Regulatory
Studies Center
THE GEORGE WASHINGTON UNIVERSITY

ALGORITHMIC PRICING AND REGULATION: A GAME-THEORETIC ANALYSIS OF FAIRNESS AND COMPETITION TRADEOFFS

Eliane Catilina

SEPTEMBER 2026

© 2026 by author(s). All rights reserved.

The GW Regulatory Studies Center is an academic center in the George Washington University and its Trachtenberg School of Public Policy and Public Administration. This working paper is part of a series on Regulatory Approaches to Emerging Technologies supported by the John Templeton Foundation. Any views expressed in this working paper are those of the author and not necessarily those of the John Templeton Foundation, the GW Regulatory Studies Center, or the Department of Transportation.

https://regulatorystudies.columbian.gwu.edu/

ABSTRACT

Artificial intelligence enables firms to estimate individual willingness to pay and implement personalized pricing at scale, raising concerns about fairness, transparency, and potential consumer harm. At the same time, these capabilities may intensify competition by enabling firms to target marginal consumers more effectively. Much of the policy debate, however, focuses on regulating algorithms or data rather than on the competitive interactions that ultimately determine pricing outcomes. This paper develops a game-theoretic framework to examine the regulation of algorithmic pricing in differentiated oligopolies. Modeling algorithmic pricing as the ability to implement consumer-specific pricing functions, the analysis embeds personalized pricing in a regulator–firm game in which policy instruments alter firms' pricing problems and equilibrium behavior.

The paper applies established results from the industrial organization literature to analyze how regulatory instruments, including pricing constraints, audits and penalties, and transparency requirements, affect markups, price dispersion, competition, and welfare through distinct economic mechanisms. The analysis highlights a central tradeoff between fairness and competition. Policies that compress price dispersion may also weaken competition for marginal consumers, while policies that preserve pricing flexibility may sustain competition at the cost of greater dispersion. The framework also yields empirical implications for evaluating pricing distributions and competitive behavior in both ex ante policy analysis and retrospective regulatory evaluation.

JEL CLASSIFICATION

D43; D82; L13; L51; O33

KEYWORDS

Algorithmic pricing; personalized pricing; price discrimination; regulatory economics; oligopoly; asymmetric information.

AUTHOR

Eliane Catilina is an economist at the U.S. Department of Transportation and an adjunct professor at the McCourt School of Public Policy at Georgetown University.

ACKNOWLEDGEMENTS

This project/ publication was made possible through the support of Grant 63641 from the John Templeton Foundation. The opinions expressed in this publication are those of the author and presented for academic discussion. They do not necessarily reflect the views of the John Templeton Foundation, the GW Regulatory Studies Center or the U.S. Department of Transportation or any of its agencies.

The author is grateful to Mary Sullivan and Susan Dudley of the George Washington University Regulatory Studies Center, Gabriel Movsesyan and Sam LaRussa from U.S. DOT, and to the participants in the GW Regulatory Studies Center workshop, *Regulatory Approaches to Emerging Technologies*, for helpful comments and suggestions. Any errors or omissions remain solely the author's responsibility.

DISCLAIMER

The views expressed in this paper are those of the author and are presented for academic discussion. They do not necessarily reflect the views or official positions of the U.S. Department of Transportation or any of its agencies.

ABOUT THE GEORGE WASHINGTON UNIVERSITY REGULATORY STUDIES CENTER

Established in 2009, the GW Regulatory Studies Center is an academic center of the George Washington University and its Trachtenberg School of Public Policy and Public Administration. The Center's mission is to improve regulatory policy through research, education, and outreach. The Center is a leading source for applied scholarship in regulatory issues, and a training ground for anyone who wants to understand the effects of regulation and ensure that regulatory policies are designed in the public interest.

I. Introduction

The increasing use of artificial intelligence in pricing has transformed how firms interact with consumers. Advances in data collection and machine learning enable firms to estimate individual willingness to pay and to adjust prices at a highly granular level. These developments have renewed policy concerns about price discrimination, particularly regarding fairness, transparency, and potential consumer harm (OECD, 2022; Zhao, 2023; Calvano et al., 2021).

Individualized pricing is no longer a purely theoretical possibility and is already observed across a range of industries. In the airline sector, firms have long relied on sophisticated revenue management systems to adjust fares based on demand conditions and booking behavior. In e-commerce, platforms such as Amazon adjust prices dynamically and have been reported to vary pricing across users and over time. Ride-hailing services like Uber implement algorithmic pricing that adjusts fares in real time based on demand, location, and user characteristics. Similarly, digital advertising markets rely extensively on individualized pricing mechanisms, where advertisers bid differently for users based on detailed profiles. These developments have been widely documented in policy and industry analyses (OECD, 2018; Federal Trade Commission, 2021; The Economist, 2019). These examples illustrate that the ability to tailor prices to individual consumers is already embedded in many sectors, reinforcing the need to understand its implications for competition and regulation.

Much of the policy debate has focused on regulating the inputs to algorithmic pricing systems, including data usage, transparency, and explainability. While these dimensions are important, they do not directly address the economic forces that determine pricing outcomes. Prices emerge from strategic interaction among firms, and algorithms operate as tools within that interaction rather than as independent decision-makers.

This observation has important implications for regulatory design. If firms use algorithmic pricing strategically, regulation must be evaluated within a framework that accounts for equilibrium responses. Policies that ignore strategic interaction risk mischaracterize both the problem and the effects of intervention.

More generally, evaluating algorithmic pricing requires understanding how regulation changes the strategic incentives firms face when competing for consumers. Policies with similar stated objectives may affect pricing behavior through very different mechanisms. Some policies directly constrain pricing strategies, while others operate indirectly by altering enforcement incentives or the informational environment surrounding pricing decisions. These distinctions are important because firms respond strategically to regulation, and equilibrium outcomes depend on how policy interventions reshape the underlying pricing game.

This paper develops a game-theoretic framework to analyze the regulation of algorithmic pricing. The model embeds personalized pricing in a differentiated oligopoly and introduces a

regulator who sets policy prior to firms' pricing decisions. Firms then choose pricing strategies in response to both market conditions and regulatory constraints, and equilibrium outcomes arise from this interaction. In this paper, algorithmic pricing is modeled as the ability of firms to implement consumer-specific pricing functions at scale.

The contribution of this paper is not to derive new results on personalized pricing itself. The effects of individualized pricing under competition are well established in the industrial organization literature (Thisse & Vives, 1988; Corts, 1998; Stole, 2007; Rhodes & Zhou, 2024). Rather, the contribution lies in embedding these results within a regulatory framework and using them to compare how different policy instruments affect firm incentives, equilibrium pricing behavior, and welfare outcomes.

More specifically, the paper develops a mechanism-based framework in which regulatory tools operate through distinct transformations of the firm's pricing problem. Constraint-based interventions, such as guardrails, directly limit feasible pricing strategies, whereas enforcement-based mechanisms alter incentives by altering expected regulatory costs. Transparency affects the informational environment that governs enforcement. By linking these mechanisms to equilibrium outcomes, the framework clarifies why policies with similar objectives may generate different effects on price dispersion, competition, and consumer welfare.

Existing research typically examines these margins separately. The industrial organization literature characterizes personalized pricing equilibria, the regulatory literature distinguishes among policy instruments, and the emerging AI policy literature emphasizes data, transparency, and accountability. This paper links these strands by asking how regulatory instruments change the mapping from consumer information to equilibrium prices, and how those changes affect competition, price dispersion, and welfare.

A central insight is that the effects of algorithmic pricing depend critically on market structure. In monopoly settings, improved information enables firms to extract more surplus (Varian, 1989). In oligopolistic markets, however, the same information may intensify competition by enabling firms to target marginal consumers more effectively. This distinction implies that regulatory interventions must be sensitive to the competitive environment in which they operate.

By modeling regulation as part of a strategic game, the paper provides a unified framework for evaluating policy tools based on how they alter firms' pricing incentives and equilibrium outcomes. This approach shifts the focus of regulation away from controlling algorithms per se and toward understanding how policy shapes the competitive use of information in pricing. The framework also generates empirical implications regarding price dispersion, markup compression, enforcement effects, and competitive responses, offering a mechanism-based perspective for both ex ante policy analysis and retrospective regulatory evaluation.

II. Regulatory Pricing Game

The analysis is based on a sequential game that captures the interaction between regulation and firm behavior.

Timing

The game unfolds in three stages:

1. A regulator selects a policy *R*, which governs the pricing environment. This policy may include pricing constraints, enforcement mechanisms, and transparency requirements.
2. Firms observe the regulatory policy and simultaneously choose pricing strategies.
3. Consumers observe prices and select the firm that provides the highest utility.

The solution concept is the subgame perfect equilibrium, ensuring that firms' pricing strategies are optimal given the regulatory policy and that the regulator anticipates these responses.

Consumers

Consumers are uniformly distributed along the unit interval $x \in [0,1]$.

A consumer located at *x* derives utility from firm *A* and firm *B* according to:

$$U_A(x) = v - tx - p_A(x),$$

$$U_B(x) = v - t(1 - x) - p_B(x),$$

where:

- v is the baseline valuation,
- $t > 0$ measures product differentiation,
- $p_i(x)$ is the price charged by firm *i* to consumer *x*.

A higher value of *t* implies weaker competition.

Market coverage

The market is assumed to be fully covered, meaning that all consumers purchase from one of the two firms. This assumption allows the analysis to focus on pricing decisions rather than participation.

Firms

Firms have constant marginal cost *c*. They may adopt:

- Uniform pricing: a single price p_i for all consumers
- Personalized pricing: a function $p_i(x)$

Key structural property

Under personalized pricing, firm profits can be written as:

$$\pi_i = \int_{D_i} (p_i(x) - c)dx,$$

where D_i is the set of consumers choosing firm i.

Because consumers make independent decisions, the pricing game decomposes into a continuum of independent subgames, one for each consumer x. Each subgame is a localized Bertrand competition with asymmetric advantages.

Regulatory policies modify these local games either by:

- constraining feasible prices, or
- altering firms' payoff structure.

This decomposition is exact under the maintained assumptions and forms the basis of the analysis in Section 6.

III. Benchmark: Uniform Pricing Competition

Under uniform pricing, each firm selects a single price that applies to all consumers. The marginal consumer is determined by the condition that the consumer is indifferent between the two firms:

$$v - tx^* - P_A = v - t(1 - x^*) - P_B.$$

Solving for x^* yields

$$x^* = \frac{1}{2} + \frac{P_B - P_A}{2t.}$$

Firm A's demand is therefore x^*, and its profit is

$$\pi_A = (P_A - c)\left(\frac{1}{2} + \frac{P_B - P_A}{2t}\right).$$

Maximizing this expression with respect to (P_A yields the first-order condition

$$\frac{1}{2} + \frac{P_B - P_A}{2t} - \frac{P_A - c}{2t} = 0.$$

Solving the system of best responses leads to the unique symmetric equilibrium

$$P_A = P_B = c + t, \qquad \pi_i = \frac{t}{2}$$

This result reflects the standard logic of spatial competition. Firms cannot perfectly discriminate among consumers, so they set a price that balances the desire to attract marginal consumers with the ability to extract surplus from less price-sensitive consumers.

IV. Personalized Pricing and Localized Competition

When firms can condition prices on consumer characteristics, the nature of competition changes fundamentally. Instead of choosing a single price, firms effectively choose a schedule of prices, one for each consumer.

For a given consumer *x*, the firm closest to the consumer has an advantage. This advantage can be expressed as

$$\Delta(x) = t(1 - 2x)$$

which is positive for consumers closer to firm *A* and negative for those closer to firm *B*.

The interaction between firms for each consumer takes the form of a Bertrand competition with asymmetric advantages. The firm with the advantage can charge a higher price, but only up to the point at which the consumer switches to the competitor.

Solving this local game yields an equilibrium in which the advantaged firm charges a markup equal to its advantage, while the disadvantaged firm prices at marginal cost. Aggregating across consumers yields total profits of

$$\pi_i = \frac{t}{4}.$$

The economic intuition behind this result is subtle yet important. Personalized pricing allows firms to extract more surplus from consumers who strongly prefer them. However, it also intensifies competition for consumers who are nearly indifferent. Because firms can target these consumers more precisely, competition becomes more aggressive at the margin.

In the covered-market case, increased competition outweighs increased extraction, leading to lower profits overall. This contrasts sharply with the monopoly case, where improved information unambiguously increases profits.

a. Market Structure and the Effects of Personalized Pricing

The contrast between monopoly and oligopoly is central to understanding the regulatory implications of algorithmic pricing. In a monopoly setting, individualized pricing allows the firm to extract the full surplus from each consumer, a standard result in the theory of price discrimination (Varian, 1989). Because the firm faces no competitive constraint, it can set prices equal to each consumer's willingness to pay, thereby eliminating consumer surplus. In this environment, improved information operates purely as a tool for surplus extraction, and concerns about exploitation are well-founded.

In an oligopoly, however, the role of individualized pricing changes fundamentally. Firms compete for each consumer, and pricing becomes a strategic interaction rather than a unilateral decision. While individualized pricing still allows firms to exploit their relative advantages, it also exposes them to more intense competition, particularly for consumers who are close to being indifferent between firms. The presence of a competitor limits how much any firm can extract surplus, since attempts to raise prices for a given consumer can be countered by rivals' targeted price reductions. As a result, individualized pricing may intensify competition and reduce equilibrium profits in certain settings, particularly in covered markets (Thisse & Vives, 1988; Stole, 2007; Rhodes & Zhou, 2024).

This distinction between monopoly and oligopoly is well established in the industrial organization literature and is not this paper's contribution. Rather, it provides the conceptual foundation for the analysis that follows. The paper's contribution lies in embedding these established results within a regulatory framework in which policy is modeled as part of a strategic game. By doing so, the analysis moves beyond characterizing pricing equilibria to examine how different regulatory instruments affect outcomes when firms respond strategically.

The implication is that the effects of algorithmic pricing depend critically on market structure, and that regulatory policy should therefore do the same. Policies grounded in monopoly intuition may be inappropriate in competitive environments. In oligopolistic markets, individualized pricing can lead to lower prices for marginal consumers by intensifying competition, even as it increases dispersion across consumers. The game-theoretic framework developed in this paper provides a way to evaluate these effects systematically, by linking regulatory interventions to equilibrium outcomes rather than to the underlying technology.

V. Regulatory Mechanisms and Pricing Incentives

The regulation of algorithmic pricing presents a distinct challenge because it lies at the intersection of technological capability and strategic firm behavior. While advances in artificial intelligence have expanded firms' ability to implement individualized pricing, the economic implications of these capabilities depend fundamentally on how firms deploy them in

competitive environments. As a result, the appropriate object of regulation is not the technology itself, but the pricing behavior and market outcomes that emerge from its use.

Much of the existing policy discussion has focused on regulating the inputs to algorithmic pricing systems, including data collection, algorithmic transparency, and explainability (OECD, 2022; Zhao, 2023). Concerns about algorithmic pricing also extend to the potential for tacit coordination among firms, as algorithms may facilitate more stable pricing strategies even in competitive environments (Calvano et al., 2021). These approaches reflect legitimate concerns about accountability and consumer protection, particularly in contexts where pricing decisions may be opaque or difficult to interpret. However, from an economic perspective, such approaches address the informational environment surrounding pricing rather than the strategic behavior that ultimately determines prices.

An alternative approach, grounded in the law and economics literature, emphasizes regulating outcomes rather than inputs. This perspective aligns with the broader regulatory tradition in industrial organization, where policy interventions are typically evaluated by their effects on prices, quantities, and welfare rather than on the specific technologies firms use (Stigler, 1964; Posner, 1975). In the context of algorithmic pricing, this suggests that regulatory efforts should focus on how pricing behavior affects competition and consumer outcomes, rather than on whether firms employ AI systems per se.

The framework developed in this paper adopts this outcome-based perspective but extends it by explicitly incorporating strategic interaction. Regulation is modeled as a first-stage decision in a sequential game, in which the regulator anticipates how firms will respond to policy constraints when setting prices. This structure captures a key feature of real-world regulation: policies do not operate in isolation but instead shape the incentives under which firms compete.

Within this framework, regulatory tools can be understood as interventions that modify the pricing game in different ways. Some policies impose direct constraints on pricing behavior, effectively limiting the set of strategies available to firms. Others operate through enforcement mechanisms, introducing expected costs that depend on firms' pricing decisions. This approach is consistent with the responsive regulation literature, which emphasizes incentive-based and enforcement-oriented interventions rather than rigid command-and-control rules (Ayres & Braithwaite, 1992). Still others affect the informational environment, thereby influencing regulators' and consumers' ability to observe and respond to pricing behavior. These distinctions reflect a central theme in regulatory economics: different policy instruments operate through distinct mechanisms and entail trade-offs between rules and standards, ex ante constraints and ex post enforcement (Kaplow, 1992; Shavell, 2004).

The advantage of embedding these tools within a game-theoretic framework is that it allows their effects to be evaluated in terms of equilibrium outcomes. Policies that pursue similar objectives may nevertheless generate different equilibrium outcomes because they alter firm incentives

through different mechanisms. For example, a pricing constraint may yield predictable but rigid outcomes, while an enforcement-based policy may offer greater flexibility but depend heavily on monitoring and compliance. Without explicitly modeling firm behavior, these differences are difficult to assess.

The distinction between monopoly and oligopoly further complicates the regulatory problem. As discussed in Section 4.1, individualized pricing has fundamentally different implications depending on the competitive environment. In monopoly settings, regulatory intervention is primarily concerned with limiting surplus extraction. In oligopolistic markets, however, regulation also interacts with competitive forces, potentially dampening or enhancing competition depending on its design. This reinforces the need for a framework that links regulatory instruments to equilibrium behavior rather than treating policy as a set of static constraints.

The contribution of this paper is to provide such a framework. By integrating established results on individualized pricing into a model of regulatory intervention, it allows for a systematic comparison of policy tools based on how they affect firm incentives, competitive dynamics, and welfare outcomes. In doing so, it bridges the gap between the theoretical literature on price discrimination and the emerging policy debate on algorithmic pricing.

More broadly, the framework highlights the importance of aligning regulatory design with the market's underlying economic mechanisms. In the context of algorithmic pricing, this means recognizing that AI does not fundamentally change the nature of competition but rather enhances firms' ability to implement existing pricing strategies. Effective regulation must therefore operate at the level of strategic interaction, shaping how firms use these tools rather than attempting to regulate the tools themselves.

To clarify how different regulatory tools operate within the pricing game, it is useful to summarize their effects in a unified representation. Each instrument can be interpreted as modifying the firm's optimization problem in a distinct way, either by constraining the feasible set of prices, altering the objective function, or changing the informational environment. Table 1 summarizes these mechanisms.

Table 1: Regulatory mechanisms

Policy Instrument	Simplified Representation	Effect on Firm Problem	Mechanism
Guardrails	Markup $\leq \kappa \times$ competitive advantage	Constrained maximization	Restricts feasible pricing

Policy Instrument	Simplified Representation	Effect on Firm Problem	Mechanism
Audits & penalties	Profit = markup − penalty(markup)	Modified objective function	Alters incentives
Transparency	Higher detection probability	Increases expected penalty	Strengthens enforcement

The table highlights a central distinction that will guide the analysis in the following section. Constraint-based policies, such as guardrails, directly limit the set of admissible pricing strategies. Incentive-based policies, such as audits and penalties, modify the tradeoffs faced by firms when choosing prices. Transparency, in turn, affects the effectiveness of enforcement by altering the informational environment. These differences imply that policies with similar objectives may produce different equilibrium outcomes once firms respond strategically. The formal analysis in Section 6 makes these mechanisms explicit.

VI. Equilibrium Pricing Under Regulation

a. Introduction

This section develops the pricing game formally and characterizes equilibrium outcomes under the regulatory interventions described in Section 5. The objective is to make explicit how regulation enters the firm's optimization problem and how it shapes equilibrium outcomes through strategic interaction.

As discussed in Section 2, firms choose pricing strategies that map consumer characteristics into prices. Because consumers are independent and the market is fully covered, firm profits can be expressed as the integral of profits across individual consumers. This implies that the pricing game decomposes into a continuum of independent subgames, one for each consumer $x \in [0,1]$.

In each subgame, firms compete for a single consumer by offering individualized prices, considering both their relative advantage and their competitor's pricing behavior. This interaction takes the form of an asymmetric Bertrand game, with the degree of asymmetry determined by differences in consumer valuation across firms. The overall equilibrium is obtained by solving these subgames pointwise and aggregating across consumers.

This formulation is consistent with the industrial organization literature on spatial price discrimination and personalized pricing (Thisse & Vives, 1988; Stole, 2007; Rhodes & Zhou, 2024), and it provides a tractable way to incorporate regulatory constraints into the pricing problem. Each regulatory instrument modifies the firm's optimization problem at the level of

individual consumers, allowing the analysis to characterize equilibrium behavior under different policy regimes.

b. Baseline Equilibrium Under Personalized Pricing

Consider first the unregulated pricing game. Firms simultaneously choose pricing schedules $p_i(x)$, and for each consumer $x \in [0,1]$, demand is determined by utility maximization. Let $m(x) = p_i(x) - c$ denote the markup.

For a given consumer *x*, firm *A* obtains demand if and only if

$$v - tx - p_A(x) \geq v - t\,(1 - x) - p_B(x),$$

which simplifies to

$$p_A(x) \leq p_B + \Delta(x),$$

where

$$\Delta(x) = t(1 - 2x).$$

For $x < ½$, firm *A* has a positive advantage, while for $x > ½$, firm *B* does. The pricing interaction for each consumer can therefore be interpreted as a localized Bertrand game with asymmetrical advantages. Fix $x < 1/2$. Firm *A* solves

$$\max_{p_{A(x)}} (p_A(x) - c)$$

subject to the constraint that the consumer does not switch:

$$p_A(x) \leq p_B + \Delta(x).$$

Firm *B*, if it were to win the consumer, would earn $p_B(x) - c$ and therefore cannot profitably set a price below marginal cost. This implies that in equilibrium $p_B(x) \geq c$, and the best response of firm *B* $p_B(x) = c$.

Substituting this into firm *A*'s constraint yields

$$p_A(x) \leq c + \Delta(x).$$

Proposition 1 (Personalized Pricing Equilibrium)

For each consumer *x*, the equilibrium prices are given by
$p_i(x) = c + |\Delta(x)|$ for the advantaged firm, $p_j(x) = c$ for the disadvantaged firm.

Firm profits are

$$\pi_i = \frac{t}{4}.$$

This equilibrium reflects the fact that the advantaged firm sets the highest price consistent with retaining the consumer, while the disadvantaged firm prices at marginal cost.

See Appendix A.1 for proof.

This equilibrium reflects the fundamental structure of competition under personalized pricing. Each consumer is modeled as a player in a localized Bertrand game in which firms compete based on their relative advantage. The advantaged firm extracts all surplus consistent with keeping the consumer, while the disadvantaged firm prices at cost.

The key economic force is the heterogeneity in competitive intensity across consumers. For infra-marginal consumers, where |Δ(x)|is large, firms retain pricing power. For marginal consumers, where $|\Delta(x)|$ is small, competition becomes extremely intense, driving prices toward marginal cost. The aggregation of these local interactions leads to lower profits relative to uniform pricing, illustrating how increased informational precision can intensify competition in oligopoly.

c. Constraint-Based Regulation: Guardrails

We now consider a regulatory regime in which the regulator imposes direct constraints on pricing behavior. These interventions resemble traditional rule-based or command-and-control forms of regulation, in which firms face ex ante constraints on admissible behavior (Kaplow, 1992; Shavell, 2004). Specifically, firms are restricted in the extent to which they can translate their competitive advantage into markups. This type of intervention captures policies that impose caps on price differentials, limit the use of personalized pricing, or otherwise restrict price dispersion.

Formally, suppose that for each consumer x, firms must satisfy the constraint

$$p_i(x) - c \leq \kappa|\Delta(x)|, \qquad \kappa \in [0,1].$$

This restriction can be interpreted as limiting the fraction of the firm's competitive advantage that can be monetized through pricing. When $\kappa = 1$, the constraint is non-binding, and the unregulated equilibrium is recovered. When $\kappa = 0$, firms are forced to price at marginal cost for all consumers.

As in the baseline model, the pricing interaction decomposes into independent subgames for each consumer. Fix $x < 1/2$, so that firm A is advantaged. In the absence of regulation, firm A's pricing is constrained by the competitive condition

$$p_A(x) \leq p_B(x) + \Delta(x).$$

Under guardrails, firm A faces an additional regulatory constraint.

$$p_A(x) - c \leq \kappa\Delta(x).$$

Combining these constraints, the firm's problem can be written as

$$\max_{p_A(x)} (p_A(x) - c)$$

subject to

$$p_A(x) \leq p_B(x) + \Delta(x), \quad p_A(x) \leq c + \kappa\Delta(x).$$

Given that firm B sets $p_B(x) = c$ in equilibrium, the first constraint reduces to $p_A(x) \leq p_B(x) + \Delta(x)$. Since $\kappa \leq 1$ the regulatory constraint is weakly tighter than the competitive constraint, and therefore determines the feasible set of prices.

Proposition 2 (Equilibrium Under Guardrails)

Under the constraint $p_A(x) - c \leq \kappa\Delta(x)$ equilibrium prices satisfy $p_i(x) = c + \kappa\Delta(x)$ for the advantaged firm, and $p_j(x) = c$ for the disadvantaged firm.

Firm profits are given by

$$\pi_i = \frac{\kappa t}{4}.$$

Since the firm's objective is linear in the markup, the regulatory constraint binds for all consumers.

See Appendix A.2 for proof.

Guardrails modify the pricing game by introducing a binding constraint that limits how much of the firm's competitive advantage can be converted into markup. Unlike the baseline case, where pricing is determined entirely by competitive constraints, equilibrium pricing is now governed by the regulatory cap.

A key feature of this regime is that the constraint binds uniformly across all consumers. Because the firm's objective is linear in the markup, it always sets the highest feasible price. As a result, the regulatory constraint directly determines equilibrium pricing, leaving no scope for interior solutions. This contrasts sharply with enforcement-based regimes, where firms optimally choose markups based on a tradeoff between profits and expected penalties.

The effect of guardrails is therefore a proportional compression of markups across all consumers. High-markup and low-markup consumers are affected in the same proportional way, preserving the relative structure of price dispersion while reducing its magnitude. This makes the policy predictable and easy to implement, but also rigid, as it does not allow firms to adjust pricing in response to local competitive conditions beyond the imposed cap.

From a strategic perspective, guardrails replace the endogenous pricing constraint arising from competition with an exogenous regulatory constraint. The pricing game remains a collection of localized Bertrand interactions, but the equilibrium outcome is no longer determined solely by competitive forces. Instead, regulation directly truncates the pricing strategies available to firms, reshaping equilibrium outcomes in a uniform manner.

d. Enforcement-Based Regulation: Audits and Penalties

We now consider a regulatory regime in which pricing behavior is not directly constrained, but firms face expected penalties that increase with the magnitude of markups. This type of intervention captures enforcement-based policies in which regulators monitor pricing behavior and impose sanctions when prices are deemed excessive.

Formally, suppose that for each consumer x, a firm chooses a markup $m = p_i(x) - c$ faces an expected penalty given by $e\gamma m^2$ where $\gamma > 0$ measures the severity of penalties and $e > 0$ captures the intensity of enforcement. The quadratic form reflects the idea that larger deviations from cost-based pricing are more likely to trigger scrutiny and sanctions, and that penalties increase disproportionately with the magnitude of the markup.

For a given consumer x, the advantaged firm solves

$$\max_{0 \le m \le |\Delta(x)|} [m - e\gamma m^2],$$

where $|\Delta(x)|$ represents the maximum feasible markup consistent with retaining the consumer in the absence of regulation.

Proposition 3 (Equilibrium Under Audits and Penalties)

Suppose that, for each consumer x, the advantaged firm chooses a markup $m = p_i(x) - c$ and faces the expected penalty $e\gamma m^2$ where $e > 0$ is enforcement intensity and $\gamma > 0$ is penalty severity. Then, for each consumer x, the firm solves

$$max_{0 \le m \le |\Delta(x)|} [m - e\gamma m^2].$$

Under the audit and penalty regime, the equilibrium markup is given by

$$m^*(x) = \min\left\{|\Delta(x)|, \frac{1}{2e\gamma}\right\}.$$

Equilibrium prices, therefore, satisfy

$$p_i(x) = c + m^*(x),$$

for the advantaged firm, while the disadvantaged firm prices at cost. Profit for firm *A* is

$$\pi_A = \int_0^{1/2} m^*(x)dx,$$

with an analogous expression for the other firm by symmetry.

The equilibrium exhibits two regimes. If the unconstrained optimum satisfies $\frac{1}{2e\gamma} \leq |\Delta x|$, the interior solution binds, and equilibrium markups are determined by the tradeoff between profits and expected penalties. If $\frac{1}{2e\gamma} > |\Delta x|$, the competitive constraint binds, and the outcome coincides with the unregulated equilibrium.

See Appendix A.3 for proof.

Audits and penalties affect pricing behavior by altering firms' incentives rather than imposing direct constraints. The firm now trades off the benefit of a higher markup against the expected enforcement costs. This generates an interior solution in which the firm voluntarily limits its pricing even when higher prices would be feasible in the absence of regulation.

A key feature of this regime is that it introduces an endogenous cap on markups. When enforcement is weak, the unconstrained optimum exceeds the competitive advantage, and the equilibrium coincides with the unregulated outcome. When enforcement is sufficiently strong, the optimal markup becomes independent of $\Delta(x)$ and is instead determined by regulatory parameters. This implies that the effect of regulation is heterogeneous across consumers: it binds more strongly for those with larger valuation differences.

Compared to guardrails, which impose a uniform proportional cap, audits generate a nonlinear compression of markups. High-markup consumers are affected more strongly than low-markup consumers, leading to a different pattern of price dispersion. This distinction has important welfare implications. Because audits and penalties primarily affect high-markup consumers while leaving low-markup interactions relatively unconstrained, they preserve stronger competitive pressure at the margin than uniform pricing constraints. Relative to guardrails, enforcement-based policies may therefore achieve reductions in price dispersion with smaller effects on competitive intensity.

From a strategic perspective, the audit regime preserves pricing flexibility while altering incentives. Firms retain the ability to respond to local competitive conditions, but they internalize the expected cost of regulatory scrutiny. This makes enforcement-based policies potentially more adaptable than rigid constraints, but also more dependent on institutional capacity and monitoring effectiveness.

e. Transparency and Enforcement

An important dimension of regulatory design concerns the role of transparency. While the previous subsections focus on constraints and incentives, transparency operates through a different channel by altering the informational environment in which firms and regulators interact.

To formalize this, suppose that the expected penalty parameter *e* depends on both the probability of detection and the observability of pricing behavior. Let

$$e = q(1 + \mu\tau),$$

where $q \in (0,1]$ represents the baseline probability of enforcement, $\tau \geq 0$ measures the degree of transparency, and $\mu > 0$ captures the sensitivity of enforcement to transparency.

Under this specification, transparency increases the effective probability that pricing behavior is detected and sanctioned. The firm's problem under audits and penalties therefore becomes

$$\max_{m \leq |\Delta(x)|} [m - q(1 + \mu\tau)\gamma m^2].$$

Proposition 4 (Transparency and Equilibrium Markups)

An increase in transparency τ weakly reduces equilibrium markups. In particular, the optimal markup is

$$m^*(x) = min\left\{|\Delta(x)|, \frac{1}{2q(1 + \mu\tau)\gamma}\right\},$$

which is decreasing whenever the enforcement constraint binds.

Since transparency increases the effective enforcement parameter *e*, the equilibrium markup decreases with the transparency parameter τ. Formally $\frac{\partial m^*(x)}{\partial \tau} < 0$.

See Appendix A.4 for proof.

Transparency affects pricing behavior indirectly by strengthening the effectiveness of enforcement. Unlike guardrails, it does not restrict the feasible set of prices, and unlike penalties,

it does not directly alter the payoff function. Instead, it increases the likelihood that pricing behavior will be detected and penalized.

This distinction is important because it implies that transparency is not a standalone regulatory tool. Its effectiveness depends on the presence of enforcement mechanisms. In the absence of audits or penalties, increasing transparency has no effect on firm behavior. When enforcement is present, however, transparency amplifies its impact by raising the expected cost of aggressive pricing.

From a strategic perspective, transparency increases the degree to which firms internalize regulatory constraints. As the probability of detection rises, firms optimally reduce markups even in the absence of explicit caps. This makes transparency a flexible policy instrument, but also one whose effectiveness depends critically on institutional capacity.

f. Mechanisms and Equilibrium Effects

The results derived above can be summarized as transformations of the equilibrium markup function. In the unregulated benchmark, the advantaged firm's markup is given by

$$m_0(x) = |\Delta(x)|.$$

Under guardrails, the markup function becomes

$$m_G(x) = \kappa|\Delta(x)|,$$

implying a proportional compression of markups across consumers. Under audits and penalties, the markup function becomes

$$m_A(x) = min\left\{|\Delta(x)|, \frac{1}{2e\gamma}\right\},$$

implying a nonlinear compression that binds primarily for high-markup consumers. Under transparency, the enforcement parameter becomes $e = q(1 + \mu\tau)$, so transparency reduces markups indirectly by strengthening enforcement. This comparison highlights the central mechanism of the paper: regulatory instruments differ not only in their objectives, but in how they transform the firm's pricing problem. Some policies change the feasible set, others change the payoff function, and still others change the informational environment in which pricing occurs.

The analysis in this section provides a formal representation of algorithmic pricing by modeling it as firms' ability to condition prices on individual consumer characteristics. In this sense, the pricing functions $p_i(x)$ capture the operational capability enabled by artificial intelligence and large-scale data systems. Rather than introducing a fundamentally new form of competition,

algorithmic pricing increases the granularity at which firms can implement pricing strategies, effectively decomposing the market into a continuum of localized interactions. The regulatory mechanisms analyzed above, therefore, correspond to different ways of constraining or shaping the use of this capability. By expressing each intervention as a transformation of the markup function, the framework links abstract policy tools—such as guardrails, audits, and transparency—to their concrete effects on algorithmic pricing behavior. This connection clarifies that the regulation of algorithmic pricing is, in essence, the regulation of how firms map consumer information into prices within a competitive environment.

VII. Welfare, Competition, and Regulatory Tradeoffs

The welfare implications of algorithmic pricing must be understood in light of the distinction between monopoly and oligopoly developed in Section 4.1 and the equilibrium characterization of pricing behavior derived in Section 6. Together, these results highlight that the effects of algorithmic pricing depend not only on the availability of information but also on how that information is used within a competitive environment.

For the purposes of the analysis, welfare is understood as the combination of consumer surplus, firm profits, and the distribution of prices across consumers. The paper does not adopt a fully specified social welfare function but instead focuses on how regulatory instruments affect three economically relevant dimensions: (i) the level of markups, (ii) the dispersion of prices across consumers, and (iii) the intensity of competition for marginal consumers. Throughout the analysis, price dispersion is treated as a tractable proxy for concerns about differential treatment and fairness, while recognizing that alternative normative interpretations are possible.

The analysis focuses primarily on static welfare effects: the distribution of prices, consumer surplus, and firm profits given the pricing technology and regulatory instrument in place. Dynamic questions concerning investment in data, prediction technologies, or information acquisition are important but lie outside the scope of the present model.

In the baseline equilibrium, markups are given by $m(x) = |\Delta(x)|$, reflecting the full extent of each firm's competitive advantage. This pricing structure generates two opposing effects. More specifically, welfare effects can be decomposed into extraction and competition effects. The extraction effect arises because firms can charge higher prices to consumers with strong preferences, reducing consumer surplus for inframarginal consumers. The competition effect arises because firms compete more aggressively for consumers located near the margin between firms, pushing prices closer to marginal cost. The overall welfare implications of personalized pricing, therefore, depend on the relative strength of these two forces. The resulting allocation is efficient in the sense that each consumer is served by the preferred firm, but it generates substantial price dispersion across consumers (Thisse & Vives, 1988; Stole, 2007).

Regulatory interventions modify this allocation by altering the mapping from competitive advantage to markups. Under guardrails, markups are given by $m(x) = \kappa|\Delta(x)|$, implying a proportional compression of prices across all consumers. This reduces price dispersion, limits surplus extraction, and preserves the relative structure of competition. Because the mapping remains linear, the allocation of consumers across firms is unaffected, and efficiency losses are limited. However, guardrails also reduce firms' ability to compete selectively for marginal consumers through targeted pricing. As a result, policies that reduce price dispersion may simultaneously reduce some of the competitive benefits generated by personalized pricing.

Under audits and penalties, markups are instead given by

$$m(x) = min\left\{|\Delta(x)|, \frac{1}{2e\gamma}\right\},$$

which introduces an endogenous cap on pricing. Unlike guardrails, this cap binds only for consumers with sufficiently large valuation differences. As a result, price compression is nonlinear: high-markup consumers are affected more strongly than low-markup consumers. This reduces dispersion while preserving competitive pressure at the margin, potentially leading to more favorable tradeoffs between efficiency and fairness. However, the effectiveness of this regime depends critically on enforcement intensity, as weak enforcement may leave equilibrium outcomes largely unchanged.

Transparency affects welfare indirectly by modifying the enforcement parameter $e = q(1 + \mu\tau)$. By increasing the probability that pricing behavior is detected, transparency strengthens the impact of enforcement-based policies and reduces equilibrium markups. Importantly, transparency has no direct effect in the absence of enforcement, highlighting its complementary nature. Its welfare implications therefore depend on the broader regulatory environment.

Taken together, these results show that regulatory instruments differ not only in their objectives but also in the mechanisms through which they operate. Guardrails impose exogenous constraints that compress prices uniformly, while audits and penalties induce endogenous adjustments that depend on enforcement parameters. Transparency, in turn, alters the informational environment that governs these adjustments. These differences imply that policies with similar goals, such as reducing price dispersion, may yield different equilibrium outcomes and welfare effects.

Proposition 5 (Informal Welfare Comparison)

Relative to guardrails, enforcement-based policies preserve greater competitive pressure for marginal consumers while achieving targeted reductions in high markups. Consequently, the welfare effects of audits and penalties depend less on uniform price compression and more on the distribution of competitive advantages across consumers.

This proposition follows directly from the nonlinear structure of equilibrium markups under enforcement-based regulation. Whereas guardrails proportionally compress all markups, audits primarily bind consumers with sufficiently large competitive advantages.

The welfare effects of regulatory interventions can be understood by examining how each policy alters the relationship between competitive advantage and equilibrium markups. Different mechanisms compress prices in different ways, thereby generating distinct effects on price dispersion, competition, consumer surplus, and firm profitability. Table 2 summarizes these comparative effects.

Table 2: Welfare and Competitive Effects of Regulatory Instruments

Policy Instrument	Markup Structure	Price Dispersion	Competition at the Margin	Consumer Surplus	Firm Profits	Main Welfare Tradeoff
No Regulation	Full advantage-based pricing	High	Strong	Mixed: lower for infra-marginal consumers, higher for marginal consumers	Lower than uniform pricing	Efficient allocation with substantial dispersion
Guardrails	Uniform proportional compression	Lower (uniformly)	Reduced	Higher for high-markup consumers; weaker marginal competition	Lower	Predictable but rigid
Audits & Penalties	Endogenous nonlinear cap	Lower (targeted)	Preserved	Improved for high-markup consumers while maintaining competition	Depends on enforcement	Flexible but enforcement-dependent
Transparency	Indirect reduction through enforcement	Conditionally lower	Depends on enforcement	Depends on monitoring effectiveness	Depends on enforcement	Complements incentive-based tools

A key implication of the comparison is that regulatory interventions should not be evaluated solely based on whether they reduce price dispersion. Policies differ in how they affect competitive pressure and the distribution of surplus across consumers and firms. As a result, welfare comparisons depend not only on the magnitude of price reductions, but also on how those reductions are generated within the pricing game. Consequently, evaluating algorithmic pricing regulation requires distinguishing between reductions in price dispersion arising from stronger competition and those arising from direct suppression of pricing flexibility.

The game-theoretical framework developed in this paper provides a unified framework for evaluating these trade-offs. By expressing regulatory interventions as transformations of the markup function, the analysis links policy tools directly to equilibrium outcomes. This approach clarifies that the regulation of algorithmic pricing is fundamentally about shaping how firms map consumer information into prices within a competitive environment.

A key implication is that regulatory design must be sensitive to market structure. In environments that approximate monopoly, the primary concern is surplus extraction, and stronger pricing constraints may be warranted (Varian, 1989). In oligopolistic markets, however, individualized pricing can intensify competition, particularly for marginal consumers. In such settings, overly restrictive policies may dampen competition and reduce consumer welfare. Effective regulation must therefore balance the goal of limiting harmful price discrimination with the need to preserve the benefits of competition and innovation.

This perspective shifts the focus of regulation from controlling algorithms to designing rules that shape the equilibrium use of information in pricing. More broadly, the analysis suggests that welfare evaluation in algorithmic pricing environments requires moving beyond static notions of price discrimination toward a framework centered on strategic interaction. The same informational capabilities that permit surplus extraction may also intensify competition, and regulatory interventions may alter both simultaneously. As a result, the welfare effects of regulation cannot be inferred solely from the reduction of price dispersion but must instead be evaluated in equilibrium.

VIII. Empirical Implications and Retrospective Evaluation

The framework developed in this paper yields several empirical implications that distinguish regulatory mechanisms by their effects on pricing behavior and competition. Because the model expresses regulation as transformations of firms' pricing problems, different policy instruments generate distinct observable patterns in markups, price dispersion, and competitive intensity.

The first implication concerns the effect of pricing constraints or guardrails. Under this regime, equilibrium markups are compressed proportionally across consumers. This implies that regulatory interventions based on direct pricing restrictions should produce relatively uniform

reductions in price dispersion. Empirically, one would expect a narrowing of the overall price distribution, with similar proportional effects across consumer segments.

In contrast, audits and penalties generate nonlinear effects. Because the equilibrium markup is determined by the minimum between the competitive advantage and the enforcement-based cap, enforcement primarily affects consumers associated with large markups while leaving low-markup interactions relatively unchanged. Empirically, this implies that enforcement-based policies should disproportionately compress the upper tail of the price distribution while preserving stronger competitive pressure for marginal consumers. Relative to guardrails, the reduction in observed dispersion should therefore be more concentrated among high-price transactions.

Transparency requirements generate a different prediction. In the framework, transparency affects outcomes indirectly by increasing the probability that pricing behavior is detected and sanctioned. As a result, the effects of transparency should depend critically on the broader enforcement environment. In settings with weak monitoring or limited enforcement capacity, transparency alone may have little measurable effect on prices. In settings with active enforcement, however, transparency may significantly amplify the impact of audits and penalties.

More broadly, the framework suggests that empirical evaluation of algorithmic pricing regulation should focus not only on average prices, but also on the structure and distribution of prices across consumers. Different regulatory tools may generate similar average price effects while producing substantially different patterns of competition and price dispersion. This distinction is particularly important for retrospective analysis, where identifying the mechanism through which a policy operates may be necessary to correctly interpret observed outcomes. This suggests that average prices alone may be insufficient for evaluating regulatory outcomes in algorithmic pricing environments.

The framework also suggests several directions for future empirical research. These include examining how algorithmic pricing affects price distribution across consumer groups, how enforcement intensity shapes pricing behavior, and how different regulatory instruments affect the upper tail of the price distribution. Such analyses may help clarify the conditions under which algorithmic pricing enhances competition versus those in which it reinforces concerns about differential treatment. These distinctions may also improve retrospective regulatory analysis by helping identify whether observed price changes result from intensified competition, stronger enforcement, or reduced pricing flexibility.

a. Mechanism-Based Retrospective Evaluation

The framework also has implications for retrospective regulatory analysis. A retrospective evaluation that focuses only on average prices may miss the mechanism through which a policy

affected market outcomes. For example, a decline in average prices may reflect stronger competition, direct suppression of markups, or increased enforcement. These mechanisms have different welfare implications even when average price effects appear similar.

A mechanism-based retrospective analysis should therefore examine the distribution of prices, markups, switching behavior, and competitive responses across market segments. This approach would allow evaluators to distinguish between price reductions driven by stronger competition and those driven by rigid constraints on pricing flexibility.

IX. Conclusion and Future Directions

The increasing use of artificial intelligence in pricing has raised important questions for policymakers, particularly regarding fairness, transparency, and the potential for consumer harm. At the same time, algorithmic pricing reflects a broader economic phenomenon: firms' ability to condition prices on increasingly detailed information about consumers. Understanding how to regulate this capability requires a framework that accounts not only for the availability of information but also for how firms use it in competitive environments.

This paper develops such a framework by modeling algorithmic pricing as the ability to implement consumer-specific pricing functions at scale and embedding this capability within a strategic pricing game. The analysis shows that equilibrium prices emerge from the interaction of firms competing for individual consumers, and that regulatory interventions affect outcomes by modifying this interaction. In this sense, the regulation of algorithmic pricing is fundamentally about how firms map consumer information into prices.

A central result of the paper is that the effects of individualized pricing depend critically on market structure. While in monopoly settings improved information facilitates surplus extraction, in oligopolistic environments it can intensify competition, particularly for marginal consumers. This distinction, well established in the literature, has important implications for policy. Regulatory approaches based on monopoly intuition may be misaligned with competitive markets, where restricting individualized pricing can reduce the benefits of competition.

The paper's contribution lies in showing how different regulatory instruments operate within a unified game-theoretic framework. By framing policies as modifications of the pricing game, the analysis distinguishes between constraint-based interventions, such as guardrails, and incentive-based mechanisms, such as audits and penalties, and examines the role of transparency in shaping enforcement. These instruments generate different transformations of the equilibrium markup function and therefore produce distinct patterns of price dispersion, competition, and welfare outcomes.

This mechanism-based perspective yields a central policy insight: regulatory tools that appear similar in their objectives can lead to different outcomes once firms respond strategically.

Effective regulation must therefore be evaluated in terms of its equilibrium effects rather than its direct constraints. In particular, policies that uniformly limit pricing may reduce harmful price dispersion but at the cost of dampening competitive forces, while incentive-based approaches may preserve flexibility but depend critically on enforcement capacity.

More broadly, the analysis suggests that regulation should focus on outcomes rather than technology. Algorithmic pricing does not fundamentally alter the nature of competition; it increases the precision with which firms can implement pricing strategies. As a result, the appropriate object of regulation is not the algorithm itself, but the pricing behavior and market outcomes that arise from its use.

The framework developed in this paper provides a foundation for further work on the regulation of AI-driven markets. One natural extension would relax the symmetric-information benchmark and allow firms to differ in their access to consumer data, predictive technologies, or algorithmic capabilities. Such an extension would examine whether algorithmic pricing intensifies competition when informational capabilities are broadly shared but reinforces market power when informational advantages are concentrated. Additional work could also examine dynamic pricing, endogenous data acquisition, and investment incentives in data-driven markets. More generally, future competition policy in AI-driven markets will increasingly depend on understanding not only how firms use information, but also how informational advantages emerge, evolve, and are regulated.

Appendix

A.1 Proof Proposition 1 (Personalized Pricing Equilibrium)

Consider $x < 1/2$, so that firm A has a competitive advantage $\Delta(x) > 0$. Given that firm B sets $p_B(x) = c$ to marginal cost, firm A solves

$$\max_{p_A(x) \le c + \Delta(x)} (p_A(x) - c).$$

The objective function is strictly increasing in p_A. Therefore, firm A sets the highest feasible price,

$$p_A(x) = c + \Delta(x).$$

Given this price, firm B cannot profitably deviate by undercutting, since doing so would require $p_B(x) < c$, yielding negative profit. Therefore, $p_B(x) = c$ is optimal.

By symmetry, for $x > 1/2$, firm B charges $c + |\Delta(x)|$, and firm A charges c.

Firm A's total profit is

$$\pi_A = \int_0^{1/2} (c + t(1 - 2x) - c)dx = \int_0^{1/2} t(1 - 2x)dx.$$

Evaluating the integral,

$$\pi_A = t[x - x^2]_0^{1/2} = t\left(\frac{1}{2} - \frac{1}{4}\right) = \frac{t}{4}.$$

Symmetry implies the same for firm B. ■

A.2 Proof Proposition 2 (Equilibrium Under Guardrails)

Fix $x < 1/2$, so that A is advantaged. Given that B sets $p_B(x) = c$, firm A solves

$$\max_{p_A \le c + \kappa\Delta x} (p_A(x) - c).$$

Since the objective function is strictly increasing in $p_A(x)$, the constraint binds, yielding

$$p_A(x) = c + \kappa\Delta(x).$$

Given this price, firm B cannot profitably deviate by undercutting it unless prices below marginal cost, which would yield negative profit. Therefore, $p_B(x) = c$ is optimal.

By symmetry, for $x > 1/2$ firm B sets $p_B(x) = c + \kappa|\Delta(x)|$, and firm A sets $p_A(x) = c$.

Firm A's total profit is

$$\pi_A = \int_0^{1/2} \kappa t(1 - 2x)dx = \kappa \int_0^{1/2} t(1 - 2x)dx.$$

Evaluating the integral yields

$$\pi_A = \frac{\kappa t}{4}.$$

Since the objective function is linear, the upper bound always binds.

Symmetry implies the same for firm B. ■

A.3 Proof Proposition 3 (Equilibrium Under Audits and Penalties)

Fix a consumer-specific contest with a match advantage $|\Delta(x)|$. Let $m = p_i(x) - c$ denote the markup chosen by the advantaged firm. This firm solves

For a given consumer x, the advantaged firm maximizes the objective function

$$\max_{0 \le m \le |\Delta(x)|} \phi(m), \qquad \text{where } \phi(m) = m - e\gamma m^2$$

The first-order condition is

$$\frac{d\phi}{dm} = 1 - 2e\gamma m = 0,$$

which yields the unconstrained optimum

$$m^{uc} = \frac{1}{2e\gamma}.$$

This solution is feasible only if $\frac{1}{2e\gamma} \le |\Delta(x)|$. Otherwise, the competitive constraint binds, and the firm sets $m(x) = |\Delta(x)|$.

The second-order condition is

$$\frac{d^2\phi}{dm^2} = -2e\gamma < 0,$$

so the objective function is strictly concave, and the first-order condition characterizes the unique maximum.

Imposing the feasibility constraint $m \leq |\Delta(x)|$, the optimal markup is

$$m^*(x) = min\left\{|\Delta(x)|, \frac{1}{2e\gamma}\right\}.$$

If $\frac{1}{2e\gamma} \geq |\Delta(x)|$, the constraint binds, and the outcome coincides with the unregulated equilibrium.

Substituting into the pricing rule yields

$$p_i(x) = c + m^*(x).$$

Firm profits are obtained by integrating the equilibrium markup over the set of consumers for which the firm is advantaged. For firm *A*, these yields

$$\pi_A = \int_0^{1/2} m^*(x)dx.$$

Symmetry implies the same for firm *B*. ■

A.4 Proof Proposition 4 (Transparency and Equilibrium Markups)

Fix a consumer-specific contest with a match advantage $\Delta(x)$. Let $m = p_i(x) - c$ denote the markup. The firm solves

$\max_{0 \leq m \leq |\Delta(x)|} \phi(m)$, where $\phi(m) = m - e\gamma m^2$.

First, compute the derivatives:

$$\phi'(m) = 1 - 2e\gamma m,$$

$$\phi''(m) = -2e\gamma.$$

Since $e > 0$ and $\gamma > 0$, we have $\phi''(m) < 0$ so $\phi \quad (m)$ is strictly concave. Therefore, any solution to the first-order condition is the unique unconstrained maximizer.

Set the first-order condition equal to zero:

$$1 - 2e\gamma m = 0$$

Solving gives

$$m^{uc} = \frac{1}{2e\gamma}.$$

This is the unconstrained optimum. However, the pricing problem imposes the constraint

$$m \leq |\Delta(x)|.$$

Thus, the constrained optimum is:

$$m^*(x) = \min\left\{|\Delta(x)|, \frac{1}{2e\gamma}\right\}.$$

Under transparency, the enforcement parameter satisfies:

$$e = q(1 + \mu\tau).$$

Substituting into the optimal markup yields:

$$m^*(x) = min\left\{|\Delta(x)|, \frac{1}{2q(1 + \mu\tau)\gamma}\right\}.$$

Since *e* is increasing in τ, the unconstrained optimum $\frac{1}{2e\gamma}$ is decreasing in τ. Therefore,

$$\frac{\partial}{\partial \tau}\left(\frac{1}{2q(1 + \mu\tau)\gamma}\right) < 0,$$

Whenever the interior solution binds. ■

References

Ayres, I., & Braithwaite, J. (1992). Responsive regulation: Transcending the deregulation debate. Oxford University Press.

Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2021). Tacit collusion on steroids: The potential risks for competition resulting from the use of algorithm technology by companies. *Sustainability, 13*(2), 951. https://doi.org/10.3390/su13020951

Corts, K. S. (1998). Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment. *The RAND Journal of Economics*, *29*(2), 306–323. https://doi.org/10.2307/2555890

Federal Trade Commission. (2021, April 19). *Aiming for truth, fairness, and equity in your company's use of AI*. https://www.ftc.gov/ai

Fudenberg, D., & Tirole, J. (2000). Customer Poaching and Brand Switching. *The RAND Journal of Economics, 31*(4), 634–657. https://doi.org/10.2307/2696352

Kaplow, L. (1992). Rules versus standards: An economic analysis. *Duke Law Journal, 42*(3), 557–629. https://scholarship.law.duke.edu/dlj/vol42/iss3/2/

OECD. (2018). *Personalised pricing in the digital era*. OECD Publishing. https://www.oecd.org/competition/personalised-pricing-in-the-digital-era.htm

OECD. (2022). *Algorithms and collusion: Competition policy in the digital age*. OECD Publishing. https://www.oecd.org/content/dam/oecd/en/publications/reports/2017/05/algorithms-and-collusion-competition-policy-in-the-digital-age_02371a73/258dcb14-en.pdf

Posner, R. A. (1975). The social costs of monopoly and regulation. *Journal of Political Economy, 83*(4), 807–827. https://doi.org/10.1086/260357

Rhodes, A., & Zhou, J. (2024). Personalized pricing and competition. *American Economic Review* 114(7):2141–70. https://doi.org/10.1257/aer.20221524

Shavell, S. (2004). *Foundations of economic analysis of law*. Harvard University Press.

Stigler, G. J. (1964). Public regulation of the securities markets. *The Journal of Business, 37*(2), 117–142. https://doi.org/10.1086/294677

Stole, L. A. (2007). Price discrimination and competition. In M. Armstrong & R. Porter (Eds.), *Handbook of industrial organization* (Vol. 3, pp. 2221–2299). Elsevier. https://doi.org/10.1016/S1573-448X(06)03034-2

The Economist. (2019, January 17). *The rise of personalised pricing*. https://www.economist.com

Thisse, J.-F., & Vives, X. (1988). On the strategic choice of spatial price policy. *The American Economic Review, 78*(1), 122–137.

Varian, H. R. (1989). Price discrimination. In R. Schmalensee & R. Willig (Eds.), *Handbook of industrial organization* (Vol. 1, pp. 597–654). Elsevier.

Zhao, Z (2023). Algorithmic Personalized Pricing with the Right to Explanation, *Journal of Competition Law & Economics*, Volume 19(3), 367–396. https://doi.org/10.1093/joclec/nhad008