

NONLINEAR PRICING WITH AVERAGE-PRICE BIAS*

DAVID MARTIMORT[†] LARS STOLE[‡]

January 7, 2020

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 with the highest marginal utility and typically feature quantity premia rather than quantity discounts. These properties arise because the bias causes consumer *information rents* to be replaced 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. D42, D82, D91.

*We thank Austan Goolsbee, Bruno Jullien, Emir Kamenica, Amin Sayedi, seminar participants at the 2016 North American Econometric Society (ASSA) winter meetings, University of Chicago, Toulouse School of Economics and UC Berkeley for comments and useful suggestions on earlier drafts. We are especially grateful to Larry Samuelson and the referees for valuable comments.

[†]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., James Mirrlees (1971)), second-degree price discrimination (e.g., Michael Mussa and Sherwin Rosen (1978), Eric Maskin and John Riley (1984)), and regulation (e.g., David Baron and Roger 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 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. In Koichiro Ito (2014), spatial discontinuities in electricity service areas in California are used to identify consumers’ perceived prices; the paper finds that average price explains short-run demand variations far better than the marginal price. 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), Charles 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. While we provide a simple justification for such behavior based on the consumer’s misspecification of the price schedule, our primary goal is not to rationalize average-price bias. Rather,

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

our contribution is to take average-price bias as an exogenous feature of consumer decision making and examine the implications. 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, 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 does not lead to information rents for the consumer. That said, average-price biased consumers still earn consumer surplus through what we call *curvature* rents. The monopolist's desire to extract curvature rents rather than information rents leads to significant differences in the shape of optimal nonlinear prices. For example, in a simple 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 whether 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 and price schedules are remarkably different. 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.

Literature review

Several papers incorporate models of decision-making from behavioral economics and psychology into the standard contracting paradigms.² Juan Carlos Carbajal and Jeffrey Ely (2016) and Jong Hee Hahn, et al (2018), for example, introduce loss aversion and reference-dependent preferences (as in Botond Köszegi and Matthew Rabin (2006)) into

²Botond Köszegi (2014) provides a survey.

a nonlinear pricing model. Our paper is in the same spirit but focuses on average-price bias. More loosely related, the literature on inattention and salience provides evidence that consumers are inattentive to the correct prices: e.g., ignoring taxes that are not salient (Raj Chetty, Adam Looney and Kory Kroft (2009)), treating shipping costs on eBay differently from product prices (Tanjim Hossain and John Morgan (2006)), ignoring fixed-price offers that are dominated by eBay auctions (Elrike Malmendier and Young Han Lee (2011)), and for consumers who routinely purchase large sized products in grocery stores, ignoring dominating price promotions on small sized packages of the same product (Sofronis Clerides and Pascal Courty (2017)). Stefano Della Vigna (2009) surveys the evidence of limited attention, salience and suboptimal heuristics in field experiments.³

Two early papers are closely related to our investigation. Joel Sobel (1984) considers naive “price-taking” buyers who are identical to our average-price biased consumers, but assumes constant average costs and, importantly, that demand becomes weakly more elastic as consumption increases. In contrast, our focus is on the more natural setting where demand becomes less elastic moving downward along the demand curve and where average costs are not necessarily constant. Unlike Sobel (1984), in our setting whether bias raises or lowers profit, for example, depends upon the economics of utility and costs. The working paper of Liebman and Zeckhauser (2004) is also closely related.⁴ In their model, a monopolist sells to a consumer with one of two possible types, but is constrained to use piecewise-linear price schedules that originate from the origin. Given this restriction, the monopolist always prefers selling to biased consumers relative to unbiased ones. Such restricted schedules are suboptimal in the unbiased setting, however, as they leave rent to the low-type consumer. We emphasize that neither Sobel (1984) nor Liebman and Zeckhauser (2004) identify the key implication of average-price bias,

³A related 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. See Emel Filiz-Obay (2012), as an example in the context of insurance. In our setting, the consumer is unaware that the price schedule is nonlinear, but we do not consider the possibility the price schedule could be used to alter this belief.

⁴Liebman and Zeckhauser refer to manifestations of average-price bias as “ironing.” Emmanuel Farhi and Xavier Gabaix (2018) develop a general framework of optimal taxation with behavioral agents which is sufficiently general that it embeds the model of Liebman and Zeckhauser (2004), as well as a list of other well-studied heuristics and biases. We emphasize the specific implications of average-price bias, for which the analysis in Liebman and Zeckhauser (2004) is more directly related.

that consumer *curvature rents* are substituted for information rents.⁵ Illuminating this substitution and its consequences for the firm’s optimal price schedule and consumer allocations are the primary contributions of this paper.

2 The model

2.1 Preferences

We adopt a textbook⁶ model of nonlinear pricing by a monopolist. The firm’s per-consumer cost is $C(q)$ – convex and continuously differentiable. The type- θ consumer’s payoff is

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

where $q \in \mathcal{Q} \equiv [0, \bar{q}]$ 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 differentiable, log-concave distribution, F , with density f , and $P(q)$ is the total payment made to the firm for q units.⁷ We make standard assumptions that u is thrice continuously differentiable, strictly concave in q , $u_\theta(q, \theta)$ 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. We normalize the consumer’s outside option to zero, $u(0, \theta) = 0$.

Given unit price p , we denote the type- θ consumer’s Marshallian demand function as

$$D(p, \theta) \equiv \arg \max_{q \in \mathcal{Q}} u(q, \theta) - pq,$$

and the corresponding inverse-demand function as $p = u_q(q, \theta)$. Throughout we assume that the implied type- θ consumer’s elasticity of demand,

$$\varepsilon(q, \theta) \equiv -\frac{u_q(q, \theta)}{u_{qq}(q, \theta)q} > 0,$$

⁵In an uncirculated working paper that was discovered late in the editorial process, Phuong Hô (2017) takes a similar approach as ours to examine the implications of bias for nonlinear Ramsey-pricing. Phuong Hô (2017) does not note the replacement of information rents for curvature rents, as we have emphasized.

⁶See, for example, Jean Tirole (1988), chapter 3.

⁷We assume \bar{q} strictly exceeds the efficient level of consumption for all types.

is nonincreasing in q .⁸

In what follows, it is useful to consider the artificial setting in which a firm faces an unbiased consumer of known type θ , but is forced to offer a linear price. We will refer to this as a case of *perfect third-degree price discrimination* because the firm can perfectly condition its linear price on θ and offer $P(q|\theta) = p(\theta)q$ to a consumer of type θ . The firm's type- θ revenue in this case is

$$R(q, \theta) = qu_q(q, \theta),$$

and the assumption that $\varepsilon(q, \theta)$ is non-increasing in q implies that the firm's type- θ profit, $R(q, \theta) - C(q)$, is quasi-concave.⁹ For simplicity, we assume throughout that marginal revenue is increasing in θ , $R_{q\theta}(q, \theta) > 0$.¹⁰

2.2 Optimal nonlinear pricing without bias

As a benchmark, consider the textbook model of nonlinear pricing. The firm offers the consumer a nonlinear price schedule, $P : \mathcal{Q} \rightarrow \mathfrak{R}_+$, in order to maximize expected profits,

$$E_\theta [P(q(\theta)) - C(q(\theta))],$$

where

$$q(\theta) \in \arg \max_{q \in \mathcal{Q}} u(q, \theta - P(q)).$$

⁸A variety of demand functions satisfy this assumption including linear demand, log-linear demand, $p = a - bq^\delta$, and $p = a - b \ln q$. The reverse assumption that elasticities weakly increase in q implausibly requires that the demand curve is everywhere elastic and nowhere intersects either axis.

⁹Specifically, $\varepsilon_q(q, \theta) \leq 0$ implies

$$u_{qqq}q \leq -u_{qq} \left(1 + \frac{1}{\varepsilon(q, \theta)} \right),$$

hence

$$R_{qq}(q, \theta) = 2u_{qq} + u_{qqq}q \leq u_{qq} \left(1 - \frac{1}{\varepsilon(q, \theta)} \right),$$

which is negative for $\varepsilon(q, \theta) > 1$. Because $C'(q) > 0$, the solution to the first-order condition occurs in the strictly concave region and thus profit is globally quasi-concave.

¹⁰Given $u_{q\theta} > 0$, it is sufficient for $R_{q\theta} > 0$ that $u_{qq\theta} \geq 0$.

Define the consumer surplus for the unbiased setting as

$$U^*(\theta) \equiv \max_{q \in \mathcal{Q}} u(q, \theta) - P(q).^{11}$$

As is well known, an allocation $q(\theta)$ is implementable by some nonlinear price schedule if and only if $q(\cdot)$ is nondecreasing and $U^{*'}(\theta)$ is equal to $u_\theta(q(\theta), \theta)$ almost everywhere. The latter implies the consumer's expected surplus is

$$\mathbb{E}_\theta [U^*(\theta)] = U^*(\theta_0) + \mathbb{E}_\theta \left[u_\theta(q(\theta), \theta) \frac{1 - F(\theta)}{f(\theta)} \right],$$

which allows us to write expected profit (total surplus minus consumer surplus) as

$$\mathbb{E}_\theta [\Pi^*(\theta)] = \mathbb{E}_\theta \left[u(q(\theta), \theta) - C(q(\theta)) - \frac{1 - F(\theta)}{f(\theta)} u_\theta(q(\theta), \theta) - U^*(\theta_0) \right].$$

Assuming a regularity condition that the integrand exhibits increasing differences in (q, θ) , the pointwise maximum of q over θ is nondecreasing and represents the firm's optimal allocation:

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

The corresponding optimal nonlinear tariff, $P^*(q)$, is uniquely determined from an envelope condition and a participation constraint. The solution exhibits three properties.

1. *No surplus at the bottom.* The lowest-type consumer obtains zero surplus.
2. *No distortion at the top.* The highest-type consumes efficiently, $u_q(q^*(\theta_1), \theta_1) = C'(q^*(\theta_1))$.
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

¹¹We use * to indicate variables related to the unbiased consumers and an overline to denote the corresponding variables for biased consumers (e.g., biased consumer surplus is denoted $\bar{U}(\theta)$).

quantity discounts.

These properties are typically absent when consumers suffer from average-price bias.

3 Average-price bias

3.1 Behavioral Assumption of Average-price Bias

We assume that consumers with average-price bias equate marginal utility with the average price. Formally, for a given price schedule, $P(q)$, we assume that a consumer chooses output to satisfy

$$(2) \quad q \in \arg \max_{\tilde{q} \in \mathcal{Q}} u(\tilde{q}, \theta) - \frac{P(q)}{q} \tilde{q},$$

when such a choice exists, and $q = 0$ otherwise.¹²

We have in mind that the consumer searches via trial and error to find a choice of q for which her marginal utility equals her perception of the marginal price, $P(q)/q$, leading to a demand which satisfies (2). Suppose, for example, that the average-price biased consumer chooses $q = D(\tilde{p}, \theta)$ given a perceived unit price of \tilde{p} . If the resulting average price, $P(q)/q$, is greater than (resp., less than) the consumer's marginal utility, $u_q(q, \theta)$, we suppose that the consumer revises their belief to a slightly higher (resp., lower) unit price and reduces (resp., increases) consumption accordingly. This process repeats itself until reaching a consistent consumption level that satisfies (2) or consumption is reduced to zero. Formally, we can model such a tâtonnement by a process with steps indexed by t , requiring at each stage of consideration for which $q(t) > 0$,

$$\frac{dq(t)}{dt} = \lambda \left(u_q(q(t), \theta) - \frac{P(q(t))}{q(t)} \right) = \frac{\lambda}{q(t)} (R(q(t), \theta) - P(q(t))), \quad \lambda > 0,$$

and $\frac{dq(t)}{dt} = 0$ if $q(t) = 0$. For this simple adjustment rule, wherever the price schedule $P(q)$ intersects the revenue function $R(q, \theta)$ from below, the solution to (2) is asymptotically

¹²We assume that if there are multiple solutions to (2), then the monopolist's preferred solution is selected. This is inconsequential because the firm's optimal price schedule never exhibits multiple, positive solutions.

stable. Given that $P(\cdot)$ is nondecreasing and $R(q, \theta)$ is decreasing in the inelastic region, at least one such stable solution exists if $u_q(0, \theta) > P'(0)$. Moreover, if only one such locally stable solution exists, then the solution is globally stable when starting from any $q > 0$.¹³ If the price schedule $P(q)$ exceeds the revenue function $R(q, \theta)$ for all $q > 0$, then the globally stable solution to the adjustment process is $q = 0$.

A similar adjustment process can be founded as the result of mis-specification bias by an otherwise rational consumer. For example, suppose the consumer repeatedly makes choices, using past, noisy consumption-price data to estimate a mis-specified model of linear pricing to inform the current choice. In a previous version of this paper, we employed a simple normal-learning model as in Bengt Holmström (1999) in which the consumer incorrectly believes $P(q_t) = (p + \eta_t)q_t$, where η_t is normally-distributed noise. Each period the consumer uses past data to estimate p and then chooses q_t to equate marginal utility to the current estimate of unit price. Under some mild assumptions, consumption converges to a choice satisfying (2). More generally, this is an example of a Berk-Nash equilibrium as developed in Ignacio Esponda and Demian Pouzo (2016).¹⁴ In a Berk-Nash equilibrium, the agent chooses an optimal strategy given beliefs, and the agent's beliefs put probability one on the (possibly misspecified) set of subjective distributions over consequences that are closest to the true distribution.¹⁵ Thus, the consumer's belief is required to be the best estimate given the linear misspecification error. Esponda and Pouzo (2016) provide a more general learning foundation for Berk-Nash equilibria which is in the spirit of the normal-learning model previously sketched. We find the notion that in the presence of misspecification bias, a consumer's choice should ultimately conform to a Berk-Nash equilibrium to be persuasive, and are unable to find an alternative behavioral rule that both encapsulates the empirical evidence of average-price bias and that is internally consistent.

¹³When there is a single $\bar{q} > 0$ for which $P(\bar{q}) = R(\bar{q}, \theta)$, it is necessarily the case that P intersects R from below at \bar{q} . Thus, $R(q, \theta) > P(q)$ for all $q \in (0, \bar{q})$ and $R(q, \theta) < P(q)$ for all $q > \bar{q}$. Consequently, convergence to \bar{q} arises for any starting point $q > 0$.

¹⁴Such an equilibrium is also known as a *restricted perceptions* learning equilibrium from the macroeconomics literature; see George Evans and Seppo Honkapohja (2001).

¹⁵The metric used is weighted Kullback-Leibler divergence.

3.2 Implementability and curvature rents

Given (2) must be satisfied, we first determine which quantity allocations, $q : \Theta \rightarrow \mathcal{Q}$, are implementable and the price schedules that implement them. Compared to the setting with unbiased consumers, the answer is surprisingly 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, $q(\cdot) : \Theta \rightarrow \mathcal{Q}$ is implementable because it is not possible 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 $q > 0$, however, there are no other restrictions. An allocation $\bar{q}(\cdot)$ is implementable under average-price bias if and only if it is fully separating over positive quantities. In such a case, the implementing price schedule is

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

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

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

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

Because u is strictly concave, $\Delta(q, \theta) > 0$ for $q > 0$. Consumer surplus, given (3), is

therefore

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

We refer to $\Delta(q, \theta)q$ as the consumer's *curvature rent* as it reflects the surplus a consumer obtains due to diminishing marginal utility.¹⁶ Because of this curvature rent, a consumer's participation constraint is trivially satisfied. Indeed, if $\bar{q}(\theta_0) > 0$, then the lowest-type agent will earn a positive rent, unlike the standard nonlinear pricing model.

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

$$E_\theta [P(q(\theta)) - C(q(\theta))] = E_\theta [u_q(q(\theta), \theta)q(\theta) - C(q(\theta))].$$

This is exactly the same program as determining the optimal allocation under perfect third-degree price discrimination. That is, it is as if the firm can price conditionally on observed θ , but the firm is restricted to offering a linear price. Given our assumption that demand elasticities $\varepsilon(q, \theta)$ are nonincreasing in q , the solution to this program is determined by pointwise maximization. Given our assumption that marginal revenue is increasing in θ , the resulting allocation is fully separating.¹⁷

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

$$(4) \quad u_q(\bar{q}(\theta), \theta) - C'(\bar{q}(\theta)) = -u_{qq}(\bar{q}(\theta), \theta)\bar{q}(\theta)$$

¹⁶We use “curvature” to evoke the idea that nonlinear utility is the source of consumer surplus. This differs from the mathematical definition of curvature used in differential geometry.

¹⁷Differentiating (4),

$$\bar{q}'(\theta) = \frac{R_{q\theta}(\bar{q}(\theta), \theta)}{C''(\bar{q}(\theta)) - R_{qq}(\bar{q}(\theta), \theta)} > 0.$$

for $\bar{q}(\theta) > 0$ and $\bar{q}(\theta) = 0$ otherwise; the optimal price schedule is

$$(5) \quad \bar{P}(q) = u_q(q, \bar{\vartheta}(q))q,$$

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

Several comparisons with the standard model are now available. First, the optimal allocation is independent of the type distribution, F . As suggested above, average-price bias eliminates the costs of incentive compatibility; information rents (indicated 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,

$$E_\theta [u(q(\theta), \theta) - C(q(\theta)) - \Delta(q(\theta), \theta)q(\theta)],$$

and thus interpret the first-order condition as a marginal tradeoff between efficient consumption and reducing such 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, in sharp contrast with the standard nonlinear pricing model. That said, the marginal consumer type who purchases is the same as in the first best. In other words, on the extensive margin there is efficiency at the bottom.

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 $\bar{U}(\theta_0) = \Delta(\bar{q}(\theta_0), \theta_0)\bar{q}(\theta_0)$. This is positive if $\bar{q}(\theta_0) > 0$ which arises if $u_q(0, \theta_0) > C'(0)$. 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 $\bar{\vartheta}(q)$ represent the inverse allocation, substituting (4) into (5) reveals

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

If utility is quadratic and $C'''(q) \geq 0$, then $\bar{P}''(q) > 0$ 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 typical. Quantity premia (or increasing-block tariffs) are a common feature in electricity pricing, although prices are typically set without a singular profit motive.¹⁸ Quantity premia are also observed in a variety of profit-maximizing settings. For example, cellular phone plans historically featured higher marginal prices for minutes over a fixed monthly allocation, Michael Grubb (2015). Kenneth Manning, David Sprott and Anthony Miyazaki (1998) find quantity surcharges in grocery stores for 27 percent of the product brands offering packages of multiple sizes. And Yuri Levin, Mikhail Nediak and Andrei Bazhanov (2014) document airlines selling multiple seats at average prices which exceed the single-priced ticket.

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 every consumer who purchases earns a curvature rent.

¹⁸See Ito (2014) and Shaffer (2018) for examples.

4 Welfare comparisons

We are interested in how profits, consumer surplus and welfare differ across the standard model and the average-price bias setting. If, for example, such a bias reduces a firm's profits, we might expect firms to mitigate its presence. If such a bias reduces social welfare, changes in policy may be advisable. Given our previous observation that optimal pricing under average-price bias is equivalent to perfect third-degree price discrimination, welfare comparisons between settings with unbiased and biased consumers are equivalent to comparisons between second-degree and perfect third-degree price discrimination environments.

We have previously noted that the lowest-type consumer may earn a strictly positive rent in the average-price bias game if $\bar{q}(\theta_0) > 0$, 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.¹⁹ In order to explore these different effects, we investigate a model with quadratic preferences and uniformly-distributed heterogeneity. We then illustrate how variations from constant unit costs lead to very different welfare conclusions.

4.1 Results for Uniform-Quadratic Preferences

Consider the following *uniform-quadratic* model: (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, assume

¹⁹Sobel (1984) assumes $u(q, \theta) = \theta v(q)$, $C(q) = cq$ and the elasticity of demand is nondecreasing in q , which imply that the firm always prefers unbiased consumers. Liebman and Zeckhauser (2004) consider a two-type consumer model with restricted price schedules, which imply that the firm always prefers biased consumers. We demonstrate below that assuming demand elasticity magnitudes weakly decrease in q and allowing for unrestricted price schedules, the firm's preference for bias depends upon marginal utility and costs.

$$C(q) = cq,$$

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

$\gamma > 0$, and θ is uniformly distributed on $[\theta_0, \theta_1]$. Consequently, the θ -type consumer's demand curve is characterized by

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

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

In the standard model without bias, using (1),

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

and the corresponding optimal nonlinear price is

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

Consumer surplus for type θ is

$$CS^*(\theta) \equiv \max_{q \in \mathcal{Q}} u(q, \theta) - P^*(q) = \frac{1}{4\gamma} (\max\{0, 2\theta - \theta_1 - c\})^2,$$

with ex ante surplus

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

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

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

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

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

Now consider the average-price bias regime. The consumption allocation is given by

(4), which specializes to

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

The corresponding consumer surpluses and profits are

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

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

As in the case of unbiased consumers, expected profits are twice the expected consumer surplus, $\overline{\Pi} = 2\overline{CS}$. More remarkably, these ex ante values are equal across the regimes. The expected benefit of eliminating information rents equals the expected cost of curvature rents.

Proposition 2. *In the uniform-quadratic model,*

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

Although profits and consumer surpluses are equal across the regimes, it is worth emphasizing that the optimal allocations in the two regimes, \bar{q} and q^* , are very different. Furthermore, the marginal non-participating consumer type under average-price bias, $\theta = c$, coincides with the first-best marginal consumer; but this is different from the marginal consumer under the second-degree price discrimination setting, $\theta = \frac{1}{2}(c + \theta_1)$ which is higher. Figure 1 illustrates these differences.

4.2 Welfare for nonlinear costs

The result in Proposition 2 is knife edge and relies on constant unit costs, among other details. By changing the cost function to either be strictly convex or to contain a fixed cost, we will arrive at very different comparisons across the regimes. If marginal costs are strictly increasing, the firm can be more easily pass on marginal costs to biased

consumers (who spread the increases across all units) than to unbiased consumers (who correctly perceive the marginal price increase). As such, the introduction of strictly increasing marginal costs leads firms to prefer biased consumers over rational ones. The introduction of a fixed cost has the opposite effect. Because a fixed cost cannot be passed along to a biased consumer without leading to the perception that the marginal price is higher, the introduction of a fixed cost will lead firms to earn lower profits with biased consumers than rational ones.

To provide the precise details for the above argument, we consider 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 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 $c_2 \in (0, \gamma)$,*

$$\begin{aligned}\frac{d}{dc_2} (\bar{\Pi} - \Pi^*) &> 0, \\ \frac{d}{dc_2} (\overline{CS} - CS^*) &< 0, \\ \frac{d}{dc_2} (\bar{\Pi} + \overline{CS} - \Pi^* - CS^*) &> 0.\end{aligned}$$

An increase in c_2 lowers profit and consumer surplus in both regimes, but such an increase adversely impacts profits less (and consumer surplus more) under average-price bias than in the unbiased setting. An increase in c_2 by 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,

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

and thus

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

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 a greater amount of the marginal cost increase is passed along:

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

and thus

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

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

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

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, $q^*(\theta)$ and $\bar{q}(\theta)$ exhibit positive jumps for the marginal consumer types so the algebra to derive Π^* and $\bar{\Pi}$ is more involved.²⁰ Nonetheless, computing a general comparative static result for c_0 is straightforward if c_0 is not too

²⁰In the general analysis of section 3, we assumed that $C(q)$ was convex and continuous through the origin – i.e., no fixed costs. Introducing a fixed cost requires a straightforward but tedious modification for the determination of the marginal consumer. The details are provided in the online appendix.

large relative to the gross profit earned for the highest type.²¹

Proposition 4. *In the quadratic-uniform setting with $C(q) = c_0 + cq$, for c_0 sufficiently small*

$$\begin{aligned}\frac{d}{dc_0} (\bar{\Pi} - \Pi^*) &< 0, \\ \frac{d}{dc_0} (\overline{CS} - CS^*) &> 0, \\ \frac{d}{dc_0} (\bar{\Pi} + \overline{CS} - \Pi^* - CS^*) &< 0.\end{aligned}$$

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 rational consumers with increases in the fixed component of a price schedule, but introducing a fixed increase in \bar{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:

$$\begin{aligned}P^*(q) &= \frac{1}{2}c_0 + \frac{q}{2}(\theta_1 + c) - \frac{\gamma}{4}q^2, \\ \bar{P}(q) &= cq + \gamma q^2.\end{aligned}$$

Half the fixed cost is passed along to rational consumers, yet it is entirely absorbed by firms serving average-price biased consumers. 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)

²¹A sufficient condition is that $(2 + \sqrt{2})\sqrt{c_0\gamma} < \theta_1 - c$. The algebraic computations are available in an online appendix.

and in recent work by Shaffer (2018) which identified distinct behavioral types. How should the firm alter its nonlinear price schedule if it must offer the same schedule to all consumers? Second, given a model of adaptive consumer learning with misspecification, how should the firm 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.

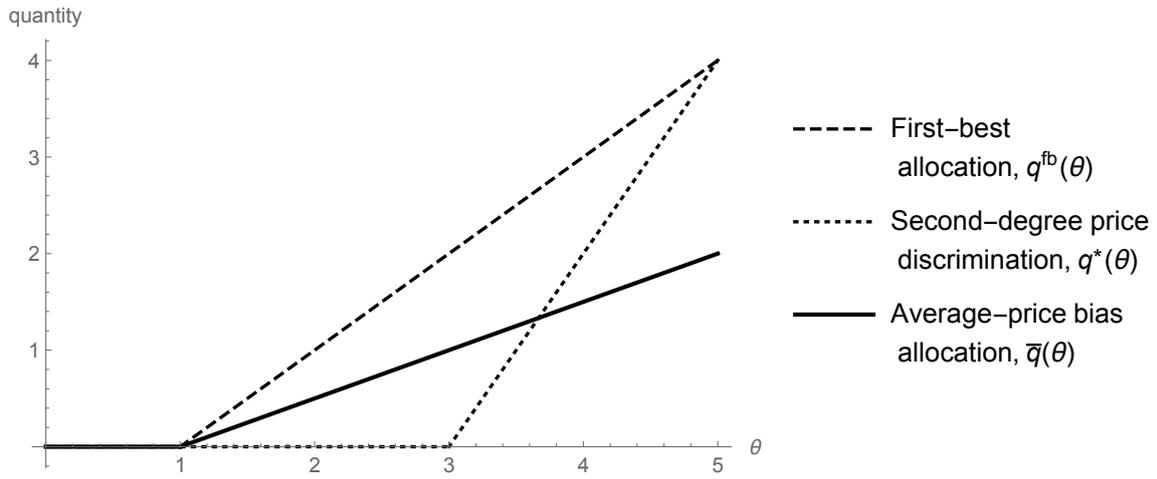


Figure 1: First-best allocation, $q^{fb}(\theta) = \max\{0, \theta - 3.25\}$; second-degree price discrimination allocation, $q^*(\theta) = \max\{0, 2\theta - 8.25\}$; “average-price bias” allocation, $\bar{q}(\theta) = \frac{1}{2} \max\{0, \theta - 3.25\}$. In all cases, $\theta_0 = 3$, $\theta_1 =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

References

- BARON, DAVID AND ROGER MYERSON (1982). A model of price discrimination under loss aversion and state-contingent reference points. *Theoretical Economics*, **11**, 455-485.
- CARBAJAL, JUAN CARLOS AND JEFFREY ELY (2016). Regulating a monopolist with unknown costs. *Econometrica*, **50**, 911-930.
- CHETTY, RAJ, ADAM LOONEY AND KORY KROFT (2009). Saliency and taxation: Theory and evidence. *American Economic Review*, **99**: 1145-1177.
- CLERIDES, SOFRONIS AND PASCAL COURTY (2017). Sales, quantity surcharge, and consumer inattention. *Review of Economics and Statistics*, **99(2)**: 357-370.
- DE BARTOLOME, CHARLES (1995). Which tax rate do people use: average or marginal? *Journal of Public Economics*, **56**, 79-96.
- DELLA VIGNA, STEFANO (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, **47(2)**: 315-72.
- ESPONDA, IGNACIO AND DEMIAN POUZO (2016). Berk-Nash equilibrium: A Framework for modeling agents with misspecified models. *Econometrica*, **84(3)**: 1093-1130.
- 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. Working paper, June 7, 2018.
- FILIZ-OBAY, EMEL (2012). Incorporating unawareness into contract theory. *Games and Economic Behavior*, **76**: 181-194.
- GERSTNER, EITAN AND JAMES HESS (1987). Why do hot dogs come in packs of 10 and buns in 8s or 12s? A demand side investigation. *Journal of Business*, **60(4)**: 491-517.
- GRUBB, MICHAEL (2015). Overconfident Consumers in the Marketplace. *Journal of Economic Perspectives*, **29(4)**: 9-36.
- HAHN, JONG-HEE, JINWOO KIM, SANG-HYUN KIM AND JOHONG LEE (2018). Price discrimination with loss averse consumers. *Economic Theory*, **65**: 681-728.
- HÔ, PHUONG (2017). Nonlinear pricing, biased consumers and regulatory policy. Working paper, May 13, 2017.
- 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.
- LEVIN, YURI; MIKHAIL NEDIAK AND ANDREI BAZHANOV (2014) Quantity Premiums and Discounts in Dynamic Pricing, *Operations Research*, **62(4)**: 846-863.
- 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.
- MANNING, KENNETH, DAVID SPROTT AND ANTHONY MIYAZAKI (1998). Consumer response to quantity surcharges: Implications for retail price setters. *Journal of Retailing*, **74(3)**: 373-399.
- MASKIN, ERIC AND JOHN RILEY (1984). Monopoly with Incomplete Information. *RAND Journal of Economics*, **15**, 171–196.
- MIRRLEES, JAMES (1971). An exploration in the theory of optimal income taxation. *Review of Economic Studies*, **38(2)**: 175-208.
- 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.
- SOBEL, JOEL (1984). Non-linear prices and price-taking behavior. *Journal of Economic Behavior and Organizations*, **5** 387-396.
- 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 $C(q) = cq + \frac{c_2}{2}q^2$.

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

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

For participating types, the inverse allocation is

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

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

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

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

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

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

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

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

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

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

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

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

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

The associated inverse allocation for participating types is therefore

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

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

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

Consumer surplus for type θ is

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

Profit is accordingly

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

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

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

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

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

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

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

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

Differentiating with respect to c_2 yields

$$\frac{d}{dc_2} (\bar{\Pi} - \Pi^*) = \frac{(2\gamma^2 - c_2^2)(\theta_1 - c)^3}{12(\theta_1 - \theta_0)(2\gamma + c_2)^2(\gamma + c_2)^2},$$

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

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

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

$$\frac{d}{dc_2} (\bar{\Pi} - \Pi^*) > 0 \iff c_2 < \sqrt{2}\gamma,$$

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

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

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

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)

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

and the inverse allocation is

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

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:

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

Solving for θ^* , we obtain

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

Thus, the optimal allocation is

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

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

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

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

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

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

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

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

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

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

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

The associated inverse allocation for participating types is therefore

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

The profit and surplus earned for participating types is given by

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

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

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

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

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

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

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

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

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

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

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

Differentiating with respect to c_0 yields

$$\frac{d}{dc_0} (\bar{\Pi} - \Pi^*) = \frac{(4 - \sqrt{2})\sqrt{c_0\gamma} - (\theta_1 - c)}{2(\theta_1 - \theta_0)},$$

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

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

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

$$\frac{d}{dc_0} (\bar{\Pi} - \Pi^*) < 0 \iff (4 - \sqrt{2})\sqrt{c_0\gamma} < \theta_1 - c,$$

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

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

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