European Journal of Operational Research 306 (2023) 795–809

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Decision Support

Price discrimination with robust beliefs

Jun Han[a], Thomas A. Weber[b,*]

[a] École Polytechnique Fédérale de Lausanne, Station 5, Lausanne CH-1015, Switzerland
[b] Chair of Operations, Economics and Strategy, École Polytechnique Fédérale de Lausanne, Station 5, Lausanne CH-1015, Switzerland

ARTICLE INFO

Article history:
Received 1 April 2021
Accepted 16 August 2022
Available online 21 August 2022

Keywords:
Ambiguity
Pricing
Relative regret
Robust optimization
Screening

ABSTRACT

This paper considers the problem of second-degree price discrimination when the type distribution is unknown or imperfectly specified by means of an ambiguity set. As robustness measure we use a performance index, equivalent to relative regret, which quantifies the worst-case attainment ratio between actual payoff and ex-post optimal payoff. We provide a simple representation of this performance index, as the lower envelope of two boundary performance ratios, relative to beliefs that lie at the boundary of the ambiguity set. A characterization of the solution to the underlying robust identification problem is given, which leads to a robust product portfolio, for which we also determine the worst-case performance over all possible consumer types. For a standard linear-quadratic specification of the robust screening model, a worst-case performance index of 75% guarantees that the robust product portfolio exhibits a profitability that lies within a 25%-band of an ex-post optimal product portfolio, over all possible model parameters and beliefs. Finally, a numerical comparison benchmarks the robust solution against a number of alternative belief heuristics.

1. Introduction

Second-degree price discrimination (or "screening") is widely used by firms to distinguish buyers of unknown types. To cater to the heterogeneous preferences of different buyer types and maximize expected profits, the seller offers a menu of products at different levels of quality (or quantity) and different prices. The performance of standard screening relies on the quality of the firm's beliefs about the distribution of buyer types. Naturally, in many cases the firm may have no precise idea about this distribution, especially for new products, or more established products in new markets. To cope with the implied belief ambiguity, we propose a robust identification technique for the practically important two-type second-degree price-discrimination model, based on relative regret, which provides performance guarantees over all possible beliefs and types. For the popular linear-quadratic parametrization of the general model a worst-case performance guarantee of 75% is obtained, with respect to all possible two-type distributions.

There are three key motivations for our study. The first is the widely known "Wilson doctrine" (1987) which calls for robust trading mechanisms that are insensitive to the informational assumptions usually required in traditional models.[1] For example, depending on the distribution of consumer types it may sometimes be ex-post optimal to not serve certain customers, yet this proves to never be best under a robust approach. Our second motivation is to devise a method which provides *relative* performance guarantees, in terms of a fraction of what would be possible in the absence of model uncertainty, for which we deploy a performance index based on relative regret. As we show, relative performance guarantees provide balanced performance ratios with worst-case scenarios attained at the boundaries of the belief spectrum. The third motivation for our approach is to preserve a homotopic connection of the robust solution to the nonrobust solution, which is achieved by means of the ambiguity set. As its diameter varies from 0 to 1, the solution changes continuously from classical to fully robust.

To the best of our knowledge, this paper is the first to address the issue of second-degree price discrimination in a relative-regret framework for a seller with ambiguous beliefs about the distribution of consumer types. We already can note at this point that the

* Corresponding author.
E-mail addresses: jun.han@epfl.ch (J. Han), thomas.weber@epfl.ch (T.A. Weber).

[1] In a more explicit manner, Wilson (1992, p. 271) notes "... significant gaps remain. The theory relies on (...) strong assumptions, such as (...) common knowledge of probability distributions, (...). These assumptions facilitate theoretical work but they hamper empirical and experimental studies, since they are never precisely true in practice and little has been done to establish the robustness of the predictions."

https://doi.org/10.1016/j.ejor.2022.08.022

simple maximin principle is not useful in our case, for it generally leads to no price-discrimination at all, whereas when pursuing relative robustness the firm retains a strong motivation to provide a nondegenerate product portfolio.

1.1. Literature

In his "theory of differential rates," Watkins (1916) recognized early on that selling two books "differing only in the quality of the paper and binding to the extent of a few cents, is a case of differentiation" (p. 693) and that "the determination of value may be either on the side of supply or on that of demand" (ibid., p. 694). He thus foreshadows the analysis by Pigou (1920, Ch. 14) of a "discriminating monopoly" with three degrees of possible price discrimination, a classification that is still in use today. While first-degree ("perfect") and third-degree price discrimination techniques rely on the seller's ability to distinguish observable characteristics of potential buyers, second-degree price discrimination corresponds to letting consumers self-select into groups. Indeed, Chandler (1938) points out that "[t]he wide differentials in the elasticities of different customers' demands for loans at an individual bank render rate discrimination profitable" (pp. 5–6). He goes on to note that "the practice of discrimination is facilitated by the high degree of secrecy surrounding the terms of contracts" (ibid., p. 6), which is one form of "non-transferability" (Pigou, 1920, p. 247). Cassady (1946) provides a first comprehensive overview of price discrimination in its different facets, and in our focus on nonobservable consumer types we are therefore concerned with his "indirect methods" that involve menus of differentiated products from which consumers individually select their most preferred alternatives. The implied individual-rationality constraints (i.e., a customer must want to purchase *something* and cannot be forced to do so) and incentive-compatibility constraints (i.e., a customer purchases only a most preferred product from the set of alternatives) correspond to Watkin's intuition that the monopolist's problem must be solved subject to the consumers' free will in their choices. Lewis (1941) discusses this self-selection problem for energy consumption where larger consumers are charged a fixed price and a low per-unit charge, and "[t]o avoid frightening off the smaller customers it is also customary to offer as an alternative to the two-part tariff a single variable charge, somewhat higher than the variable charge of the two-part tariff" (pp. 262–263). As a generalization of two-part tariffs, one can use "block tariffs" by adjusting the variable charge for different quantity blocks consumed. The latter gives rise to the idea of nonlinear pricing, introduced by Mussa & Rosen (1978)—based on the techniques developed by Mirrlees (1971) in the context of optimal taxation. To systematically cope with the self-selection issue, and thus to "screen" (i.e., separate) consumers according to their privately known characteristics (or "types") (Rothschild & Stiglitz, 1976; Stiglitz, 1975) it is necessary to propose a menu of options to consumers (Maskin & Riley, 1984).[2] Ever since, second-degree price-discrimination techniques have been used extensively in applications for the optimal design of product portfolios (see, e.g., Anderson & Dana, 2009; Moorthy, 1984; Villas-Boas, 1998).

The revenue-management literature has addressed pricing problems with unknown demand functions (see, e.g., Besbes & Zeevi, 2009; Doan et al., 2020). In the field of robust mechanism design, two objectives are commonly examined: maximin outcome and minimax regret.[3] As suggested by Wald (1939, 1945) who uses the logic of zero-sum games by Neumann & Morgenstern (1944), a distribution-free maximization of the worst-case payoff is extensively studied, among others, by Pınar & Kızılkale (2017). Carrasco et al. (2018) examine an optimal selling mechanism assuming that certain moments of the type distribution are available; similarly, under the assumption of known marginals of a multidimensional type distribution, Carroll (2017) shows that the optimal robust multidimensional screening mechanism for consumers with additively separable utilities leads to a decomposition approach along the different type dimensions. This maximin criterion is known to be conservative. For example, Bergemann & Schlag (2008) reveal that the price, which maximizes the minimum profit, is equal to the lowest consumer valuation, provided it lies above the firm's marginal cost. The well-established minimax-regret approach was introduced by Savage (1951). In the literature on monopoly pricing, Bergemann & Schlag (2008) propose a random and deterministic monopoly pricing policy for a single product based on minimax regret, while Caldentey et al. (2017) also use this criterion to examine dynamic pricing. For a recent survey of robust mechanism design we refer the reader to Carroll (2019).

In view of making the price-discrimination mechanism robust with respect to distributional assumptions, our approach is predicated upon minimizing the maximal *relative regret*, or equivalently, maximizing the attainment ratio, since the approach is not as conservative as the maximin-outcome criterion and naturally implies the relative performance guarantee. In a few instances, minimax relative regret has appeared in the pricing literature before: Eren & Maglaras (2010) examine monopoly pricing, as well as two-period dynamic pricing, without knowing the consumers' value distribution. In computer science, a "competitive ratio"—akin to relative regret—was first used for the performance evaluation of online vs. offline algorithms (Ben-David & Borodin, 1994; Boyar et al., 2015; Sleator & Tarjan, 1985). Relative regret has also been used for evaluating supply-chain performance across multiple scenarios in operations management; see, e.g., Kouvelis & Yu (1997). Goel et al. (2009) use relative regret to determine fair allocations with respect to all symmetric concave and increasing social welfare functions. The key there is a simple representation of the relative fairness measure as a function of extremal prefix functions. Under quite different assumptions, notably in the absence of any type of convexity assumptions, we derive a similar simple representation of the performance index, together with a characterization of the optimal robust solution. The latter provides performance guarantees with respect to all possible beliefs the firm might hold over the type distribution. More recently, Weber (2022) studies a general relative-regret framework for decision-making in the absence of distributional assumptions, obtaining a representation of the performance index as the minimum of two boundary performance ratios based on a semi-ordering of the state space; however, the assumptions on the objective function required there are not satisfied in our setting. There is a stream of literature studying prior-independent auctions to obtain a relative performance guarantee or bounds of the competitive ratio compared with optimal revenues such as Azar & Micali (2013); Dhangwatnotai et al. (2015) and Fu et al. (2015). In particular,

[2] By the "revelation principle" the designer of this "mechanism" (in the form of an optimal menu of options, each of which is characterized by a price and by an attribute such as quality or quantity) can restrict attention to "truth-revealing" mechanisms whereby each agent's choice is consistent with his type. An important tacit assumption is that the mechanism designer (or "principal") can fully commit to the proposed scheme (characterized by an allocation function mapping a message space to an outcome space), so that any agent's ex-post allocation is determined entirely by an agent's message (i.e., choice) communicated *ex ante* to the principal, without the possibility of opportunistic renegotiation *ex post*.

[3] Distributional variants of regret, in the form of "ambiguity aversion" and "regret aversion"—the former induced by the possibility of an incorrect belief and the latter measured by an ex-post cost based on (absolute) regret—have been used for nonlinear pricing with continuous types (Zheng et al., 2015) and discrete types (Wong, 2020), respectively. Destan & Yılmaz (2020) consider "inequity aversion" in nonlinear pricing.

Allouah & Besbes (2020) derive upper and lower bounds for the maximin ratio for some classes of distributions such as regular distributions.[4] Finally, our paper is also related to robust optimization, which addresses parameter uncertainties (Ben-Tal et al., 2009), for example, with the aid of primal-dual methods (see, e.g., Caprari et al., 2019).

1.2. Outline

The remainder of this paper is organized as follows. Section 2 presents the two-type price-discrimination problem and its ex-post optimal solution, in terms of an optimal menu of products and prices. Section 3 introduces the performance index as the minimal performance ratio of a candidate belief over all possible beliefs. A simple representation of the performance index is followed by a characterization of the solution to the firm's robust identification problem. The worst-case performance of our belief-robust solution leads to a-priori performance guarantees over all possible demand realizations. Finally, we provide an extension of the model for general ambiguity sets. Section 4 illustrates our results using a standard parametrization of the model, for which closed-form solutions are obtained. We also provide a numerical comparison of the solution behavior against several standard belief heuristics the firm might pursue. Section 5 concludes.[5]

2. Model

A monopolist faces a heterogeneous population of consumers (or "agents"), whose preferences are characterized by their respective "types" which belong to the agents' private information. The firm's goal is to design a portfolio of product offerings, characterized by the quality (or a scalar aggregate of multiple features including, e.g., quantity) of each product and its corresponding price, so as to maximize expected profits. The particular complication in this otherwise standard screening problem is that the type distribution is not known, requiring a robust approach to tackle the resulting model uncertainty.

2.1. Agents

Agents can be of two different types $\theta \in \Theta = \{\theta_L, \theta_H\}$, where θ_L, θ_H with $0 < \theta_L < \theta_H < \infty$ are given. A type-θ agent's willingness-to-pay for a given product of quality $q \in \mathcal{Q} = \mathbb{R}_+$ is described by his gross utility $u(q,\theta)$; the twice continuously differentiable function $u : \mathcal{Q} \times \mathbb{R}_+ \to \mathbb{R}_+$ is such that [6]

$$\big(\theta q = 0 \quad \Leftrightarrow \quad u(q,\theta) = 0\big) \quad \text{and} \quad \big(u(q,\theta) > 0 \quad \Rightarrow \quad u_\theta(q,\theta), u_q(q,\theta), u_{q\theta}(q,\theta) > 0\big), \tag{1}$$

for all $(q,\theta) \in \mathcal{Q} \times \mathbb{R}_+$. The preceding properties include the standard assumption that for $\theta = 0$ an agent's willingness-to-pay vanishes, and that it also vanishes if the product has zero quality. The assumption of a positive partial derivative of the agent's utility $u(q,\theta)$ with respect to θ (as long as $\theta q \neq 0$) implies that for any given product (with a positive quality level) the type-θ_H agent always has a strictly higher value than the type-θ_L agent. This persistent difference in valuations drives the firm's motivation to apply price discrimination, with the aim of extracting different values from different agents. Since this cannot be achieved by offering a single product, the firm offers a menu of products at different prices and quality levels. Finally, it is assumed that the agent's gross utility features increasing differences in (q,θ), which means that as the agent's type increases, his marginal utility for additional quality goes up,[7] which implies the following "sorting condition:"

$$q < \hat{q} \quad \Rightarrow \quad u(\hat{q},\theta_L) - u(q,\theta_L) < u(\hat{q},\theta_H) - u(q,\theta_H),$$

for all $q, \hat{q} \in \mathcal{Q}$. This relation enables the seller to implement a menu of product offerings in such a way that the agents self-select into the options designed for their respective types. This self-selection takes place based on an agent's net utility, $u(q,\theta) - p$, for a given product offering (p,q) at price $p \in \mathbb{R}$ and quality $q \in \mathcal{Q}$.

2.2. Cost

Each product of quality $q \in \mathcal{Q}$ costs the firm $C(q)$ to provide, where the continuously differentiable function $C : \mathcal{Q} \to \mathbb{R}_+$, with $C(0) = 0$, is such that

$$\big(\exists\, q > 0 : \; u(q,\theta_L) - C(q) > 0\big) \quad \text{and} \quad \big(\exists\, \bar{q} \in \mathcal{Q} : \; u(q,\theta_H) - C(q) \leq 0, \;\; \forall q > \bar{q}\big), \tag{2}$$

which implies that the firm would never want to increase quality levels indefinitely, and there is no fixed cost.[8] The first part of Eq. (2) means that for some quality level the economic surplus (i.e., the gross value minus the cost) is positive for type-θ_L consumers, while its second part states that when quality levels become too large this is no longer possible, even for type-θ_H consumers. This means that the firm can restrict attention to the interval $[0,\bar{q}]$ in its search for optimal quality levels, which by the extreme value theorem implies the existence of a finite optimum. In particular, as shown in Remark 3, the firm will always be able to guarantee itself a positive expected profit.

By Eq. (2) there exist $q_1, q_2 \in (0,\bar{q})$ such that $u_q(q_1,\theta_H) - C'(q_1) < 0 < u_q(q_2,\theta_H) - C'(q_2)$. To ensure good behavior of the optimal quality in the standard screening (cf. Corollary 1) we assume that marginal surplus $u_q(\cdot,\theta_H) - C'(\cdot)$ changes sign at most twice, so the nonempty set

$$\mathcal{I} = \{q \in \mathcal{Q} : u_q(q,\theta_H) - C'(q) > 0\}$$

is an interval. This "regularity condition" is implied if the marginal surplus $u_q(\cdot,\theta_H) - C'(\cdot)$ is quasiconcave.

Remark 1 (Relaxation of Standard Convexity Assumptions)**.** Our "nondegenerate-surplus properties" in Eq. (2) and the regularity condition are weaker than the usual requirements which include weak concavity of the agent's utility (so $u_{qq}(q,\theta) \leq 0$) and a twice differentiable cost function $C : \mathcal{Q} \to \mathbb{R}_+$ such that

$$\big(q = 0 \Leftrightarrow C(q) = 0\big) \text{ and } \big(C(q) > 0 \Rightarrow C'(q) > 0, C''(q) > 0\big),$$

together with the "Inada conditions" $C'(0) = 0$ and $\lim_{q\to\infty} C'(q) = \infty$, in the spirit of Inada (1963) and Uzawa (1963). In particular, the weaker properties of "positive surplus" and "surplus coercivity" in Eq. (2) allow for utility functions (resp., cost functions) that are nonconcave (resp., nonconvex) in quality.

Remark 2 (LQ Model)**.** A linear-quadratic (LQ) parametrization of utility and cost functions is $u(q,\theta) = \theta q$ as in Mussa & Rosen (1978), for $\theta \in \{\theta_L, \theta_H\}$ with $0 < \theta_L < \theta_H < \infty$, and $C(q) = \gamma q^2/2$,

[4] We thank an anonymous referee for noting the link to approximately optimal mechanism design, aiming at guaranteeing an objective (e.g., expected revenue) to remain close to a full-information benchmark across scenarios.

[5] Appendix A provides all proofs, and Appendix B contains a summary of the notation used in this paper.

[6] For convenience, partial derivatives of the utility function are denoted by subscripts (e.g., $u_q = \partial_q u$ or $u_{q\theta} = \partial^2_{q\theta} u$). For other partial derivatives, the standard operator notation is maintained (e.g., $\partial_{\theta_L} \Pi^*_L(\hat{\mu}^*)$).

[7] Topkis (1968) showed that the increasing-differences property of $u(q,\theta)$ implies that the quality choice by the customers must be nondecreasing in their type, thus allowing the firm to effectively separate the market.

[8] Positive fixed costs imply absolute viability constraints for the firm that are of no particular interest here; cf. Remark 3.

for any quality level $q \geq 0$, where $\gamma > 0$ is a given constant. The resulting LQ model satisfies our assumptions in Eqs. (1) and (2). Section 4 uses this parametrization to illustrate the results.

2.3. Standard screening

Assume that the firm's beliefs about the distribution of agents is such that $\mu = \mathbb{P}(\tilde{\theta} = \theta_H) \in [0,1]$ denotes the probability of a randomly drawn consumer (of random type $\tilde{\theta}$) to be of type θ_H. With this, the firm's objective is to design a menu of options $(\boldsymbol{p}, \boldsymbol{q})$, with $\boldsymbol{p} = (p_L, p_H) \in \mathbb{R}_+^2$ and $\boldsymbol{q} = (q_L, q_H) \in \mathcal{Q}^2$, so as to maximize expected profits,

$$\Pi(\boldsymbol{p}, \boldsymbol{q}) = (1-\mu)(p_L - C(q_L)) + \mu(p_H - C(q_H)),$$

subject to the constraints implied by the fact that agents can freely choose among all available options, so

$$\begin{aligned} u(q_L, \theta_L) - p_L &\geq u(q_H, \theta_L) - p_H, \\ u(q_H, \theta_H) - p_H &\geq u(q_L, \theta_H) - p_L, \end{aligned} \tag{3}$$

and that they have the option to walk away without purchasing any product, so

$$u(q_L, \theta_L) - p_L \geq 0 \quad \text{and} \quad u(q_H, \theta_H) - p_H \geq 0. \tag{4}$$

The constraints in Eq. (3) ensure "incentive compatibility" of the available options, while the constraints in Eq. (4) characterize the agents' "participation" (or "individual rationality").[9]

The firm's "standard screening problem" is therefore

$$\mathcal{M}(\mu) = \arg \max_{(\boldsymbol{p}, \boldsymbol{q}) \in \mathbb{R}_+^2 \times \mathcal{Q}^2} \Pi(\boldsymbol{p}, \boldsymbol{q}), \quad \text{s.t.} \quad (3),\ (4), \tag{5}$$

for any $\mu \in [0,1]$. The set-valued solution $\mathcal{M} : [0,1] \rightrightarrows \mathbb{R}_+^2 \times \mathcal{Q}^2$ maps the firm's belief μ to the set of optimal menus (as a subset of $\mathbb{R}_+^2 \times \mathcal{Q}^2$) for that belief. In the following lemma, we state a classical result of the solution to the standard screening problem (a proof of which is included for completeness).

Lemma 1 (Standard Screening; see, e.g., Maskin & Riley, 1984)**.** *Let $\mu \in [0,1]$. The optimal menu, defined by Eq. (5), is $\mathcal{M}(\mu) = \{(\boldsymbol{p}^*(\mu), \boldsymbol{q}^*(\mu))\}$, where*

$$\boldsymbol{p}^* = (p_L^*, p_H^*(\mu)) = \big(u(q_L^*, \theta_L), u(q_H^*, \theta_H) - (u(q_L^*, \theta_H) - u(q_L^*, \theta_L))\big), \tag{6}$$

and $\boldsymbol{q}^* = (q_L^*(\mu), q_H^*)$, with [10]

$$\begin{aligned} q_L^*(\mu) &\in \mathcal{Q}_L^*(\mu) = \arg\max_{q_L \geq 0}\{F(q_L, \mu)\}, \\ q_H^* &\in \arg\max_{q_H \geq 0}\{u(q_H, \theta_H) - C(q_H)\}, \end{aligned} \tag{7}$$

where

$$F(q_L, \mu) = (1-\mu)\big(u(q_L, \theta_L) - C(q_L)\big) - \mu\big(u(q_L, \theta_H) - u(q_L, \theta_L)\big),$$

and $\mathcal{Q}_L^* : [0,1] \rightrightarrows \mathcal{Q}$ is a correspondence.

Lemma 1 shows that Eqs. (3) and (4) reduce to two binding constraints which can be used to determine the optimal price vector $\boldsymbol{p}^*$ in Eq. (6) as a function of the optimal quality vector $\boldsymbol{q}^*$. The latter is characterized by Eq. (7). It is remarkable that the optimal quality q_H^* for the "high type" (θ_H) does not depend on the characteristics of the other type, nor on μ. By contrast, the optimal quality for the "low type" (θ_L) depends on the full type vector $\boldsymbol{\theta} = (\theta_L, \theta_H)$, as well as on the firm's belief μ, leading to a downwards distortion relative to a "first-best" pricing solution which the firm would choose if the types were observable.[11] Conversely, the optimal price p_L^* corresponds to the low-type agents' gross surplus, and thus depends on the full type vector. The optimal price p_H^* is the gross utility of the high type *minus* a nonnegative "information rent" (or discount) which depends on both agents' types, as well as on the firm's belief μ. In addition, by our assumption in Eq. (2), for any belief μ, the optimal quality $\boldsymbol{q}^*$ is finite, and in particular, $q_L^*(0)$ and q_H^* are both positive. Eq. (6) implies that the firm's profit contingent on either serving a type-θ_L agent or a type-θ_H agent as $\Pi_L(q_L) = u(q_L, \theta_L) - C(q_L)$ and $\Pi_H(\boldsymbol{q}) = u(q_H, \theta_H) - C(q_H) - \big(u(q_L, \theta_H) - u(q_L, \theta_L)\big)$, respectively. The optimal expected profit $\Pi^*(\mu) = \Pi(\boldsymbol{p}^*(\mu), \boldsymbol{q}^*(\mu))$ can be written as

$$\Pi^*(\mu) = (1-\mu)\Pi_L^*(\mu) + \mu \Pi_H^*(\mu),$$

where $\Pi_L^*(\mu) = \Pi_L(q_L^*(\mu))$, $\Pi_H^*(\mu) = \Pi_H(q_L^*(\mu), q_H^*)$, and $q_L^*(\cdot) \in \mathcal{Q}_L^*(\cdot)$ is a selection of the compact-valued image of $\mathcal{Q}_L^* : [0,1] \rightrightarrows \mathcal{Q}$. The following corollary summarizes the behavior of the optimal quality, the firm's optimal type-contingent payoffs, and its optimal profit with respect to changes in the belief μ.

Corollary 1 (Comparative Statics)**.**

(i) *There exists a critical belief,*

$$\mu_0 = \inf\{\mu \in [0,1] : 0 \in \mathcal{Q}_L^*(\mu)\} > 0, \tag{8}$$

such that $0 \in \mathcal{Q}_L^*(\mu)$, the corresponding $\Pi_L^* = 0$, and q_H^*, Π_H^* are constant in $\mu \in [\mu_0, 1]$. For $\mu < \mu_0$, any selection $q_L^*(\mu) \in \mathcal{Q}_L^*(\mu)$ is strictly decreasing, and so is $\Pi_L^*(\mu)$; the quality q_H^* stays constant, while $\Pi_H^*(\mu)$ is strictly increasing.[12] *Finally, it is always: $0 \leq \Pi_L^* \leq \Pi_H^*$.*

(ii) *The compact-valued correspondence $\mathcal{Q}_L^*(\mu)$ is single-valued in $\mu \in [0, \mu_0]$ except countably many points. $\Pi^*(\mu)$ is continuous in $\mu \in [0,1]$, while for $\mu \in [0, \mu_0]$, it is differentiable (except countably many points) and is increasing.*

Corollary 1 establishes a belief threshold μ_0 beyond which for one selection q_L^* we obtain a "shut-down solution" where only the high type is served with a positive quality level. That is, for $\mu \geq \mu_0$, the firm is able to extract the full surplus from all high-type consumers with $p_H^* = u(q_H^*, \theta_H)$, while selling no (positive-quality) product to low-type consumers. This part also describes the comparative statics, namely that the quality of the low-type product decreases as the likelihood of high-type agents in the population goes up, while q_H^* remains unaffected by the type distribution. Under the regularity condition, $\mathcal{Q}_L^*$ is well behaved and the optimal profit $\Pi^*(\mu)$ is intuitively increasing. An additional insight, particularly important in our subsequent analysis, is that when μ increases in $[0, \mu_0)$, more agents of the high type lead to a higher Π_H^* and a lower Π_L^*.

Remark 3 (Viability)**.** The assumptions in Eqs. (1) and (2) imply that the firm's optimal profit is always positive (i.e., $\Pi^*(\mu) >$

[9] The incentive-compatibility constraint (3) ensures truthful revelation of the agents' types, which means that each agent prefers the product designed for his type (L, H), at least weakly. This constraint comes without loss of generality, as a consequence of the revelation principle (Gibbard, 1973; Myerson, 1979). The latter guarantees that any mechanism can also be implemented as a 'direct revelation mechanism' where all participating agents reveal their true types. Since the firm has the option of inaction, that is, to provide a zero-quality product at zero cost (and zero price), the individual-rationality constraint (4) can always be satisfied, and all agents always participate (albeit by potentially buying a zero-quality/zero-price product).

[10] For $\mu = 0$ (i.e., in the absence of high-type agents), the high-type quality q_H^* and the corresponding price p_H^* cannot be determined by the optimization problem in Eq. (5) and in principle arbitrary—as long as they satisfy constraints of incentive compatibility and individual rationality. Here for convenience of presentation, we assume continuous completion to still define q_H^* as shown in Eq. (7) when $\mu = 0$. We note that at $\mu = 0$, the definition of q_H^* does not affect any other result in this paper.

[11] The "first-best" menu is $(\boldsymbol{p}^{**}, \boldsymbol{q}^{**})$ with $\boldsymbol{p}^{**} = (u(q_L^*(0), \theta_L), u(q_H^*, \theta_H))$ and $\boldsymbol{q}^{**} = (q_L^*(0), q_H^*)$, with $q_L^*(0)$ and q_H^* as in Eq. (7); it can be implemented in the absence of the incentive-compatibility constraint (3), extracting all the agents' surplus.

[12] A selection $q_L^*(\cdot) \in \mathcal{Q}_L^*(\cdot)$ is strictly decreasing if $\mu_1 < \mu_2 \Rightarrow q_L^*(\mu_2) < q_L^*(\mu_1)$. The monotonicity of $\Pi_L^*(\cdot) = u(q_L^*(\cdot), \theta_L) - C(q_L^*(\cdot))$ and $\Pi_H^*(\cdot) = u(q_H^*, \theta_H) - C(q_H^*) - \big(u(q_L^*(\cdot), \theta_H) - u(q_L^*(\cdot), \theta_L)\big)$ applies to *any* such selection.

0, for all $\mu \in [0,1]$). Thus, the performance ratio introduced in Section 3 is well defined; see Appendix A for details, including a nontrivial strictly positive lower bound.

3. Robust screening

While the firm knows the type values θ_L and θ_H, it does not know their likelihoods; that is, its belief $\mu = P(\tilde{\theta} = \theta_H)$ is "ambiguous" (i.e., unknown). In this section, we first introduce a robust identification problem to find a candidate belief which implies a product menu with the best performance relative to all possible beliefs, as captured by a "performance index." This problem is solved by simplifying the representation of the performance index, so it appears as the lower performance envelope with respect to extreme beliefs. We then examine the properties of the optimal performance ratio as a function of the type parameters to obtain general performance guarantees. Finally, we generalize our findings by allowing for a bounded ambiguity set, with a more focused belief structure which provides a homotopic connection between the optimal robust solution to the firm's price-discrimination problem and the solution to the standard screening problem with fully determined beliefs.

3.1. Identification problem

With the knowledge that an agent's type can be either θ_L or θ_H, to arrive at a "robust belief" about the distribution of consumer types, the firm needs to examine the performance of the menu $(\boldsymbol{p}^*, \boldsymbol{q}^*) = (\boldsymbol{p}^*(\hat{\mu}), \boldsymbol{q}^*(\hat{\mu}))$ in Lemma 1 for candidate beliefs $\hat{\mu}$, given that the underlying true type distribution is consistent with the unknown belief μ. Under this mismatch between the firm's menu (geared towards $\hat{\mu}$) and the actual type distribution (consistent with μ), the firm's (worst-case) expected profit becomes [13]

$$\hat{\Pi}(\hat{\mu}, \mu) = \min_{q_L^*(\hat{\mu}) \in \mathcal{Q}_L^*(\hat{\mu})} \{(1-\mu)\Pi_L(q_L^*(\hat{\mu})) + \mu \Pi_H(q_L^*(\hat{\mu}), q_H^*)\},$$
$$(\hat{\mu}, \mu) \in [0,1] \times [0,1], \tag{9}$$

where $\mathcal{Q}_L^*(\cdot)$ and q_H^* are specified in Eq. (7). The *performance ratio*,

$$\varphi(\hat{\mu}, \mu) = \frac{\hat{\Pi}(\hat{\mu}, \mu)}{\Pi^*(\mu)} \in [0,1], \tag{10}$$

provides a lower bound for the achievement of a product menu $\mathcal{M}(\hat{\mu}) = \{\boldsymbol{p}^*(\hat{\mu}), \boldsymbol{q}^*(\hat{\mu})\}$ for the candidate belief $\hat{\mu}$—relative to the "ex-post optimal profit,"

$$\Pi^*(\mu) = \hat{\Pi}(\mu, \mu) > 0, \tag{11}$$

which would have been obtained in the absence of any ambiguity about the type distribution. We note that the performance ratio is continuous in μ, while it is generically discontinuous in $\hat{\mu}$. The *performance index*,

$$\rho(\hat{\mu}) = \inf_{\mu \in [0,1]} \varphi(\hat{\mu}, \mu), \quad (\hat{\mu}, \mu) \in [0,1] \times [0,1], \tag{12}$$

corresponds to the worst-case performance guarantee. For example, a performance index $\rho(\hat{\mu}) = 75\%$ means that by selecting the product menu $\mathcal{M}(\hat{\mu})$ the firm's performance in terms of profitability will be at least three quarters of what it could have been with full information about the type distribution. The firm's *robust identification problem* is to determine a "robust belief" $\hat{\mu}^*$ with the best possible performance index:[14]

$$\hat{\mu}^* \in \arg \max_{\hat{\mu} \in [0,1]} \rho(\hat{\mu}). \tag{13}$$

The solution to this (generally nonconvex) optimization problem is discussed next.

Remark 4 (Relative Regret)**.** The performance index ρ in Eq. (12) is related to the notion of maximum relative regret,

$$r(\hat{\mu}) = \sup_{\mu \in [0,1]} \frac{\Pi^*(\mu) - \hat{\Pi}(\hat{\mu}, \mu)}{\Pi^*(\mu)} = 1 - \rho(\hat{\mu}).$$

Thus, maximizing the relative performance index ρ, as in Eq. (13), is equivalent to minimizing the maximum relative regret r, with $\sup \rho([0,1]) = 1 - \inf r([0,1])$.

3.2. Robust beliefs

To tackle the robust identification problem (13), we now provide a representation of the performance index as lower envelope of the "boundary performance ratios," $\varphi_0(\cdot) = \varphi(\cdot, 0)$ and $\varphi_1(\cdot) = \varphi(\cdot, 1)$.

Proposition 1 (Representation of Performance Index)**.** *The performance ratio and the performance index are lower semicontinuous in $\hat{\mu}$. Moreover, the firm's performance index in Eq. (12) is given by*

$$\rho(\hat{\mu}) = \min\{\varphi_0(\hat{\mu}), \varphi_1(\hat{\mu})\}, \tag{14}$$

for all candidate beliefs $\hat{\mu} \in [0,1]$.

The preceding result is driven by the fact that each performance ratio $\varphi(\hat{\mu}, \mu)$ is quasiconcave in μ. The lower semicontinuity of the performance index guarantees that the supremum can be attained by the limit of a sequence $(\hat{\mu}_k)_{k=1}^{\infty}$; cf. footnote 14. The representation of the performance index removes the need to determine the selection for all $\mu \in [0,1]$, and only two selections to minimize $\Pi_L^*(\hat{\mu}), \Pi_H^*(\hat{\mu})$, respectively, are required (cf. footnote 13). It also illustrates the balancedness of any robust belief with respect to extreme beliefs, and simplifies the procedure of finding the optimal robust belief in Eq. (13) substantially. Consider now the difference between the boundary performance ratios,

$$\Delta(\hat{\mu}) = \varphi_1(\hat{\mu}) - \varphi_0(\hat{\mu}), \quad \hat{\mu} \in [0,1].$$

An optimal robust action should be implemented for a belief at which Δ is about to change sign.

Proposition 2 (Characterization of Robust Beliefs)**.** *The firm's robust belief $\hat{\mu}^*$ is unique and such that*

$$\hat{\mu}^* = \sup\{\hat{\mu} : \Delta(\hat{\mu}) < 0\} \leq \mu_0, \tag{15}$$

while the optimal performance index $\rho(\hat{\mu}^*)$ is positive. Moreover, if the correspondence $\mathcal{Q}_L^*$ defined in Eq. (7) is single-valued for $\mu \leq \mu_0$, then the robust belief $\hat{\mu}^* \in (0, \mu_0)$ satisfies

$$\Delta(\hat{\mu}^*) = 0. \tag{16}$$

The characterization of the firm's robust belief in Eq. (15) turns out to be quite simple, especially when the standard screening problem in Lemma 1 produces a unique optimal menu. In this case, it is enough to balance the performance ratios at the boundary by

[13] The minimum is achieved, since $\mathcal{Q}_L^*$ is compact-valued by Corollary 1. An alternative equivalent way is to take the infimum in Eq. (12) jointly with respect to μ and the selection $q_L^*(\hat{\mu}) \in \mathcal{Q}_L^*(\hat{\mu})$, while dropping the minimization in Eq. (9).

[14] In view of the generic discontinuity of the performance index, a technically more correct definition of an optimal robust belief is for $\hat{\mu}^*$ to be in the (nonempty) set $\{\hat{\mu} \in [0,1] : \exists (\hat{\mu}_k)_{k=1}^{\infty} \subset [0,1]$ *such that* $\lim_{k\to\infty} \hat{\mu}_k = \hat{\mu}$ *and* $\lim_{k\to\infty} \rho(\hat{\mu}_k) = \sup \rho([0,1])\}$, which merely requires that there exists a converging sequence of candidate beliefs with performance indices converging towards the supremum of $\rho([0,1])$.

finding candidate beliefs $\hat{\mu}$ such that $\varphi_0(\hat{\mu}) = \varphi_1(\hat{\mu})$. It is interesting to note that $\hat{\mu}^* < \mu_0$ when the menu is unique, implying that it is best—from a robustness standpoint—to serve both consumer types. We also assert that a positive optimal performance index is achieved, which means that the firm is guaranteed a nontrivial relative performance.

3.3. Performance bounds

The robust identification problem (13) yields a robust belief $\hat{\mu}^*$, together with an optimal performance index,

$$\rho^* = \rho(\hat{\mu}^*) = \sup \rho([0,1]).$$

The latter implies a relative performance guarantee, which is now examined with respect to unfavorable scenarios in terms of consumer types $\boldsymbol{\theta} \in \hat{\Theta} = \{(\theta_L, \theta_H) \in \mathbb{R}^2_{++} : \theta_L < \theta_H\}$. The resulting worst-case performance index ρ^*_{WC} is significant in the sense that no matter what the demand environment might be, an optimal product portfolio (determined with respect to a robust belief) will be at least within the factor ρ^*_{WC} of any ex-post optimum. The corresponding result for our example (cf. Section 4) guarantees that no matter what demand environment the firm might encounter, by choosing its unique robust belief, the firm always achieves at least $\rho^*_{\mathrm{WC}} = 75\%$. That is, it attains at least three quarters of its ex-post optimal performance in the absence of ambiguity about the type distribution.

To simplify our presentation, we assume in this subsection that for $\mu \le \mu_0$, q^*_L—as solution to the quality-optimization problem for the low type in Eq. (7)—is unique, i.e., $\mathcal{Q}^*_L$ is single-valued. By Propositions 1 and 2, the optimal performance ρ^* is such that

$$\rho^* = \varphi_0(\hat{\mu}^*) = \varphi_1(\hat{\mu}^*), \tag{17}$$

where the second equality determines the firm's robust belief $\hat{\mu}^*$. Both ρ^* and $\hat{\mu}^*$ depend on the type vector $\boldsymbol{\theta} \in \hat{\Theta}$. Denote $\hat{q}^*_L = q^*_L(\hat{\mu}^*)$ and consider any $\boldsymbol{\theta}^m = (\theta^m_L, \theta^m_H)$ which attains the worst-case optimal performance index in the open domain $\hat{\Theta}$:

$$\rho^*_{\mathrm{WC}} = \inf_{\boldsymbol{\theta} \in \hat{\Theta}} \rho^*.$$

In order to use the first-order condition in Eq. (17), in this subsection we make the additional technical assumption that $\hat{q}^*_L$ is differentiable with respect to θ_L, θ_H at $\boldsymbol{\theta}^m \in \hat{\Theta}$, and q^*_H is uniquely determined by Eq. (7). By Cor. 4 in Milgrom & Segal (2002), all terms in Eq. (17) are differentiable with respect to θ_L, θ_H. Thus, the standard first-order condition $(\partial_{\theta_L}\rho^*, \partial_{\theta_H}\rho^*) = 0$ must hold, which leads to a characterization of the worst-case demand scenario.

Lemma 2 (Worst-Case Demand)**.** *If $\boldsymbol{\theta}^m \in \hat{\Theta}$ with $\rho^*|_{\boldsymbol{\theta}=\boldsymbol{\theta}^m} = \rho^*_{\mathrm{WC}}$, then* [15]

$$(a_1) \quad \frac{\partial_{\theta_L}\Pi^*_L(\hat{\mu}^*)}{\Pi^*_L(\hat{\mu}^*)} = \frac{\partial_{\theta_L}\Pi^*_L(0)}{\Pi^*_L(0)} \quad \text{and} \quad (a_2) \quad \partial_{\theta_H}\Pi^*_L(\hat{\mu}^*) = 0,$$

as well as

$$(b_1) \quad \partial_{\theta_L}\Pi^*_H(\hat{\mu}^*) = 0 \quad \text{and} \quad (b_2) \quad \frac{\partial_{\theta_H}\Pi^*_H(\hat{\mu}^*)}{\Pi^*_H(\hat{\mu}^*)} = \frac{\partial_{\theta_H}\Pi^*_H(1)}{\Pi^*_H(1)}.$$

Conditions (a) and (b) are equivalent, in the sense that $(a_1) \Leftrightarrow (b_1)$ and $(a_2) \Leftrightarrow (b_2)$.

[15] The equivalence of conditions (a) and (b) holds, provided that $\Delta(\hat{\mu}^*) = 0$; cf. Proposition 2.

Conditions (a_1) and (b_2) describe a situation where the type-elasticity of profits at the robust belief is as if that type occurs exclusively (i.e., when $\mu = 0$ for θ_L, and $\mu = 1$ for θ_H). Equivalently, conditions (b_1) and (a_2) state that the profits for a given type do not change (at the margin) for any small variation in the other type. The latter is quite natural—in view of the "type exclusivity" prescribed in (a_1) and (b_2). In practice, as illustrated by our example in Section 4, the low-type profit Π^*_L has a simpler form than the high-type profit Π^*_H, so that it may be advantageous to use conditions (a_1) and (a_2). Using Lemma 2, it is now possible to characterize the lowest possible performance bound.

Proposition 3 (Worst-Case Performance Index)**.** *If $\boldsymbol{\theta}^m = (\theta^m_L, \theta^m_H) \in \hat{\Theta}$ with $\rho^*|_{\boldsymbol{\theta}=\boldsymbol{\theta}^m} = \rho^*_{\mathrm{WC}}$, then*

$$\rho^*_{\mathrm{WC}} = \frac{1}{1-\hat{\mu}^*}\frac{u_\theta(\hat{q}^*_L, \theta^m_L)}{u_\theta(q^*_L(0), \theta^m_L)} = 1 - \frac{u_\theta(\hat{q}^*_L, \theta^m_H)}{u_\theta(q^*_H, \theta^m_H)} > 0, \tag{18}$$

where $\hat{q}^*_L$, q^*_H, and $\hat{\mu}^*$ are determined by Eqs. (7) and (16), respectively, for $\boldsymbol{\theta} = \boldsymbol{\theta}^m$.

Proposition 3 provides an alternative way of determining $\boldsymbol{\theta}^m$ (after having computed $\hat{\mu}^*$ by means of Eq. (16) in Proposition 2). Eq. (18) provides two alternative expressions for the worst-case performance index.[16] The first expression is the ratio of the type-gradient of the low type's utility for $\hat{q}^*_L$ conditional on serving the low type (with $1 - \hat{\mu}^* = \mathrm{Prob}(\tilde{\theta} = \theta_L)$ at the robust belief $\hat{\mu}^*$) and his type-gradient at $q^*_L(0)$ when only low-type consumers are present. The second expression corresponds to the relative increase of the type-gradient in the high type's utility when augmenting quality from $\hat{q}^*_L$ to q^*_H. One of the most remarkable aspects of Proposition 3 (which was already contained in Proposition 2) is that the worst-case performance index in Eq. (18) is positive, thus guaranteeing the firm a certain nontrivial performance relative to ex-post optimal profits, no matter what the demand scenario might be. As already mentioned, the worst-case performance guarantee may be quite substantial; see Section 4.

Remark 5 (Nondegeneracy)**.** If the type vector $\boldsymbol{\theta}$ lies at the boundary of the open set $\hat{\Theta}$ (i.e., if $\boldsymbol{\theta} \in \partial\hat{\Theta}$, so $\theta_L \in \{0, \theta_H\}$), then by Eq. (7) the performance ratio must be maximal: $\varphi(\hat{\mu}, \mu) = 1$ (except for $\theta_L = \mu = 0$, when $\varphi(\hat{\mu}, \mu)$ is not defined). Thus, $\rho^* = 1$ which further implies that the worst-case performance ratio is attained at an interior $\boldsymbol{\theta} \in \hat{\Theta}$, as long as $\theta_H < \infty$, so that Lemma 2 and Proposition 3 can be viewed as characterization results.[17]

3.4. Extension: general ambiguity

Expanding on the preceding developments, we now consider a situation where the firm faces only limited ambiguity, based on a nuanced prior knowledge about the type distribution. Specifically, we assume that all beliefs the firm considers possible are elements of the (nonempty) compact *ambiguity set* $\mathcal{A} \subset [0,1]$. Using the same arguments and proof techniques as before, the firm can restrict attention to the performance ratios at the boundary of the convex hull of the ambiguity set.

[16] A straightforward interpretation of Eq. (18) is that the worst-case performance index ρ^*_{WC} corresponds to the ratio between the gradient of the type-contingent robust profit and the gradient of the type-exclusive profit, as spelled out by Eq. (31) in the proof of Proposition 3: $\rho^*_{\mathrm{WC}} = [\partial_{\theta_L}\Pi^*_L(\hat{\mu}^*)]/[\partial_{\theta_L}\Pi^*_L(0)] = [\partial_{\theta_H}\Pi^*_H(\hat{\mu}^*)]/[\partial_{\theta_H}\Pi^*_H(1)]$.

[17] The characterization is notwithstanding a potential multiplicity of local extrema in the performance index, due to the lack of higher-order curvature assumptions on the model primitives.

Proposition 4 (Robust Beliefs under General Ambiguity)**.** *Let $\mu_1 = \min \mathcal{A}$ and $\mu_2 = \max \mathcal{A}$ denote the largest lower bound and smallest upper bound of the firm's ambiguity set $\mathcal{A}$, respectively, and let $\rho(\cdot|\mathcal{A}) = \inf \varphi(\cdot, \mathcal{A})$ be the corresponding ("$\mathcal{A}$-conditional") performance index. Then:*

(i) *The $\mathcal{A}$-conditional performance index $\rho(\cdot|\mathcal{A})$ is lower semicontinuous, and*

$$\rho(\hat{\mu}|\mathcal{A}) = \min\{\varphi(\hat{\mu}, \mu_1), \varphi(\hat{\mu}, \mu_2)\}, \quad \hat{\mu} \in [0, 1]. \tag{19}$$

(ii) *If $\mathcal{Q}_L^*(\cdot)$ is single-valued, then any solution $\hat{\mu}^*$ to the robust identification problem (13) is such that*

$$\varphi(\hat{\mu}^*, \mu_1) = \varphi(\hat{\mu}^*, \mu_2). \tag{20}$$

Parts (i) and (ii) of Proposition 4 generalize the earlier representation results in Propositions 1 and 2. It is important to note that Eq. (20) will generically yield a robust belief in the convex hull of the ambiguity set, and thus potentially a belief that the firm may have considered impossible. Yet, by pragmatically designing its product portfolio in accordance with a robust belief $\hat{\mu}^*$, the firm achieves the best possible performance guarantee.

Remark 6 (Wasserstein Distance)**.** For any constant $p \geq 1$ and any type vector $(\theta_L, \theta_H) \in \hat{\Theta}$, the p-Wasserstein distance between the distributions $\hat{f}(\theta) = (1 - \hat{\mu})\delta(\theta - \theta_L) + \hat{\mu}\delta(\theta - \theta_H)$ and $f(\theta) = (1 - \mu)\delta(\theta - \theta_L) + \mu\delta(\theta - \theta_H)$, defined for $\theta > 0$ and the given beliefs $\hat{\mu}, \mu \in \mathcal{A}$, with the Dirac distribution $\delta(\cdot)$ and the standard absolute-value metric $|\cdot|$ on the real line, is

$$W_p(\hat{f}, f) = |\hat{\mu} - \mu|^{1/p}(\theta_H - \theta_L).$$

Thus, for $p = 1$ that distance becomes directly proportional to $|\hat{\mu} - \mu|$, implying that the Wasserstein diameter of $\mathcal{A}$ is $(\mu_2 - \mu_1)(\theta_H - \theta_L)$.

Remark 7 (Homotopic Connection)**.** As the diameter of the ambiguity set tends to zero, that is, as $\mu_2 \to \mu_1$, the firm's pricing problem collapses to the standard screening problem discussed in Section 2; see Lemma 1.[18]

4. Example: a standard specification

To illustrate our results, we now discuss the LQ model as a classical instance of the price discrimination-problem, with the linear utility and quadratic cost as introduced in Remark 2. Models of this type are in widespread use to this day (see, e.g., Wong et al., 2021; Zou et al., 2020). In what follows, we first provide an explicit solution to the firm's price-discrimination problem with robust beliefs, and then compare the performance of this solution to several common belief heuristics.

4.1. Closed-form solution

Using Eqs. (6) and (7) in Lemma 1, it is straightforward to determine the (in this case unique) solution $\mathcal{M}(\mu) = \{(\boldsymbol{p}^*(\mu), \boldsymbol{q}^*(\mu))\}$ to the firm's standard screening problem (5) for any given belief $\mu \in [0, 1]$:

$$\boldsymbol{p}^*(\mu) = (p_L^*(\mu), p_H^*(\mu)) = \big(\theta_L q_L^*(\mu), \theta_H q_H^* - (\theta_H - \theta_L) q_L^*(\mu)\big),$$

and [19]

$$\boldsymbol{q}^*(\mu) = (q_L^*(\mu), q_H^*) = \left(\left[\frac{\theta_L}{\gamma} - \frac{\mu}{1-\mu}\,\frac{\theta_H - \theta_L}{\gamma}\right]_+, \frac{\theta_H}{\gamma}\right).$$

The nature of the solution depends on the magnitude of the beliefs about the prevalence of high types in the consumer population. For $\mu \in [0, \mu_0)$, where the belief threshold $\mu_0 = \theta_L/\theta_H \in (0, 1)$ is obtained from Eq. (8), the quality $q_L^*(\mu)$ provided to the low type is positive. For $\mu \geq \mu_0$, the firm prefers a "shutdown solution" by not serving the low type and providing zero quality (at zero cost) while offering a socially efficient ("first-best") quality level q_H^* to the high type. This first-best, undistorted quality level q_H^* served to the high type is constant in μ. However, the information rent (i.e., the price discount over full revenue extraction),[20]

$$\mathrm{IR}(q_L^*(\mu)) = (\theta_H - \theta_L) q_L^*(\mu),$$

is only positive when the low type is not shut down, i.e., for $\mu \in [0, \mu_0)$. On the other hand, the firm's price charged for the low-type product extracts all surplus (and is "first-best" in that sense). The quality delivered to the low type,

$$q_L^*(\mu) = \frac{\theta_L}{\gamma}\left[1 - \frac{\mu}{\mu_0}\,\frac{(1-\mu_0)}{(1-\mu)}\right]_+, \tag{21}$$

is distorted downwards from the first-best level θ_L/γ. We note that the optimal price-quantity tuple $(\boldsymbol{p}^*(\mu), \boldsymbol{q}^*(\mu))$ is continuous in μ.

We now turn our attention to the firm's fundamentally ambiguous beliefs, where the type distribution is in fact unknown. For any candidate belief $\hat{\mu}$ the profit $\hat{\Pi}(\hat{\mu}, \mu)$ as a function of the true belief μ is given in Eq. (9), with type-contingent profits

$$\Pi_L^*(\hat{\mu}) = \frac{\theta_L^2}{2\gamma}\left[1 - \left(\frac{\hat{\mu}}{\mu_0}\right)^2 \left(\frac{1-\mu_0}{1-\hat{\mu}}\right)^2\right]_+$$

and

$$\Pi_H^*(\hat{\mu}) = \frac{\theta_H^2}{2\gamma} - (\theta_H - \theta_L) q_L^*(\hat{\mu}),$$

with the underlying assumption that the firm's optimal menu $(\boldsymbol{p}^*(\hat{\mu}), \boldsymbol{q}^*(\hat{\mu}))$ is designed according to the candidate belief $\hat{\mu}$. This implies the performance ratio $\varphi(\hat{\mu}, \mu) = \hat{\Pi}(\hat{\mu}, \mu)/\Pi^*(\mu)$ in Eq. (10) relative to the firm's ex-post optimal profits,

$$\Pi^*(\mu) = (1 - \mu)\Pi_L^*(\mu) + \mu \Pi_H^*(\mu), \quad \mu \in [0, 1],$$

as in Eq. (11). It is clear that because of the multiplicative dependence of both type-contingent profits on the cost parameter, the performance ratio φ (and *a fortiori* also the performance index ρ) must be independent of γ. Given the performance ratios,

$$\varphi_0(\hat{\mu}) = \left[1 - \left(\frac{\hat{\mu}}{\mu_0}\right)^2 \left(\frac{1-\mu_0}{1-\hat{\mu}}\right)^2\right]_+, \qquad \varphi_1(\hat{\mu}) = 1 - 2(1-\mu_0)\left[\frac{\mu_0 - \hat{\mu}}{1-\hat{\mu}}\right]_+, \tag{22}$$

at the boundary of the firm's ambiguity set $\mathcal{A} = [0, 1]$, Proposition 1 now allows for a compact representation of the performance index: $\rho(\hat{\mu}) = \min\{\varphi_0(\hat{\mu}), \varphi_1(\hat{\mu})\}$, for all $\hat{\mu} \in [0, 1]$. By Proposition 2, the firm's robust belief $\hat{\mu}^*$ is such that the difference of the boundary performance ratios must vanish, as shown in Fig. 1. The following result provides a closed-form solution for $\hat{\mu}^*(\mu_0)$ (depicted in Fig. 2), together with its first-order and second-order monotonicity properties.

[18] This homotopic connection to the original problem is much in the spirit of Zangwill & Garcia (1981).

[19] For any real number $x \in \mathbb{R}$, we set $[x]_+ = \max\{0, x\}$ to denote its nonnegative part.

[20] The high type's "information rent" can be inferred from Eq. (6): it is the difference between the high type's willingness-to-pay and the price charged by the firm: $\mathrm{IR}(\hat{q}_L^*) = u(\hat{q}_L^*, \theta_H) - u(\hat{q}_L^*, \theta_L) = u(q_H^*, \theta_H) - p_H^*$.

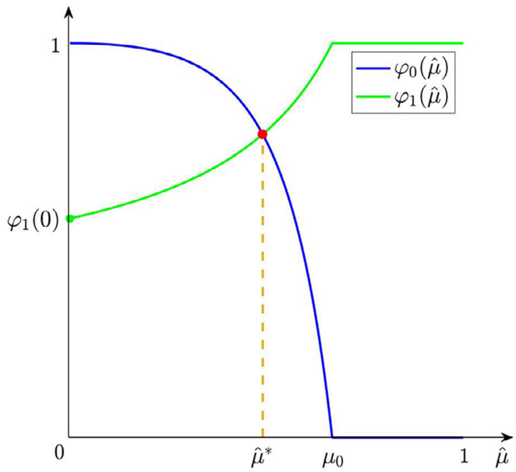

Fig. 1. Performance ratios $\varphi_0(\hat{\mu})$, $\varphi_1(\hat{\mu})$ and performance index $\rho(\hat{\mu})$, for $\hat{\mu} \in [0, 1]$.

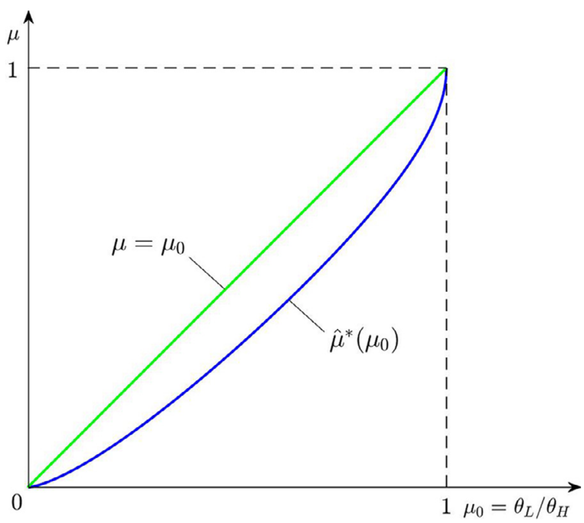

Fig. 2. Robust beliefs $\hat{\mu}^*(\mu_0)$, for $\mu_0 \in (0, 1)$.

Lemma 3 (Example: Robust Beliefs)**.** *For any type vector* $(\theta_L, \theta_H) \in \hat{\Theta}$*, the solution* $\hat{\mu}^* = \hat{\mu}^*(\mu_0)$ *to the robust identification problem (13) (for* $\mu_0 = \theta_L/\theta_H \in (0, 1)$*) is unique and given by*

$$\hat{\mu}^*(\mu_0) = \mu_0\left(1 - \frac{\sqrt{\mu_0^2(1-\mu_0)^2 + 2\mu_0(1-\mu_0)} - \mu_0(1-\mu_0)}{\sqrt{\mu_0^2(1-\mu_0)^2 + 2\mu_0(1-\mu_0)} + \mu_0(1+\mu_0)}\right). \quad (23)$$

The two endpoints of the type-threshold domain are fixed points in the limit, in the sense that $\hat{\mu}^*(0^+) = 0$ *and* $\hat{\mu}^*(1^-) = 1$*. Finally, the solution is continuously differentiable, increasing and convex, with* $\hat{\mu}^{*\prime}(0^+) = 0$ *and* $\hat{\mu}^{*\prime}(1^-) = \infty$*.*

As already formalized in Proposition 2, Eq. (23) reflects the interiority of the optimal robust belief in the sense that $\hat{\mu}^*(\mu_0) \in (0, \mu_0)$. Thus, each consumer type is offered a product of nonzero quality—no matter what the nature of demand might be (i.e., independent of the type vector $\boldsymbol{\theta} \in \hat{\Theta}$). Additionally, the firm's robust belief is independent of its cost parameter, and so is the optimal performance index,

$$\rho^*(\mu_0) = 1 - 2\mu_0(1-\mu_0) \cdot \cdot \left(1 + \mu_0(1-\mu_0) - \sqrt{\left(1 + \mu_0(1-\mu_0)\right)^2 - 1}\right),$$

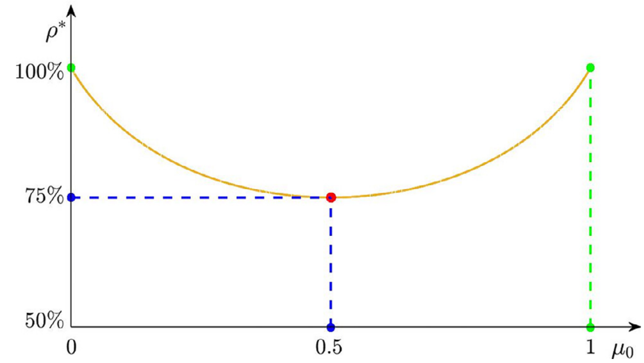

Fig. 3. Optimal performance index $\rho^*(\mu_0)$, for $\mu_0 \in (0, 1)$.

for all $\mu_0 \in (0, 1)$; see Fig. 3. It becomes readily apparent that the latter is fully symmetric in the type threshold, i.e.,

$$\rho^*(\mu_0) = \rho^*(1-\mu_0).$$

We now turn our attention to determining a tight lower bound for the performance index, as discussed in Section 3.3. Indeed, Eq. (18) in Proposition 3 yields [21]

$$\rho^*_{WC} = \frac{1 - (\hat{\mu}^*/\mu_0^m)}{(1-\hat{\mu}^*)^2} = \frac{1-\mu_0^m}{1-\hat{\mu}^*},$$

where $\mu_0^m = \mu_0|_{\boldsymbol{\theta}=\boldsymbol{\theta}^m}$ and $\boldsymbol{\theta}^m = (\theta_L^m, \theta_H^m) \in \arg\min_{\boldsymbol{\theta}\in\hat{\Theta}} \rho^*$ is such that

$$\frac{\theta_L^m}{\theta_H^m} = \frac{1}{2} = \mu_0^m.$$

This results in the worst-case performance index:[22]

$$\rho^*_{WC} = 75\%.$$

Hence, no matter what the demand characteristics may be (in terms of the type vector $\boldsymbol{\theta} \in \hat{\Theta}$), when choosing the *robust product menu*,

$$(\hat{\boldsymbol{p}}^*, \hat{\boldsymbol{q}}^*) = (\boldsymbol{p}^*(\hat{\mu}^*), \boldsymbol{q}^*(\hat{\mu}^*)),$$

the firm's profits are guaranteed to remain within a 25%-band of the ex-post optimal profits that could be achieved if the type distribution was perfectly known, irrespective of $(\mu, \boldsymbol{\theta}) \in [0, 1] \times \hat{\Theta}$.

4.2. Performance comparison

We now benchmark our robust solution against four alternative belief heuristics,

$$\hat{\mu} \in \{0, 1/2, \mu_0/2, \mu_0\}.$$

The *worst-case heuristic*, $\hat{\mu} = 0$, is the most conservative approach where the firm assumes that there are no high types at all. It implies a minimax approach and therefore a "worst-case product portfolio." Following the "principle of insufficient reason" Laplace (1825) proposes to assign uniform beliefs over outcomes in the absence of better knowledge. This yields the *Laplacian* (*equiprobable*) *belief heuristic*—with $\hat{\mu} = 1/2$. Similarly, when the type threshold μ_0 for the shutdown solution is known, then the belief $\hat{\mu} = \mu_0$

[21] Alternatively, one can either directly minimize the closed-form expression $\rho^*(\mu_0)$ or else use Lemma 2; the latter yields: $\hat{\mu}^{*\prime} = [\hat{\mu}^*(1-\hat{\mu}^*)]/[\mu_0(1-\mu_0)]$ (by condition (a_1) or (a_2)) and $\hat{\mu}^{*\prime} = [(1-\hat{\mu}^*)(\hat{\mu}^* - 2\mu_0 + 1)]/[(1-\mu_0)^2]$ (by condition (b_1) or (b_2)), and thus again the same solution: $\mu_0^m = 1/2$.

[22] The robust belief compatible with the worst-case performance index is $\hat{\mu}^*|_{\boldsymbol{\theta}=\boldsymbol{\theta}^m} = 1/3$.

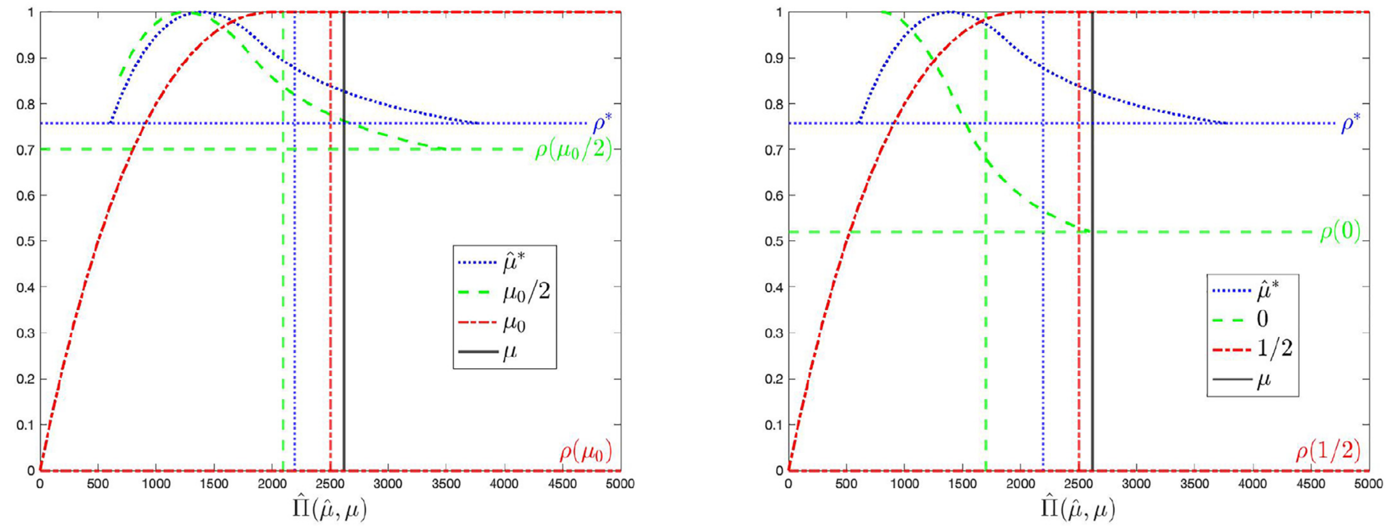

(a) $\hat{\mu} \in \{\mu_0/2, \mu_0, \hat{\mu}^*\}$. (b) $\hat{\mu} \in \{0, 1/2, \hat{\mu}^*\}$.

Fig. 4. Profit realizations and performance index for different belief heuristics.

Table 1
Profit and performance index for different belief heuristics.

	Belief heuristic ($\hat{\mu}$)					
	0	$\mu_0/2$	μ_0	1/2	$\hat{\mu}^*$	μ
Average profit	1700	2093.8	2500	2500	2194.5	2619.8
Average relative profit [%]	72.75	84.81	88.03	88.03	87.03	100
Performance index [%]	52.00	70.00	0	0	75.67	100

can be termed the *most-conservative shutdown heuristic*.[23] Finally, a candidate belief $\hat{\mu} = \mu_0/2$ corresponds to a *Laplacian full-coverage heuristic*, which is consistent with the firm's decision to serve both consumers (the latter being in turn consistent with $\mu < \mu_0$ without any additional distributional assumptions). The ex-post optimal profit $\Pi^*(\mu)$ with perfect information about the type distribution μ serves as a baseline for our profit comparisons.

When the type vector is known (in our simulation it is fixed at the nominal value $\boldsymbol{\theta}^0 = (\theta_L^0, \theta_H^0) = (40, 100)$), then the robust product portfolio depends only on the type threshold μ_0. For $\gamma = 1$ and the true belief $\mu \in [0, 1]$, we examine the performance ratio implied by the firm's expected profit $\hat{\Pi}(\hat{\mu}, \mu)$.

Fig. 4 shows the (otherwise deterministic) performance comparison against the four belief heuristics introduced earlier. The vertical lines show the average profit for each heuristic, with the ex-post optimal profits (where $\hat{\mu} = \mu$) in black. We first note that the least-conservative shutdown heuristic ($\hat{\mu} = \mu_0$) and the Laplacian heuristic ($\hat{\mu} = 1/2$), corresponding to the red curves in Fig. 4(a) and (b), perform well on average because they correctly limit the product portfolio for bullish type distributions. On the flipside, however, the performance ratio can drop arbitrarily low as soon as high types become rare in the economy (so $\rho(\mu_0) = \rho(1/2) = 0$). The worst-case belief heuristic ($\hat{\mu} = 0$) and the Laplacian full-coverage heuristic tend to do well for type distributions that favor the occurrence of the low type (i.e., when μ is fairly small). Finally, the robust product portfolio, with $\hat{\boldsymbol{p}}^* = (1984 - 256\sqrt{21}, 7024 + 384\sqrt{21}) \approx (816.86, 8783.71)$ and $\hat{\boldsymbol{q}}^* = ((248 - 32\sqrt{21})/5, 100) \approx (20.27, 100)$, shown by the blue curves, exhibits a "balanced performance," in the sense that the boundary performance ratios are the same (at $\rho^*_{\mathrm{WC}} = 75\%$). Table 1 compares the numerical performance, including the performance index (which is indicated by horizontal lines in Fig. 4). The average relative profit is the average over performance ratios. While the robust performance guarantee makes sure that the firm can never do worse than 75% of the ex-post optimum, the average robust outcome exceeds 85% of the ex-post optimal profit.

To check the worst-case performance of the proposed robust belief and the various alternative belief heuristics over possible type realizations $\boldsymbol{\theta} \in \hat{\Theta}$, we generate a random sample $\hat{\boldsymbol{\theta}} = (\boldsymbol{\theta}^k)_{k=1}^N$ of size $N = 10{,}000$, drawn from a uniform distribution on the rectangle

$$[\theta_L^0 - \varepsilon, \theta_L^0 + \varepsilon] \times [\theta_H^0 - \delta, \theta_H^0 + \delta] \subset \hat{\Theta},$$

for the nominal type vector $\boldsymbol{\theta}^0 = (\theta_L^0, \theta_H^0) = (40, 100)$ and possible type realizations characterized by the dispersion vector $(\varepsilon, \delta) = (10, 20)$.

In addition, we draw $M = 1{,}000$ random samples of μ from a uniform distribution on [0,1]. Fig. 5 shows the performance ratios on the vertical axis, and the corresponding profit $\hat{\Pi}(\hat{\mu}, \mu|\boldsymbol{\theta})$ on the horizontal axis. Vertical lines indicate average profits, while horizontal lines mark the worst-case performance index over the random sample $\hat{\boldsymbol{\theta}}$. The shaded areas cover 90% of the data points for each belief heuristic $\hat{\mu}$, around the respective average relative-performance curves. We note an interesting nonmonotonicity of these curves at the right end for $\hat{\mu} \in \{0, \mu_0/2, \hat{\mu}^*\}$, which results from the fact that large profits are realized at shutdown solutions for relatively large θ_H and small θ_L, leading to full revenue extraction and thus an improving relative performance. Table 2 juxtaposes the average profit (absolute and relative) and the worst-case performance index for the different belief heuristics $\hat{\mu}$, including the proposed optimal robust belief $\hat{\mu}^*$.

In terms of relative performance, the proposed optimal robust belief (as solution to the robust identification problem (13)) features a performance guarantee of 75% and yet attains over 87%

[23] We disregard the *least-conservative shutdown heuristic* $\hat{\mu} = 1$, as it would effectively assume that all consumers are of high type, which runs counter to any idea of performance robustness.

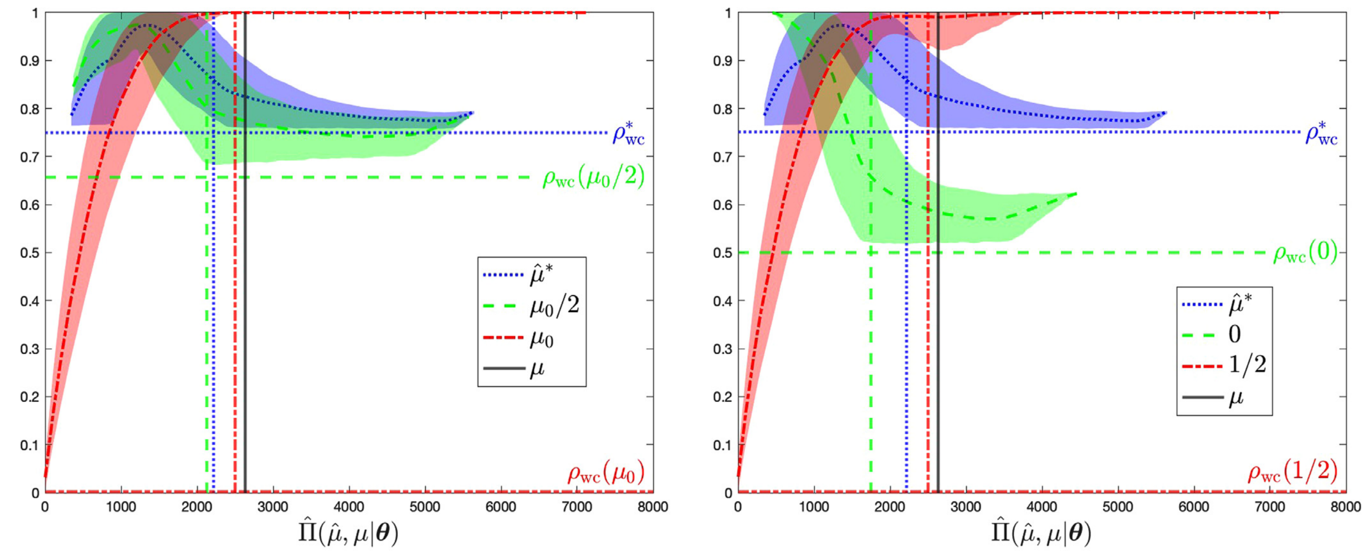

(a) $\hat{\mu} \in \{\mu_0/2, \mu_0, \hat{\mu}^*\}$.

(b) $\hat{\mu} \in \{0, 1/2, \hat{\mu}^*\}$.

Fig. 5. Performance distributions for different belief heuristics [Sample size $M = 1{,}000$ (in μ) and $N = 10{,}000$ (in $\boldsymbol{\theta}$)].

Table 2
Performance comparison across belief heuristics in the presence of type variation.

	Belief heuristic ($\hat{\mu}$)					
	0	$\mu_0/2$	μ_0	1/2	$\hat{\mu}^*$	μ
Average profit	1745.5	2126.6	2497.6	2500.1	2215.2	2635.2
Average relative profit [%]	74.00	85.37	86.84	87.34	87.21	100
Worst-case performance index [%]	50.06	65.76	0.16	0.25	75.04	100

of the average ex-post optimal profit. The reason that the worst-case performance guarantee may seem quite conservative by comparison is that it holds with respect to all possible type distributions, in particular with respect to a distribution that has all mass on types $\boldsymbol{\theta}^m = (\theta_L^m, \theta_H^m)$ with $\theta_L^m/\theta_H^m = 1/2 = \mu_0^m$ as introduced in Section 4.1. In such cases, the robust-belief estimator $\hat{\mu}^*$ would also outperform belief heuristics that under uniform sampling may provide higher average profits, but without any attractive relative (or absolute) performance guarantees.

5. Conclusion

In order to provide a belief-robust solution to the second-degree price-discrimination problem, we use a performance index—computed as the worst-case performance ratio relative to an ex-post optimal solution with perfectly known type distribution—which is consistent with the notion of relative regret. The firm's profit using an assumed synthetic belief is thereby evaluated as a fraction of the attainable optimal expected profits that would have been obtained, had that belief been correct in view of the actual distribution of consumer types. The performance index, as a function of the firm's candidate robust belief, can be represented as the lower envelope of two monotonic boundary performance ratios, allowing for a simple computation of the firm's optimal robust belief. Our characterization of robust beliefs is fairly simple, despite the generically nonconvex nature of underlying optimization problem. An additional worst-case performance benchmark with respect to all possible type values provides an ex-ante guarantee for the firm's relative profit attainment using the proposed robust product portfolio over the different demand scenarios. Interestingly it is possible to obtain expressions for the worst-case performance ratio directly from the necessary optimality conditions related to the minimization of the standard performance ratio. And for the LQ model, as an important parametrization, this leads to a performance guarantee of 75% for the optimal robust belief with respect to all possible type values. More generally, our method implies that the robust solution is continuously connected to the standard screening solution without model uncertainty. Fully integrating the type values ex ante into the robustness framework adds substantial combinatorial complications, thus left as a challenge for future work. A generalization of our model to an arbitrary number of types and instruments presents another interesting avenue for further research.

Appendix A. Proofs

Proof of Lemma 1. Recall that $\theta_H > \theta_L > 0$. We start by showing that the first constraint in Eq. (3) and the second constraint in Eq. (4) are redundant, while the two other constraints are binding, which yields

$$u(q_H, \theta_H) - p_H = u(q_L, \theta_H) - p_L, \quad \text{and} \quad u(q_L, \theta_L) - p_L = 0. \tag{24}$$

Indeed, if $u(q_L, \theta_L) - p_L > 0$, then—using the second inequality in Eq. (3)—it is

$$u(q_H, \theta_H) - p_H \geq u(q_L, \theta_H) - p_L \geq u(q_L, \theta_L) - p_L > 0,$$

since by hypothesis $u_\theta \geq 0$. This indicates that p_L and p_H can be increased (by equal amounts) without affecting the constraints; yet, this cannot be true at any profit-maximizing solution. Hence, the first constraint in Eq. (4) must be binding. Now, if the second constraint in Eq. (3) is slack, then

$$u(q_H, \theta_H) - p_H > u(q_L, \theta_H) - p_L \geq u(q_L, \theta_L) - p_L = 0, \tag{25}$$

by virtue of the fact that $u_\theta \geq 0$. But this means p_H can be increased without violating any constraints, which in turn contradicts profit maximization. Thus, the second constraint in Eq. (3) must also be binding, which together with Eq. (25) implies that

$$p_H - p_L = u(q_H, \theta_H) - u(q_L, \theta_H) \geq u(q_H, \theta_L) - u(q_L, \theta_L),$$

where the last inequality follows from the hypothesis that $u_{q\theta} \geq 0$. As an immediate consequence of the preceding inequality the first constraint in Eq. (3) must hold. The second constraint in Eq. (4) is redundant, since (by taking into account the two binding constraints and $u_\theta \geq 0$):

$$u(q_H, \theta_H) - p_H = u(q_L, \theta_H) - p_L \geq u(q_L, \theta_L) - p_L = 0.$$

Hence, the binding constraints in Eq. (24) are an equivalent replacement of Eqs. (3) and (4); they imply that for any $\boldsymbol{q} = (q_L, q_H) \in \mathcal{Q}$:

$$\boldsymbol{p} = (p_L, p_H) = \Big(u(q_L, \theta_L), u(q_H, \theta_H) - \big(u(q_L, \theta_H) - u(q_L, \theta_L)\big)\Big), \tag{26}$$

and thus also Eq. (6) for $\boldsymbol{q} = \boldsymbol{q}^*(\mu)$. By substituting Eq. (26) into the firm's standard screening problem (5) the objective function becomes independent of $\boldsymbol{p}$ and additively separable in q_L and q_H, implying the two distinct optimization problems in Eq. (7). Thus, the solution to the screening problem (5) is given by $(\boldsymbol{p}^*(\mu), \boldsymbol{q}^*(\mu))$, for all $\mu \in [0, 1]$. □

Proof of Corollary 1.

(i) Fix $\boldsymbol{\theta} = (\theta_L, \theta_H) \in \hat{\Theta}$. Consider the first optimization problem in Eq. (7), with the objective

$$F(q, \mu) = (1 - \mu)\big(u(q, \theta_L) - C(q)\big) - \mu\big(u(q, \theta_H) - u(q, \theta_L)\big),$$

where $(q, \mu) \in \mathcal{Q} \times [0, 1]$.

By Eq. (2) and the assumption that $u_\theta > 0$, it is

$$q > \bar{q} \quad \Rightarrow \quad u(q, \theta_L) - C(q) < u(q, \theta_H) - C(q) \leq 0,$$

which yields, for all $\mu \in [0, 1]$ and $q > \bar{q}$:

$$F(q, \mu) = u(q, \theta_L) - C(q) - \mu\big(u(q, \theta_H) - C(q)\big) < 0 = F(0, \mu),$$

Hence, the correspondence $\mathcal{Q}_L^*$ (which, by the maximum theorem, has compact values) is such that $\mathcal{Q}_L^*(\mu) \subset [0, \bar{q}]$, for all $\mu \in [0, 1]$.

Next we prove that if shutting down the low-type consumer is optimal for one belief $\bar{\mu} \in [0, 1]$, then it is also optimal to shut down the low-type consumer for all higher beliefs. That is,

$$0 \in \mathcal{Q}_L^*(\bar{\mu}) \quad \Rightarrow \quad 0 \in \mathcal{Q}_L^*(\mu), \;\; \forall \mu \in [\bar{\mu}, 1]. \tag{27}$$

For this, note first the equivalence

$$0 \in \mathcal{Q}_L^*(\mu) \quad \Leftrightarrow \quad F(q_L, \mu) \leq F(0, \mu) = 0, \;\; \forall q_L \in \mathcal{Q}, \tag{28}$$

for any $\mu \in [0, 1]$. Therefore, $0 \in \mathcal{Q}_L^*(\bar{\mu})$, together with $u(q, \cdot)$ being increasing, implies that

$$\begin{aligned} u(q_L, \theta_L) - C(q_L) &\leq \frac{\bar{\mu}}{1 - \bar{\mu}}\big(u(q_L, \theta_H) - u(q_L, \theta_L)\big) \\ &\leq \frac{\mu}{1 - \mu}\big(u(q_L, \theta_H) - u(q_L, \theta_L)\big), \end{aligned}$$

for all $\mu \in [\bar{\mu}, 1)$. Additionally, it is $F(q_L, 1) \leq 0$. Thus, one obtains

$$0 \in \mathcal{Q}_L^*(\bar{\mu}) \quad \Rightarrow \quad F(q_L, \mu) \leq 0, \;\; \forall (q_L, \mu) \in \mathcal{Q} \times [\bar{\mu}, 1],$$

which, by virtue of (28), establishes the implication (27).

Since by Eq. (1) we have $u_{q\theta}(q, \theta) > 0$ whenever $\theta q \neq 0$, it is $F(q, 1) < 0$, for all $q > 0$. Thus,

$$\mathcal{Q}_L^*(1) = \{0\}. \tag{29}$$

This implies that $\mu_0 \in [0, 1]$ in Eq. (8) is well-defined. Moreover, by the upper hemicontinuity of $\mathcal{Q}_L^*$, necessarily

$$0 \in \mathcal{Q}_L^*(\mu_0). \tag{30}$$

It follows from Eq. (2) that $0 \notin \mathcal{Q}_L^*(0)$, so Eq. (30) implies that necessarily $\mu_0 > 0$.

Now we establish the monotonicity of Π_L^* and the monotonicity of any selection $q_L^*(\cdot) \in \mathcal{Q}_L^*(\cdot)$. For this, we focus on the nontrivial case where $\mu \in [0, \mu_0)$, so $q_L^*(\mu) > 0$. Differentiating the objective with respect to q and μ yields

$$\partial^2_{q\mu} F(q, \mu) = -\big(u_q(q, \theta_H) - C'(q)\big), \quad (q, \mu) \in \mathcal{Q} \times [0, 1].$$

Recall that q_L^* and q_H^* are both finite, due to Eq. (2). Thus, by the first-order necessary optimality condition for $q = q_L^*(\mu) \in \mathcal{Q}_L^*(\mu)$ it is $\partial_q F(q_L^*(\mu), \mu) = 0$, so

$$\begin{aligned} &u_q(q_L^*(\mu), \theta_L) - C'(q_L^*(\mu)) \\ &= \frac{\mu}{1 - \mu}\big(u_q(q_L^*(\mu), \theta_H) - u_q(q_L^*(\mu), \theta_L)\big) > 0, \end{aligned}$$

where the strict inequality follows from the fact that, by Eq. (1), $u_{q\theta}(q, \theta) > 0$ for all (q, θ) with $\theta q > 0$. Since $u_q(q, \theta) - C'(q)$ is continuous, there exists a neighborhood of $(q_L^*(\mu), \mu)$, in which (q_L, μ) still satisfies $u_q(q_L, \theta_L) - C'(q_L) > 0$ and where $\Pi_L(q_L) = u(q_L, \theta_L) - C(q_L)$ is increasing in q_L. Moreover, in a neighborhood of $(q_L^*(\mu), \mu)$, the objective function features strictly decreasing differences in (q_L, μ):

$$\partial^2_{q\mu} F(q, \mu) = -\big(u_q(q, \theta_H) - C'(q)\big) < 0.$$

By the regularity condition introduced in Section 2.2, $F(q, \mu)$ has strictly decreasing differences in $(q, \mu) \in \mathcal{I} \times [0, \mu_0)$, and for any $\mu < \mu_0$, $\mathcal{Q}_L^*(\mu)$ is in the interval $\mathcal{I}$. We conclude from Edlin & Shannon (1998, Cor. 1) that any selection $q_L^*(\mu)$ is strictly decreasing in μ for $\mu < \mu_0$ (cf. footnote 12).

From the monotonicity of $\Pi_L(q_L)$ in q_L it follows that $\Pi_L^*(\mu)$ is strictly decreasing in μ on any connected neighborhood. For $\mu < \mu_0$, if there exist $q_1, q_2 \in \mathcal{Q}_L^*(\mu)$ satisfying $q_1 < q_2$, by the assumption of $u_{q\theta} > 0$, the sorting condition holds:

$$u(q_1, \theta_H) - u(q_1, \theta_L) < u(q_2, \theta_H) - u(q_2, \theta_L),$$

which implies that $u(q_1, \theta_L) - C(q_1) < u(q_2, \theta_L) - C(q_2)$, since $F(q_1, \mu) = F(q_2, \mu)$. It further guarantees $\Pi_L(q_1) < \Pi_L(q_2)$, so that $\Pi_L^*(\mu)$ is also strictly decreasing. For $\mu \geq \mu_0$, the selection $q_L^* = 0$ yields $\Pi_L^* = 0$.

The second optimization problem in Eq. (7) yields a solution q_H^* which is constant in $\mu \in [0, 1]$. Hence, for $\mu \geq \mu_0$ and the selection $q_L^*(\mu) = 0$, the type-contingent profit $\Pi_H^*(\mu)$ is constant. Moreover, by the representation of $\Pi_H(\boldsymbol{q})$ and the assumption $u_{q\theta} > 0$, $\Pi_H(\boldsymbol{q})$ is strictly decreasing in $q_L > 0$. For $\mu < \mu_0$, since q_L^* is strictly decreasing, $\Pi_H^*(\mu)$ increases strictly in μ.

To show that Π_L^* is nonnegative, note first that by Eq. (29) it is $\mathcal{Q}_L^*(1) = \{0\}$, so $\Pi_L^*(1) = 0$. Let $\mu < 1$. Since $u_{q\theta} > 0$ and $F(q_L^*, \mu) \geq 0$, we have $\Pi_L^* = u(q_L^*, \theta_L) - C(q_L^*) \geq 0$, for any selection q_L^*. Moreover, by the definition of q_H^* in Eq. (7) it is

$$\Pi_H^* = \max_{q_H \in \mathcal{Q}} \big\{u(q_H, \theta_H) - C(q_H)\big\} \geq u(q_L^*, \theta_H) - C(q_L^*) \geq \Pi_L^* \geq 0;$$

that is, Π_H^* is also nonnegative and not smaller than Π_L^*.

(ii) Now we prove that for $\mu \in [0, \mu_0]$, $\mathcal{Q}_L^*(\mu)$ is single-valued except for countably many points. Since any selection $q_L^* \in \mathcal{Q}_L^*$ is strictly decreasing in $\mu \in [0, \mu_0)$, q_L^* is discontinuous only at countably many points (see, e.g., Rudin, 1976, p. 96). Let $\mathcal{D}$ denote the (countable) set of discontinuity points in $[0, \mu_0]$. Suppose that $\mathcal{Q}_L^*$ is not single-valued at $\check{\mu} \in [0, \mu_0] \backslash \mathcal{D}$, so that there exists $q_0 \in \mathcal{Q}_L^*(\check{\mu}) \setminus \{q_1\}$, where $q_1 = q_L^*(\check{\mu}) \in \mathcal{Q}_L(\check{\mu})$. Since q_L^* is continuous at $\check{\mu}$, for all $\varepsilon > 0$, there exists $\delta > 0$ such that for all μ' satisfying $|\mu' - \tilde{\mu}| < \delta$, it is $|q_L^*(\mu') - q_1| < \varepsilon$. If $q_0 < q_1$, we fix $\varepsilon = (q_1 - q_0)/2$, so

$$q_L^*(\mu') - q_1 > \frac{q_0 - q_1}{2} \quad \Rightarrow \quad q_L^*(\mu') > q_0.$$

Then for the selection $\check{q}_L(\mu) = q_L^*(\mu) + \mathbf{1}_{\{\mu = \tilde{\mu}\}}(q_0 - q_1)$ which modifies q_L^* only at a single point, we have $\check{q}_L(\mu') = q_L^*(\mu') > q_0$ for $\mu' \in (\tilde{\mu}, \tilde{\mu} + \delta)$, so that $\check{q}_L$ is not decreasing, leading to a contradiction. The case where $q_0 > q_1$ can be treated in a similar manner.

$\Pi^*(\mu)$ is continuous due to the maximum theorem. For $\mu \in [0, \mu_0]\backslash\mathcal{D}$, $\mathcal{Q}_L^*(\mu)$ is single-valued and then $\{\partial_\mu F(q_L^*, \mu) | q_L^* \in \mathcal{Q}_L^*\}$ is a singleton. Since $F(q, \mu)$ and $\partial_\mu F(q, \mu)$ are continuous in (q, μ), by Milgrom & Segal (2002, Cor. 4) the total derivative of $\max_{q \in [0, \bar{q}]} F(q, \mu)$ with respect to μ is $-\big(u(q_L^*, \theta_H) - C(q_L^*)\big)$. Thus, $\Pi^*(\mu)$ is differentiable at any $\mu \in [0, \mu_0]\backslash\mathcal{D}$, and

$$\frac{d\Pi^*(\mu)}{d\mu} = \Pi_H^*(\mu) - \Pi_L^*(\mu) \geq 0.$$

By the Goldowsky-Tonelli Theorem (see, e.g., Saks, 1937, p. 206) $\Pi^*(\mu)$ is increasing on $[0, \mu_0]$.

This completes our proof. □

Proof of Remark 3. We begin by considering two particular menus, namely $(\boldsymbol{p}^1, \boldsymbol{q}^1)$, with

$$\boldsymbol{p}^1 = (u(q_L^*(0), \theta_L), u(q_L^*(0), \theta_L)) \quad \text{and} \quad \boldsymbol{q}^1 = (\mathrm{q}_L^*(0), \mathrm{q}_L^*(0)),$$

and $(\boldsymbol{p}^2, \boldsymbol{q}^2)$, with

$$\boldsymbol{p}^2 = (0, u(q_H^*, \theta_H)) \quad \text{and} \quad \boldsymbol{q}^2 = (0, \mathrm{q}_H^*),$$

where the positive quality levels $q_L^*(0)$ and q_H^* are given by Eq. (7) in Lemma 1. Both menus are feasible, in the sense that by virtue of Eq. (1) they both satisfy Eqs. (3) and (4). The corresponding expected profits are

$$\begin{aligned}\Pi(\boldsymbol{p}^1, \boldsymbol{q}^1) &= u(q_L^*(0), \theta_L) - C(q_L^*(0)) \\ &= \max_{q_L \geq 0} \{u(q_L, \theta_L) - C(q_L)\} = \Pi^*(0) > 0,\end{aligned}$$

and

$$\begin{aligned}\Pi(\boldsymbol{p}^2, \boldsymbol{q}^2) &= \mu\big(u(q_H^*, \theta_L) - C(q_H^*)\big) \\ &= \mu \max_{q_H \geq 0} \{u(q_H, \theta_H) - C(q_H)\} = \mu\Pi^*(1) \geq 0,\end{aligned}$$

respectively, taking into account that by Eq. (2) the optimal profit $\Pi^*(0)$ is strictly positive. As a result, we can conclude that

$$\Pi^*(\mu) \geq \max\{\Pi^*(0), \mu\Pi^*(1)\} > 0, \qquad \mu \in [0, 1],$$

which yields our claim.[24] □

Proof of Proposition 1. Since the utility and cost functions are single-valued and continuous in q, they are also upper hemicontinuous in q by definition and so are Π_L, Π_H. Due to upper hemicontinuity of $\mathcal{Q}_L^*$, the composite correspondences $\Pi_L(\mathcal{Q}_L^*(\cdot))$ and $\Pi_H(\mathcal{Q}_L^*(\cdot), q_H^*)$ are upper hemicontinuous (see, e.g., Aliprantis & Border, 2006, p. 566). Given the assumption of Eq. (2), $\Pi_L(\mathcal{Q}_L^*(\cdot))$ and $\Pi_H(\mathcal{Q}_L^*(\cdot), q_H^*)$ are both compact-valued so that $(1 - \mu)\Pi_L(\mathcal{Q}_L^*(\cdot)) + \mu\Pi_H(\mathcal{Q}_L^*(\cdot), q_H^*)$ is upper hemicontinuous (see, e.g., Aliprantis & Border, 2006, p. 571). Since $\mathcal{Q}_L^*$ is compact-valued, $\hat{\Pi}(\hat{\mu}, \mu)$ as the minimum of $(1 - \mu)\Pi_L(\mathcal{Q}_L^*(\cdot)) + \mu\Pi_H(\mathcal{Q}_L^*(\cdot), q_H^*)$ is lower semicontinuous in $\hat{\mu}$ (see, e.g., Moore, 1999, p. 132). Thus, the performance ratios for any fixed μ are lower semicontinuous. Then the performance index, as the infimum of lower semicontinuous functions over a compact interval [0,1], is also lower semicontinuous (see, e.g., Aubin, 1998, p. 14).

Now we prove the representation of the performance index. Note that by Eq. (9) it is

$$\frac{d\hat{\Pi}(\hat{\mu}, \mu)}{d\mu} = \Pi_H^*(\hat{\mu}) - \Pi_L^*(\hat{\mu}), \quad (\hat{\mu}, \mu) \in [0, 1] \times [0, 1].$$

By Corollary 1, the identity also holds for $\hat{\mu} = \mu \in [0, \mu_0]$ except countably many points. Thus, differentiating the performance ratio $\varphi(\hat{\mu}, \mu)$ in Eq. (10) with respect to μ leads to

$$\partial_\mu \varphi(\hat{\mu}, \mu) = \frac{\Pi_L^*(\mu)\Pi_H^*(\hat{\mu}) - \Pi_H^*(\mu)\Pi_L^*(\hat{\mu})}{(\Pi^*(\mu))^2}, \quad \hat{\mu} \in [0, 1],$$

which holds for $\mu \in [0, \mu_0]$ except countably many points. Since by Eq. (11) it is $\Pi^*(\mu) = \hat{\Pi}(\mu, \mu) > 0$, $\varphi(\hat{\mu}, \cdot)$ attains its global maximum (of 1) at $\mu = \hat{\mu}$. Then for any $\hat{\mu} < \mu_0$ such that $\varphi(\hat{\mu}, \cdot)$ is differentiable at $\mu = \hat{\mu}$, we have $\partial_\mu \varphi(\hat{\mu}, \hat{\mu}) = 0$. On the other hand, Remark 3 implies that either Π_H^* or Π_L^* is positive. By Corollary 1, Π_L^* is strictly decreasing and Π_H^* is strictly increasing when $\mu < \mu_0$, which ensures that $\Pi_L^*(\cdot)\Pi_H^*(\hat{\mu}) - \Pi_H^*(\cdot)\Pi_L^*(\hat{\mu})$ is strictly decreasing on $\mu < \mu_0$. Thus, $\partial_\mu \varphi(\hat{\mu}, \cdot)$ is equal to 0, as long as it exists, for at most one point $\mu = \hat{\mu}$ on $[0, \mu_0)$. Then $\partial_\mu \varphi(\hat{\mu}, \cdot)$ is positive for $\mu < \min\{\hat{\mu}, \mu_0\}$ except countably many points. Since the performance ratio is continuous in μ, by the Goldowsky–Tonelli Theorem again, $\varphi(\hat{\mu}, \cdot)$ is increasing for $\mu \leq \min\{\hat{\mu}, \mu_0\}$.

For $\mu \geq \mu_0$, the performance ratio is equal to $(\Pi_H^*(\hat{\mu}) + (1/\mu - 1)\Pi_L^*(\hat{\mu}))/\Pi^*(1)$, which is decreasing in μ. Let $\hat{\mu} < \mu_0$. For $\hat{\mu} < \mu < \mu_0$, wherever $\varphi(\hat{\mu}, \cdot)$ is differentiable, $\partial_\mu \varphi(\hat{\mu}, \mu)$ is negative, so that $\varphi(\hat{\mu}, \cdot)$ is decreasing.

As a result, the performance ratio is always increasing in μ when $\mu \leq \min\{\hat{\mu}, \mu_0\}$ and decreasing otherwise, so for any $\hat{\mu} \in [0, 1]$ it is

$$\varphi(\hat{\mu}, \mu) \geq \min\{\varphi_0(\hat{\mu}), \varphi_1(\hat{\mu})\},$$

implying that the performance ratio is quasiconcave in μ and attains its minimum with respect to μ necessarily at the boundary of its domain [0,1],

$$\inf_{\mu \in [0,1]} \varphi(\hat{\mu}, \mu) = \min\{\varphi_0(\hat{\mu}), \varphi_1(\hat{\mu})\},$$

concluding our proof. □

Proof of Proposition 2. By Eqs. (10) and (14), the firm's performance index is of the form

$$\rho(\hat{\mu}) = \min\{\varphi_0(\hat{\mu}), \varphi_1(\hat{\mu})\} = \min\left\{\frac{\Pi_L^*(\hat{\mu})}{\Pi_L^*(0)}, \frac{\Pi_H^*(\hat{\mu})}{\Pi_H^*(1)}\right\},$$

for $\hat{\mu} \in [0, 1]$.

Moreover, by Corollary 1, $\Pi_H^*(\cdot)$ is strictly increasing, whereas $\Pi_L^*(\cdot)$ is strictly decreasing on $[0, \mu_0)$. By choosing a selection to attain the minimum of $\Pi_L(\mathcal{Q}_L^*(\hat{\mu}))$ and $0 \in \mathcal{Q}_L^*(\mu_0)$, it is $\rho(\hat{\mu}) = \varphi_0(\hat{\mu}) = 0$, for all $\hat{\mu} \geq \mu_0$. As a result, the difference between the boundary performance ratios,

$$\Delta(\hat{\mu}) = \varphi_1(\hat{\mu}) - \varphi_0(\hat{\mu}), \quad \hat{\mu} \in [0, \mu_0),$$

is strictly increasing, with [25]

$$\Delta(0) = \varphi_1(0) - 1 < 0 \leq \Delta(\mu_0).$$

Since $\varphi_1(\cdot)$ is strictly increasing and $\varphi_0(\cdot)$ is strictly decreasing for $\mu < \mu_0$, there exists $\hat{\mu} \in (0, \mu_0)$ (Eq. (8) shows $\mu_0 > 0$) such that $\varphi_0(\hat{\mu})\varphi_1(\hat{\mu}) > 0$, which implies that the optimal performance index is positive. Furthermore, $\sup \rho([0, 1])$ in Eq. (13) is achieved

[24] The intuition is that the optimal expected profit can never be less than what the firm could obtain by choosing the best option among serving either both types together (which yields $\Pi^*(0)$ using the menu $(\boldsymbol{p}^1, \boldsymbol{q}^1)$) or concentrating on the type-θ_H consumers (which yields $\mu\Pi^*(1)$ using the menu $(\boldsymbol{p}^2, \boldsymbol{q}^2)$). The resulting lower bound is tight for $\mu \in \{0, 1\}$.

[25] It is $\varphi_1(0) = 1 - \big(u(q_L^*(0), \theta_H) - u(q_L^*(0), \theta_L)\big)/\big(u(q_H^*, \theta_H) - C(q_H^*)\big) < 1$ and $\varphi_0(\cdot) = 0$ on $[\mu_0, 1]$.

at $\hat{\mu}^* \le \mu_0$ given in Eq. (15) (possibly by a converging sequence of candidate beliefs as described in footnote 14).

If $\mathcal{Q}_L^*(\cdot)$ is single-valued on $[0, \mu_0]$, it is also continuous, implying the continuity of $\rho(\cdot)$, which in turn means $\hat{\mu}^* = \{\hat{\mu} \in (0, \mu_0) : \Delta(\hat{\mu}) = 0\}$. In that case, any solution $\hat{\mu}^*$ to the robust identification problem (13) is such that $\Delta(\hat{\mu}^*) = 0$. Indeed, by Eq. (30) it is $\mathcal{Q}_L^*(\mu_0) = \{0\}$, so $\varphi_0(\mu_0) = 0$ and $\varphi_1(\mu_0) = 1$. Then the difference $\Delta(\mu_0)$ is positive and $\hat{\mu}^* < \mu_0$, which concludes our proof. □

Proof of Lemma 2. Since $\boldsymbol{\theta}^m = (\theta_L^m, \theta_H^m)$ lies in the open set $\hat{\Theta}$, it is an interior minimizer of the performance index ρ, so the first-order necessary optimality condition $(\partial_{\theta_L}\rho^*, \partial_{\theta_H}\rho^*) = 0$ has to be satisfied at $\boldsymbol{\theta} = \boldsymbol{\theta}^m$. By definition, the boundary performance ratios are

$$\varphi_0(\hat{\mu}^*) = \frac{\Pi_L^*(\hat{\mu}^*)}{\Pi_L^*(0)} \quad \text{and} \quad \varphi_1(\hat{\mu}^*) = \frac{\Pi_H^*(\hat{\mu}^*)}{\Pi_H^*(1)}.$$

Eq. (17) leads to two alternative representations of the first-order necessary optimality condition, namely

$$\partial_{\theta_L}\rho^* = \frac{\Pi_L^*(\hat{\mu}^*)}{\Pi_L^*(0)}\left(\partial_{\theta_L}\log\Pi_L^*(\hat{\mu}^*) - \partial_{\theta_L}\log\Pi_L^*(0)\right) = 0,$$
$$\partial_{\theta_H}\rho^* = \frac{\Pi_L^*(\hat{\mu}^*)}{\Pi_L^*(0)}\partial_{\theta_H}\log\Pi_L^*(\hat{\mu}^*) = 0,$$

which yields conditions (a), or alternatively,

$$\partial_{\theta_L}\rho^* = \frac{\Pi_H^*(\hat{\mu}^*)}{\Pi_H^*(1)}\partial_{\theta_L}\log\Pi_H^*(\hat{\mu}^*) = 0,$$
$$\partial_{\theta_H}\rho^* = \frac{\Pi_H^*(\hat{\mu}^*)}{\Pi_H^*(1)}\left(\partial_{\theta_H}\log\Pi_H^*(\hat{\mu}^*) - \partial_{\theta_H}\log\Pi_H^*(1)\right) = 0,$$

which gives conditions (b). Identifying the components of the gradient vector across the two representations, we further obtain that $(a_1) \Leftrightarrow (b_1)$ and $(a_2) \Leftrightarrow (b_2)$, as claimed. □

Proof of Proposition 3. By the definition of $\Pi_L(\boldsymbol{q})$ and $\Pi_H(\boldsymbol{q})$, the partial derivatives of $\Pi_L^*(\hat{\mu}^*)$ with respect to θ_L, θ_H are

$$\partial_{\theta_L}\Pi_L^*(\hat{\mu}^*) = \left(u_q(\hat{q}_L^*, \theta_L) - C'(\hat{q}_L^*)\right)(\partial_{\theta_L}\hat{q}_L^*) + u_\theta(\hat{q}_L^*, \theta_L),$$
$$\partial_{\theta_H}\Pi_L^*(\hat{\mu}^*) = \left(u_q(\hat{q}_L^*, \theta_L) - C'(\hat{q}_L^*)\right)(\partial_{\theta_H}\hat{q}_L^*),$$

or equivalently, using the relevant first-order condition in Eq. (7),

$$\partial_{\theta_L}\Pi_L^*(\hat{\mu}^*) = \tfrac{\hat{\mu}^*}{1-\hat{\mu}^*}\left(u_q(\hat{q}_L^*, \theta_H) - u_q(\hat{q}_L^*, \theta_L)\right)(\partial_{\theta_L}\hat{q}_L^*) + u_\theta(\hat{q}_L^*, \theta_L),$$
$$\partial_{\theta_H}\Pi_L^*(\hat{\mu}^*) = \tfrac{\hat{\mu}^*}{1-\hat{\mu}^*}\left(u_q(\hat{q}_L^*, \theta_H) - u_q(\hat{q}_L^*, \theta_L)\right)(\partial_{\theta_H}\hat{q}_L^*).$$

Meanwhile, the partial derivatives of $\Pi_H^*(\hat{\mu}^*)$ are

$$\partial_{\theta_L}\Pi_H^*(\hat{\mu}^*) = -\left(u_q(\hat{q}_L^*, \theta_H) - u_q(\hat{q}_L^*, \theta_L)\right)(\partial_{\theta_L}\hat{q}_L^*) + u_\theta(\hat{q}_L^*, \theta_L),$$
$$\partial_{\theta_H}\Pi_H^*(\hat{\mu}^*) = \left(u_q(\hat{q}_L^*, \theta_L) - u_q(\hat{q}_L^*, \theta_H)\right)(\partial_{\theta_H}\hat{q}_L^*) + u_\theta(q_H^*, \theta_H) - u_\theta(\hat{q}_L^*, \theta_H),$$

where we have used that by Eq. (7) the optimal quality q_H^* does not depend on θ_L. At the boundary of the belief domain [0,1] we obtain

$$(\partial_{\theta_L}\Pi_L^*(0), \partial_{\theta_H}\Pi_L^*(0)) = (u_\theta(q_L^*(0), \theta_L), 0),$$
$$(\partial_{\theta_L}\Pi_H^*(1), \partial_{\theta_H}\Pi_H^*(1)) = (0, u_\theta(q_H^*, \theta_H)).$$

Differentiating Eq. (16) yields that $(\partial_{\theta_L}\Delta(\hat{\mu}^*), \partial_{\theta_H}\Delta(\hat{\mu}^*))$ is equal to

$$\tfrac{\partial_{\theta_L}\Pi_H^*(\hat{\mu}^*)}{\Pi_H^*(1)} - \tfrac{\partial_{\theta_L}\Pi_L^*(\hat{\mu}^*)}{\Pi_L^*(0)} + \tfrac{\partial_{\theta_L}\Pi_L^*(0)}{\Pi_L^*(0)}\rho^* = 0,$$
$$\tfrac{\partial_{\theta_H}\Pi_H^*(\hat{\mu}^*)}{\Pi_H^*(1)} - \tfrac{\partial_{\theta_H}\Pi_L^*(\hat{\mu}^*)}{\Pi_L^*(0)} - \tfrac{\partial_{\theta_H}\Pi_H^*(1)}{\Pi_H^*(1)}\rho^* = 0.$$

Combining this with conditions (a_2) and (b_1) gives:

$$\rho_{\mathrm{WC}}^* = \frac{\partial_{\theta_L}\Pi_L^*(\hat{\mu}^*)}{\partial_{\theta_L}\Pi_L^*(0)} = \frac{\partial_{\theta_H}\Pi_H^*(\hat{\mu}^*)}{\partial_{\theta_H}\Pi_H^*(1)}. \tag{31}$$

Note that by our expression of $\partial_{\theta_H}\Pi_L^*(\hat{\mu}^*)$ and $\partial_{\theta_L}\Pi_H^*(\hat{\mu}^*)$, conditions (a_2) and (b_1) imply,

$$\partial_{\theta_H}\hat{q}_L^* = 0 \quad \text{and} \quad \partial_{\theta_L}\hat{q}_L^* = \frac{u_\theta(\hat{q}_L^*, \theta_L)}{u_q(\hat{q}_L^*, \theta_H) - u_q(\hat{q}_L^*, \theta_L)},$$

respectively, whence by substitution into Eq. (31) (using the relevant expressions derived earlier in this proof) we obtain:

$$\rho_{\mathrm{WC}}^* = \frac{1}{1-\hat{\mu}^*}\frac{u_\theta(\hat{q}_L^*, \theta_L^m)}{u_\theta(q_L^*(0), \theta_L^m)} = 1 - \frac{u_\theta(\hat{q}_L^*, \theta_H^m)}{u_\theta(q_H^*, \theta_H^m)}.$$

Note that $\hat{q}_L^*$ must be smaller than q_H^* (otherwise all agents would purchase the product of higher quality at a lower price). Also, $\hat{q}_L^*$ cannot vanish, as otherwise $\varphi_1(\hat{\mu}^*) = 1$ and $\varphi_0(\hat{\mu}^*) = 0$, in contradiction to Proposition 2. We therefore find that $u_\theta(\hat{q}_L^*, \theta_H^m) < u_\theta(q_H^*, \theta_H^m)$, since by assumption $u_{q\theta} > 0$. Thus, we conclude that ρ_{WC}^* is positive. □

Proof of Proposition 4. Since the ambiguity set is by assumption nonempty and compact, its bounds are achieved: $\mu_1 = \min\mathcal{A}$ and $\mu_2 = \max\mathcal{A}$. They are such that $0 \le \mu_1 \le \mu_2 \le 1$.

(i) The proof of Proposition 1 establishes the fact that for any given candidate belief $\hat{\mu} \in [0,1]$ the function $\varphi(\hat{\mu}, \cdot)$ is nondecreasing on $[0, \min\{\hat{\mu}, \mu_0\}]$ and nonincreasing on $[\min\{\hat{\mu}, \mu_0\}, 1]$. As a result, its minimum on the compact ambiguity set $\mathcal{A}$ must be achieved at μ_1 or μ_2, so

$$\rho(\hat{\mu}|\mathcal{A}) = \min\{\varphi(\hat{\mu}, \mu_1), \varphi(\hat{\mu}, \mu_2)\}, \quad \hat{\mu} \in [0,1],$$

which, similar to the proof of Proposition 1, is lower semicontinuous, thus establishing Eq. (19).

(ii) Since $\mathcal{Q}_L^*(\cdot)$ is single-valued, $\mathcal{Q}_L^*(\hat{\mu}) = \{0\}$ for all $\hat{\mu} \ge \mu_0$ by Corollary 1, and if $\hat{\mu} \ge \mu_0$, the performance ratio is equal to $\mu(u(q_H^*, \theta_H) - C(q_H^*))/\Pi^*(\mu)$, which is constant in $\hat{\mu}$. Furthermore, q_L^* and the performance ratio are continuous in $\hat{\mu}$. We claim that $\hat{\Pi}(\hat{\mu}, \mu)$ is increasing in $\hat{\mu}$ when $\hat{\mu} < \min\{\mu, \mu_0\}$, and is decreasing when $\mu \le \hat{\mu} < \mu_0$. We note that the partial derivative of $(1-\mu)\Pi_L(q_L) + \mu\Pi_H(\boldsymbol{q})$ with respect to q_L evaluated at $q_L = q_L^*(\hat{\mu})$ is equal to

$$\partial_q F(q_L^*(\hat{\mu}), \mu) = (1-\mu)(u_q(q_L^*(\hat{\mu}), \theta_L) - C'(q_L^*(\hat{\mu}))) - \mu(u_q(q_L^*(\hat{\mu}), \theta_H) - u_q(q_L^*(\hat{\mu}), \theta_L)),$$

which must be 0 when $\mu = \hat{\mu} < \mu_0$ due to the first-order condition of Eq. (7). As shown in the proof of Corollary 1, $u_q(q_L^*, \theta_L) - C'(q_L^*)$ is positive, implying that $\partial_q F(q_L^*(\hat{\mu}), \mu)$ is strictly decreasing in μ. Thus, when $\hat{\mu} < \min\{\mu, \mu_0\}$, the partial derivative of $(1-\mu)\Pi_L(q_L) + \mu\Pi_H(\boldsymbol{q})$ with respect to q_L evaluated at $q_L = q_L^*(\hat{\mu})$ must be negative, and then in a neighborhood of $q_L^*(\hat{\mu})$, $(1-\mu)\Pi_L(q_L) + \mu\Pi_H(\boldsymbol{q})$ is decreasing in q_L, which implies that $\hat{\Pi}(\hat{\mu}, \mu)$ is increasing in $\hat{\mu}$ since q_L^* is decreasing and continuous. Similarly, we can prove that $\hat{\Pi}(\hat{\mu}, \mu)$ is decreasing when $\mu \le \hat{\mu} < \mu_0$. $\varphi(\hat{\mu}, \mu)$ has the same monotonicity as $\hat{\Pi}(\hat{\mu}, \mu)$ with respect to $\hat{\mu}$.

Consider the difference of boundary performance ratios, $\Delta(\cdot|\mathcal{A}) = \varphi(\cdot, \mu_2) - \varphi(\cdot, \mu_1)$, which is increasing on the interval $[\mu_1, \mu_2]$. It holds that

$$\Delta(\mu_1|\mathcal{A}) = \varphi(\mu_1, \mu_2) - 1 \le 0 \le 1 - \varphi(\mu_2, \mu_1) = \Delta(\mu_2|\mathcal{A}).$$

Hence, there exists $\hat{\mu} \in [\mu_1, \mu_2]$ such that $\Delta(\hat{\mu}|\mathcal{A}) = 0$. Since two boundary performance ratios are increasing in $\hat{\mu} \le \mu_1$ and constant in $\hat{\mu} \in [\mu_0, 1]$, any optimal robust belief $\hat{\mu}^*$ must satisfy $\Delta(\hat{\mu}^*|\mathcal{A}) = 0$, i.e., Eq. (20).

This completes our proof. □

Proof of Lemma 3. Fix $\boldsymbol{\theta} = (\theta_L, \theta_H) \in \hat{\Theta}$, inducing the belief threshold $\mu_0 = \theta_L/\theta_H$ in the open interval (0,1). The performance ratios $\varphi_0(\cdot)$ and $\varphi_1(\cdot)$ are both continuous on [0,1]. To obtain $\hat{\mu}^* \in$

$(0, \mu_0)$ from $\Delta(\hat{\mu}^*) = 0$ by Proposition 2, we introduce the variables $\hat{\eta} = \hat{\mu}/(1-\hat{\mu})$ and $\eta_0 = \mu_0/(1-\mu_0)$ with values in $(0, \infty)$, so $\hat{\eta}^* = \hat{\mu}^*/(1-\hat{\mu}^*)$. In the transformed variable $\hat{\eta}$, the boundary performance ratios $\hat{\varphi}_i(\hat{\eta}) = \varphi_i(\hat{\eta}/(1+\hat{\eta}))$, for $i \in \{0, 1\}$, become

$$\hat{\varphi}_0(\hat{\eta}) = 1 - \frac{\hat{\eta}^2}{\eta_0^2} \quad \text{and} \quad \hat{\varphi}_1(\hat{\eta}) = \mu_0^2\left(1 + \frac{1+2\hat{\eta}}{\eta_0^2}\right),$$

respectively. These two boundary performance ratios are equal if and only if $\hat{\eta} = \hat{\eta}^*$ is a root of a quadratic polynomial,

$$\hat{\eta}^2 + 2\mu_0^2\hat{\eta} + \mu_0^2 - (1-\mu_0^2)\eta_0^2 = 0,$$

whence

$$\hat{\eta}^* = -\mu_0^2 \pm \sqrt{\mu_0^4 - (\mu_0^2 - \eta_0^2(1-\mu_0^2))}.$$

Since $\hat{\eta}^*$ cannot be negative, we can restrict attention to the corresponding unique nonnegative root:

$$\begin{aligned}\hat{\eta}^* &= -\mu_0^2 + \sqrt{\mu_0^4 - (\mu_0^2 - \eta_0^2(1-\mu_0^2))} \\ &= \mu_0^2\left(\sqrt{\frac{(2-\mu_0)(1+\mu_0)}{\mu_0(1-\mu_0)}} - 1\right).\end{aligned}$$

Using transformation $\hat{\mu}^* = 1/(1+1/\hat{\eta}^*)$, we can represent $\hat{\mu}^*$ in the form

$$\begin{aligned}\hat{\mu}^* &= \frac{2\mu_0^2}{\mu_0(1+\mu_0) + \sqrt{\mu_0(1+\mu_0)(1-\mu_0)(2-\mu_0)}} \\ &= \mu_0\left(1 - \frac{\sqrt{\mu_0^2(1-\mu_0)^2 + 2\mu_0(1-\mu_0)} - \mu_0(1-\mu_0)}{\sqrt{\mu_0^2(1-\mu_0)^2 + 2\mu_0(1-\mu_0)} + \mu_0(1+\mu_0)}\right),\end{aligned}$$

for all $\mu_0 \in (0, 1)$, which establishes Eq. (23). In addition, setting $x = \mu_0(1-\mu_0) \in (0, 1)$ we have [26]

$$0 < \frac{\hat{\mu}^*(\mu_0)}{\mu_0} < 1 - \frac{\sqrt{x^2+2x} - x}{\sqrt{x^2+2x} + 1 - \sqrt{1-4x}} < 1, \quad \mu_0 \in (0, 1/2],$$

so that—by taking the limit for $\mu_0 \to 0^+$ (and thus, $x \to 0^+$): $\hat{\mu}^*(0^+) = 0$, and (by the definition of a derivative) $\hat{\mu}_0^{*\prime}(0^+) = 0$. On the other hand, substituting $\mu_0 = 1$ in Eq. (23) implies $\hat{\mu}^*(1^-) = \hat{\mu}^*(1) = 1$. Furthermore, by direct differentiation,

$$\hat{\mu}^{*\prime}(\mu_0) = \left(\frac{3 - \mu_0(1+\mu_0) + \sqrt{\mu_0^2(1-\mu_0)^2 + 2\mu_0(1-\mu_0)}}{\sqrt{\mu_0^2(1-\mu_0)^2 + 2\mu_0(1-\mu_0)}\left(\sqrt{\mu_0^2(1-\mu_0)^2 + 2\mu_0(1-\mu_0)} + \mu_0(1+\mu_0)\right)}\right)\hat{\mu}^*(\mu_0),$$

for all $\mu_0 \in (0, 1)$, which implies that $\hat{\mu}^{*\prime}(\mu_0) \to \infty$ for $\mu_0 \to 1^-$ and $\hat{\mu}^{*\prime}(\mu_0) \geq 0$. Finally, since $\hat{\mu}^{*\prime\prime}(\mu_0) \geq \hat{\mu}^{*\prime\prime}(1/2) = 16/27 > 0$, for all $\mu_0 \in (0, 1)$, the function $\hat{\mu}^*(\cdot)$ is strictly convex on (0,1). □

[26] Here we use the fact that $\mu_0(1+\mu_0) \leq 2\mu_0 = 1 - \sqrt{1-4x}$, as $x = \mu_0(1-\mu_0) \in [0, 1/4]$, for $\mu_0 \in [0, 1/2]$.

Appendix B. Notation

Symbol	Description	Domain/Definition
$\mathcal{A}$	Ambiguity set of beliefs (compact)	$\mathcal{A} \subseteq [0, 1]$
$C(\cdot)$	Cost function	$C : \mathcal{Q} \to \mathbb{R}_+$
IR	Information rent	$\mathbb{R}_+$
M, N	Sample size (beliefs, types)	$\mathbb{Z}_{++}$
$\mathcal{M}(\cdot)$	Solution to standard screening problem (5) (at μ)	$\mathcal{M} : [0, 1] \rightrightarrows \mathbb{R}_+^2 \times \mathcal{Q}^2$
$\boldsymbol{p} = (p_L, p_H)$	Price vector (p_L for low type, p_H for high type)	$\mathbb{R}_+^2$
$\boldsymbol{q} = (q_L, q_H)$	Quality vector (q_L for low type, q_H for high type)	$\mathcal{Q}^2$
$\mathcal{Q}$	Quality domain	$\mathbb{R}_+$
$\mathcal{Q}_L^*(\cdot)$	Optimal low-type quality correspondence (standard screening)	$\mathcal{Q}_L^* : [0, 1] \rightrightarrows \mathcal{Q}$
$u(\cdot, \cdot)$	Utility function	$u : \mathcal{Q} \times \Theta \to \mathbb{R}$
γ	Coefficient for cost function	$\mathbb{R}_{++}$
$\Delta(\cdot)$	Difference of boundary performance ratios (at $\hat{\mu}$)	$\Delta = \varphi_1 - \varphi_0$
$\boldsymbol{\theta} = (\theta_L, \theta_H)$	Type vector	$\hat{\Theta}$
Θ	Type space (for given $\boldsymbol{\theta}$)	$\{\theta_L, \theta_H\}$
$\hat{\Theta}$	Domain for type vectors $\boldsymbol{\theta}$, with boundary $\partial\hat{\Theta}$	$\{(\theta_L, \theta_H) \in \mathbb{R}_{++}^2 : 0 < \theta_L < \theta_H\}$
μ	Belief; subjective probability that random $\tilde{\theta}$ in Θ is high	$\text{Prob}(\tilde{\theta} = \theta_H) \in [0, 1]$
μ_0	Threshold for vanishing low-type quality	$\hat{\Theta}$
$\hat{\mu}$	Candidate belief	[0,1]
$\hat{\mu}^*(\cdot)$	Robust belief (as solution to Eq. (13)) (at μ_0)	$\hat{\mu}^* : (0, 1) \to (0, 1)$
$\Pi(\cdot, \cdot)$	Profit function (at $(\boldsymbol{p}, \boldsymbol{q})$)	$\Pi : \mathbb{R}_+^2 \times \mathcal{Q}^2 \to \mathbb{R}_+$
$\Pi_L(\cdot), \Pi_H(\cdot)$	Type-contingent profit functions (at $\boldsymbol{q}$)	$\Pi_L, \Pi_H : \mathcal{Q}^2 \to \mathbb{R}_+$
$\rho(\cdot)$	Performance index (at $\hat{\mu}$)	$\rho : [0, 1] \to [0, 1]$
ρ^*	Optimal performance index	[0,1]
$\varphi(\cdot, \cdot)$	Performance ratio (at $(\hat{\mu}, \mu)$)	$\varphi : [0, 1] \times [0, 1] \to [0, 1]$
$\varphi_0(\cdot), \varphi_1(\cdot)$	Boundary performance ratio (at $\mu \in \{0, 1\}$)	$\varphi(\cdot, 0), \varphi(\cdot, 1)$

References

Aliprantis, C. D., & Border, K. C. (2006). *Infinite dimensional analysis: A Hitchhiker's guide* (3rd ed.). Berlin, Heidelberg, DE: Springer.

Allouah, A., & Besbes, O. (2020). Prior-independent optimal auctions. *Management Science, 66*(10), 4417–4432.

Anderson, E. T., & Dana, D. D. (2009). When is price discrimination profitable? *Management Science, 55*(6), 980–989.

Aubin, J. P. (1998). *Optima and equilibria: An introduction to nonlinear analysis* (2nd ed.). Berlin, Heidelberg, DE.: Springer.

Azar, P., & Micali, S. (2013). Parametric digital auctions. In *Proceedings of the 4th conference on innovations in theoretical computer science (ITCS '13)* (pp. 231–232). New York, NY: Association for Computing Machinery (ACM).

Ben-David, S., & Borodin, A. (1994). A new measure for the study of on-line algorithms. *Algorithmica, 11*(1), 73–91.

Ben-Tal, A., El Ghaoui, L., & Nemirovski, A. (2009). *Robust optimization*. Princeton, NJ: Princeton University Press.

Bergemann, D., & Schlag, K. H. (2008). Pricing without priors. *Econometrica, 6*(2–3), 560–569.

Besbes, O., & Zeevi, A. (2009). Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research, 57*(6), 1407–1420.

Boyar, J., Irani, S., & Larsen, K. S. (2015). A comparison of performance measures for online algorithms. *Algorithmica, 72*(4), 969–994.

Caldentey, R., Liu, Y., & Lobel, I. (2017). Intertemporal pricing under minimax regret. *Operations Research, 65*(1), 104–129.

Caprari, E., Baiardi, L. C., & Molho, E. (2019). Primal worst and dual best in robust vector optimization. *European Journal of Operational Research, 275*(3), 830–838.

Carrasco, V., Luz, V. F., Kos, N., Messner, M., Monteiro, P., & Moreira, H. (2018). Optimal selling mechanisms under moment conditions. *Journal of Economic Theory, 177*, 245–279.

Carroll, G. (2017). Robustness and separation in multidimensional screening. *Econometrica, 85*(2), 453–488.

Carroll, G. (2019). Robustness in mechanism design and contracting. *Annual Review of Economics, 11*(1), 139–166.

Cassady, R. (1946). Techniques and purposes of price discrimination. *Journal of Marketing, 11*(2), 135–150.

Chandler, L. V. (1938). Monopolistic elements in commercial banking. *Journal of Political Economy, 46*(1), 1–22.

Destan, C. G., & Yılmaz, M. (2020). Nonlinear pricing under inequity aversion. *Journal of Economic Behavior and Organization, 169*, 223–244.

Dhangwatnotai, P., Roughgarden, T., & Yan, Q. (2015). Revenue maximization with a single sample. *Games and Economic Behavior, 91*, 318–333.

Doan, X. V., Lei, X., & Shen, S. (2020). Pricing of reusable resources under ambiguous distributions of demand and service time with emerging applications. *European Journal of Operational Research, 282*(1), 235–251.

Edlin, A. S., & Shannon, C. (1998). Strict monotonicity in comparative statics. *Journal of Economic Theory, 81*(1), 201–219.

Eren, S. S., & Maglaras, C. (2010). Monopoly pricing with limited demand information. *Journal of Revenue and Pricing Management, 9*(1), 23–48.

Fu, H., Immorlica, N., Lucier, B., & Strack, P. (2015). Randomization beats second price as a prior-independent auction. In *Proceedings of the sixteenth ACM conference on economics and computation* (p. 323).

Gibbard, A. (1973). Manipulation of voting schemes: A general result. *Econometrica, 41*(4), 587–601.

Goel, A., Meyerson, A., & Weber, T. A. (2009). Fair welfare maximization. *Economic Theory, 41*(3), 465–494.

Inada, K. I. (1963). On a two-sector model of economic growth: Comments and a generalization. *Review of Economic Studies, 30*(2), 119–127.

Kouvelis, P., & Yu, G. (1997). *Robust discrete optimization and its applications*. New York, NY: Springer.

Laplace, P. S. (1825). *Essai philosophique sur les probabilités*. Bachelier, Paris, France. [Reprinted by Cambridge University Press, Cambridge, UK, in 2009.]

Lewis, W. A. (1941). The two-part tariff. *Economica, New Series, 8*(31), 249–270.

Maskin, E., & Riley, J. (1984). Monopoly with incomplete information. *RAND Journal of Economics, 15*(2), 171–196.

Milgrom, P., & Segal, I. (2002). Envelope theorems for arbitrary choice sets. *Econometrica, 70*(2), 583–601.

Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *Review of Economic Studies, 38*(2), 175–208.

Moore, J. C. (1999). *Mathematical methods for economic theory 2*. Berlin, Heidelberg, DE: Springer.

Moorthy, K. S. (1984). Market segmentation, self-selection, and product line design. *Marketing Science, 3*(4), 288–307.

Mussa, M., & Rosen, S. (1978). Monopoly and product quality. *Journal of Economic Theory, 18*(2), 301–317.

Myerson, R. B. (1979). Incentive compatibility and the bargaining problem. *Econometrica, 47*(1), 61–74.

Neumann, J. v., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.

Pigou, A. C. (1920). *The economics of welfare*. London, UK: Macmillan.

Pınar, M. c., & Kızılkale, C. (2017). Robust screening under ambiguity. *Mathematical Programming, 163*(1–2), 273–299.

Rothschild, M., & Stiglitz, J. E. (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *Quarterly Journal of Economics, 90*(4), 629–649.

Rudin, W. (1976). *Principles of mathematical analysis* (3rd ed.). New York, NY: McGraw-Hill.

Saks, S. (1937). *Theory of the integral* (2nd ed.). New York, NY: Hafner Publishing.

Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical Association, 46*(253), 55–67.

Sleator, D. D., & Tarjan, R. E. (1985). Amortized efficiency of list update and paging rules. *Communications of the ACM, 28*(2), 202–208.

Stiglitz, J. E. (1975). The theory of 'screening,' education, and the distribution of income. *American Economic Review, 65*(3), 283–300.

Topkis, D. M. (1968). *Ordered optimal solutions*. Stanford University, Stanford, CA Doctoral dissertation.

Uzawa, H. (1963). On a two-sector model of economic growth II. *Review of Economic Studies, 30*(2), 105–118.

Villas-Boas, J. M. (1998). Product line design for a distribution channel. *Marketing Science, 17*(2), 156–169.

Wald, A. (1939). Contributions to the theory of statistical estimation and testing hypotheses. *Annals of Mathematics, 10*(4), 299–326.

Wald, A. (1945). Statistical decision functions which minimize the maximum risk. *Annals of Mathematics, 46*(2), 265–280.

Watkins, G. P. (1916). The theory of differential rates. *Quarterly Journal of Economics, 30*(4), 682–703.

Weber, T. A. (2022). Relatively robust decisions. *Theory and Decision*. https://doi.org/10.1007/s11238-022-09866-z. In Press

Wilson, R. (1987). Game-theoretic analyses of trading processes. In T. F. Bewley (Ed.), *Advances in economic theory: Fifth world congress (econometric society monographs)* (pp. 33–70). Cambridge, UK: Cambridge University Press.

Wilson, R. (1992). Strategic analysis of auctions. In R. J. Aumann, & S. Hart (Eds.), *Handbook of game theory with economic applications: vol. 1* (pp. 227–279). Amsterdam, NL: Elsevier.

Wong, H., Kim, K., & Chhajed, D. (2021). Reducing channel inefficiency in product line design. *International Journal of Production Economics, 232*. Art. 107964

Wong, K. P. (2020). Optimal nonlinear pricing by a regret-averse monopoly. *Managerial and Decision Economics, 41*(7), 1156–1161.

Zangwill, W. I., & Garcia, C. B. (1981). *Pathways to solutions, fixed points, and equilibria*. Englewood Cliffs, NJ: Prentice-Hall.

Zheng, M., Wang, C., & Li, C. (2015). Optimal nonlinear pricing by a monopolist with information ambiguity. *International Journal of Industrial Organization, 40*, 60–66.

Zou, T., Zhou, B., & Jiang, B. (2020). Product-line design in the presence of consumers' anticipated regret. *Management Science, 66*(12), 5665–5682.