, PAUL AND BOTOND KÖSZEGI (2017). Naïveté-based Discrimination, *Quarterly Journal of Economics*, **132**, 1019-1054.

HOLMSTRÖM, BENGT (1999). Managerial incentive problems: a dynamic perspective. *Review of Economic Studies*, **66**, 169-182.

HOSSAIN, TANJIM AND JOHN MORGAN (2006). .... Plus shipping and handling: Revenue (non) equivalence in field experiments on eBay. *Advances in Economic Analysis and Policy*, **6(2)**: Article 3.

ITO, KOICHIRO (2014). Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. *American Economic Review*, **104** 537–563.

KÖSZEGI, BOTOND AND MATTHEW RABIN (2006). A model of reference-dependent preferences. *Quarterly Journal of Economics*, **121**, 1133-1165.

KÖSZEGI, BOTOND (2013). Behavioral contract theory. *Journal of Economic Literature*, **52(4)**: 1075-1118.

LIEBMAN, JEFFREY AND RICHARD ZECKHAUSER (2004). Schmeduling. Working paper. October 2004.

MALMENDIER, ELRIKE AND YOUNG HAN LEE (2011). The bidder's curse. *American Economic Review*, **101**: 749-787.

MASKIN, ERIC AND JOHN RILEY (1984). Monopoly with Incomplete Information. *RAND Journal of Economics*, **15**, 171–196.

MUSSA, MICHAEL AND SHERWIN ROSEN (1978). Monopoly and Product Quality. *Journal of Economic Theory*, **18** 301–317.

SHAFFER, BLAKE (2018). Rational, lazy or confused? Evidence of misperception in consumer responsiveness to nonlinear prices. Working paper, September 2018.

TIROLE, JEAN (1988). *The Theory of Industrial Organization*, MIT Press, Cambridge, MA.

Online Appendix

Proof of Proposition 3: We first compute the formulas for the profits and consumer surpluses under each regime under the assumption that $C(q) = cq + \frac{\gamma}{2}q^2$.

Profits and surplus in the standard model: The optimal allocation is given by pointwise maximization of the virtual surplus function. We obtain

$$q^n(\theta) = \max\left\{0, \frac{2\theta - c - \theta_1}{\gamma + c_2}\right\}.$$

For participating types, the inverse allocation is

$$\vartheta^n(q) = \frac{c + q(\gamma + c_2) + \theta_1}{2}.$$

The marginal consumer type is determined at the greatest type for which $q^n(\theta^n) = 0$ if such a type exists, and $\theta^n = \theta_0$, otherwise.

$$\theta^n = \max\left\{\theta_0, \frac{c + \theta_1}{2}\right\}.$$

The optimal nonlinear price schedule can be found by integrating the consumer's marginal utility function:

$$P^n(q) = \int_0^q u_q(x, \vartheta^n(x))dx = \frac{q}{2}\left(\theta_1 + c + \frac{c_2 - \gamma}{2}q\right).$$

From here, we can directly compute type-θ consumer surplus and profit.

$$U^n(\theta) = u(q^n(\theta), \theta) - P^n(q^n(\theta)) = \frac{(c + \theta_1 - 2\theta)^2}{4(\gamma + c_2)}.$$

$$\pi^n(\theta) = u(q^n(\theta), \theta) - C(q^n(\theta)) - U^n(\theta) = \frac{(c + 2\theta - 3\theta_1)(c + \theta_1 - 2\theta)}{4(\gamma + c_2)}.$$

Integrating these expressions over the interval $[\theta^n, \theta_1]$ yields (after simplification)

$$CS^n = \int_{\theta^n}^{\theta_1} U^n(\theta)d\theta = \frac{(\theta_1 - c)^3}{24(\gamma + c_2)(\theta_1 - \theta_0)},$$

$$\Pi^n = \int_{\theta^n}^{\theta_1} \pi^n(\theta)d\theta = \frac{(\theta_1 - c)^3}{12(\gamma + c_2)(\theta_1 - \theta_0)}.$$

Profits and surplus in the average-price bias model: The optimal allocation under average-price bias satisfies

$$\overline{q}(\theta) = \max\left\{0, \frac{\theta - c}{2\gamma + c_2}\right\}.$$

The associated inverse allocation for participating types is therefore

$$\overline{\vartheta}(q) = \max\left\{0, c + (c_2 + 2\gamma)q\right\}.$$

The marginally-participating type is $\overline{\theta} = \max\{\theta_0, c\}$. Following Proposition 1, the optimal nonlinear price is

$$\overline{P}(q) = cq + (\gamma + c_2)q^2.$$

Consumer surplus for type θ is

$$\overline{cs}(\theta) = u(\overline{q}(\theta), \theta) - P(\overline{q}(\theta)) = \frac{\gamma(\theta - c)^2}{2(2\gamma + c_2)^2}.$$

Profit is accordingly

$$\overline{\pi}(\theta) = P(\overline{q}(\theta)) - C(\overline{q}(\theta)) = \frac{(\theta - c)^2}{2(2\gamma + c_2)}.$$

Integrating these expressions over the interval $[\overline{\theta}, \theta_1]$ yields (after simplification)

$$\overline{CS} = \int_{\overline{\theta}}^{\theta_1} \overline{cs}(\theta) d\theta = \frac{\gamma(\theta_1 - c)^3}{6(\theta_1 - \theta_0)(2\gamma + c_2)},$$

$$\overline{\Pi} = \int_{\overline{\theta}}^{\theta_1} \overline{\pi}(\theta) d\theta = \frac{(\theta_1 - c)^3}{6(\theta_1 - \theta_0)(2\gamma + c_2)}.$$

Combining terms: We combine the above expressions into the objects under study and simplify:

$$\overline{\Pi} - \Pi^n = \frac{c_2(\theta_1 - c)^3}{12(\theta_1 - \theta_0)(2\gamma + c_2)(\gamma + c_2)},$$

$$\overline{CS} - CS^n = -\frac{c_2^2(\theta_1 - c)^3}{24(\theta_1 - \theta_0)(2\gamma + c_2)^2(\gamma + c_2)},$$

$$(\overline{\Pi} + \overline{CS}) - (\Pi^n + CS^n) = \frac{c_2(4\gamma + c_2)(\theta_1 - c)^3}{24(\theta_1 - \theta_0)(2\gamma + c_2)^2(\gamma + c_2)}.$$

Differentiating with respect to c_2 yields

$$\frac{d}{dc_2}\left(\overline{\Pi} - \Pi^n\right) = \frac{(2\gamma^2 - c_2^2)(\theta_1 - c)^3}{12(\theta_1 - \theta_0)(2\gamma + c_2)^2(\gamma + c_2)^2},$$

$$\frac{d}{dc_2}\left(\overline{CS} - CS^n\right) = \frac{c_2(c_2^2 - 2c_2\gamma - 4\gamma^2)(\theta_1 - c)^3}{24(\theta_1 - \theta_0)(2\gamma + c_2)^3(\gamma + c_2)^2},$$

$$\frac{d}{dc_2}\left(\overline{\Pi} + \overline{CS} - \Pi^n - CS^n\right) = \frac{(8\gamma^3 - 6c_2^2\gamma - c_2^3)(\theta_1 - c)^3}{24(\theta_1 - \theta_0)(2\gamma + c_2)^3(\gamma + c_2)^2}.$$

Noting that $\theta_1 - \theta_0 > 0$ and $\theta_1 - c > 0$, we have the following relationships:

$$\frac{d}{dc_2}\left(\overline{\Pi} - \Pi^n\right) > 0 \iff c_2 < \sqrt{2}\gamma,$$

$$\frac{d}{dc_2}\left(\overline{CS} - CS^n\right) < 0 \iff c_2^2 - 2c_2\gamma - 4\gamma^2 < 0,$$

$$\frac{d}{dc_2}\left(\overline{\Pi} + \overline{CS} - \Pi^n - CS^n\right) > 0 \iff 8\gamma^3 > c_2^2(c_2 + 6\gamma).$$

By assumption, $c_2 < \gamma$, so all three inequalities are satisfied. □

Proof of Proposition 4: We first compute the formulas for the profits and consumer surpluses under each regime under the assumption that $C(q) = c_0 + cq$.

Profits and surplus in the standard model: Participating consumers will choose (as in the linear-cost case)

$$q^n(\theta) = \frac{2\theta - c - \theta_1}{\gamma},$$

and the inverse allocation is

$$\vartheta^n(q) = \frac{c + q\gamma + \theta_1}{2}.$$

Now, however, the allocation for the marginal consumer will exhibit a discontinuous jump. The marginal consumer type is determined at the point where virtual surplus is zero:

$$u(q^n(\theta^n), \theta^n) - cq^n(\theta^n) - (\theta_1 - \theta)q^n(\theta^n) = c_0.$$

Solving for θ^n, we obtain

$$\theta^n = \frac{c + \theta_1 + \sqrt{2\gamma c_0}}{2}.$$

Thus, the optimal allocation is

$$q^n(\theta) = \begin{cases} \frac{2\theta - c - \theta_1}{\gamma} & \text{if } \theta \geq \theta^n \\ 0 & \text{otherwise.} \end{cases}$$

Because the marginal consumer will obtain zero rent, we can determine the optimal price schedule by integrating the consumer's marginal utility accordingly:

$$P^n(q) = u(q^n(\theta^n), \theta^n) + \int_{q^n(\theta^n)}^{q} u_q(x, \vartheta^n(x))dx = \frac{1}{2}c_0 + \frac{q}{2}(\theta_1 + c) - \frac{\gamma}{4}q^2.$$

From here, we can directly compute type-θ consumer surplus and profit.

$$U^n(\theta) = u(q^n(\theta), \theta) - P^n(q^n(\theta)) = \frac{(2\theta - \theta_1 - c)^2}{4\gamma} - \frac{c_0}{2},$$

$$\pi^n(\theta) = u(q^n(\theta), \theta) - C(q^n(\theta)) - U^n(\theta) = \frac{1}{4\gamma}\left((3\theta_1 - 2\theta - c)(2\theta - \theta_1 - c)\right) - \frac{c_0}{2}.$$

Integrating these expressions over the interval $[\theta^n, \theta_1]$ yields (after simplification)

$$CS^n = \int_{\theta^n}^{\theta_1} U^n(\theta)d\theta = \frac{4\sqrt{2}(c_0\gamma)^{\frac{3}{2}} + (\theta_1 - c)^3 - 6c_0\gamma(\theta_1 - c)}{24\gamma(\theta_1 - \theta_0)},$$

$$\Pi^n = \int_{\theta^n}^{\theta_1} \pi^n(\theta)d\theta = \frac{4\sqrt{2}(c_0\gamma)^{\frac{3}{2}} + (\theta_1 - c)^3 - 6c_0\gamma(\theta_1 - c)}{12\gamma(\theta_1 - \theta_0)}.$$

Profits and surplus in the average-price bias model: Because c_0 doesn't impact the intensive margin, the optimal allocation for participating consumers under average-price bias satisfies

$$\overline{q}(\theta) = \frac{\theta - c}{2\gamma}.$$

The associated inverse allocation for participating types is therefore

$$\overline{\vartheta}(q) = c + 2\gamma q.$$

The profit and surplus earned for participating types is given by

$$\overline{\pi}(\theta) = \frac{(\theta - c)^2}{4\gamma} - c_0,$$

$$\overline{cs}(\theta) = \frac{(\theta - c)^2}{8\gamma}.$$

The marginal consumer is determined by finding the type for which $\overline{\pi}(\overline{\theta}) = 0$, if one exists, and θ_0 otherwise:

$$\overline{\theta} = \max\left\{\theta_0, c + 2\sqrt{c_0\gamma}\right\}.$$

Integrating these expressions over the interval $[\overline{\theta}, \theta_1]$ yields (after simplification)

$$\overline{CS} = \int_{\overline{\theta}}^{\theta_1} \overline{cs}(\theta)d\theta = \frac{(\theta_1 - c)^3 - 8(c_0\gamma)^{\frac{3}{2}}}{24\gamma(\theta_1 - \theta_0)},$$

$$\overline{\Pi} = \int_{\overline{\theta}}^{\theta_1} \overline{\pi}(\theta)d\theta = \frac{16(c_0\gamma)^{\frac{3}{2}} + (\theta_1 - c)^3 - 12c_0\gamma(\theta_1 - c)}{12\gamma(\theta_1 - \theta_0)}.$$

Combining terms: We now combine the above expressions into the objects under study and simplify:

$$\overline{\Pi} - \Pi^n = \frac{c_0((8 - 2\sqrt{2})\sqrt{c_0\gamma} - 3(\theta_1 - c))}{6(\theta_1 - \theta_0)},$$

$$\overline{CS} - CS^n = \frac{c_0(3(\theta_1 - c) - (4 + 2\sqrt{2})\sqrt{c_0\gamma})}{12(\theta_1 - \theta_0)},$$

$$(\overline{\Pi} + \overline{CS}) - (\Pi^n + CS^n) = \frac{c_0((4 - 2\sqrt{2})\sqrt{c_0\gamma} - (\theta_1 - c))}{4(\theta_1 - \theta_0)}.$$

Differentiating with respect to c_0 yields

$$\frac{d}{dc_0}\left(\overline{\Pi} - \Pi^n\right) = \frac{(4 - \sqrt{2})\sqrt{c_0\gamma} - (\theta_1 - c)}{2(\theta_1 - \theta_0)},$$

$$\frac{d}{dc_0}\left(\overline{CS} - CS^n\right) = \frac{(\theta_1 - c) - (2 + \sqrt{2})\sqrt{c_0\gamma}}{4(\theta_1 - \theta_0)},$$

$$\frac{d}{dc_0}\left(\overline{\Pi} + \overline{CS} - \Pi^n - CS^n\right) = \frac{3(2 - \sqrt{2})\sqrt{c_0\gamma} - (\theta_1 - c)}{4(\theta_1 - \theta_0)}.$$

Noting that $\theta_1 - \theta_0 > 0$ and $\theta_1 - c > 0$, we have the following relationships:

$$\frac{d}{dc_0}\left(\overline{\Pi} - \Pi^n\right) < 0 \iff (4 - \sqrt{2})\sqrt{c_0\gamma} < \theta_1 - c,$$

$$\frac{d}{dc_0}\left(\overline{CS} - CS^n\right) > 0 \iff (2 + \sqrt{2})\sqrt{c_0\gamma} < \theta_1 - c,$$

$$\frac{d}{dc_0}\left(\overline{\Pi} + \overline{CS} - \Pi^n - CS^n\right) < 0 \iff 3(2 - \sqrt{2})\sqrt{c_0\gamma} < \theta_1 - c.$$

By assumption, $(2 + \sqrt{2})\sqrt{c_0\gamma} < \theta_1 - c$, so these three inequalities are satisfied. □