



Why Does the Perfectly Discriminating Monopolist Look Like Perfect Competition and Why Doesn't It? A Pedagogical Treatment for Intermediate Microeconomics

Intermediate microeconomics students comfortable with perfect competition and monopoly often struggle with the perfectly discriminating monopolist, a firm that charges every consumer exactly their willingness to pay. Such a firm produces the same total output as perfect competition yet captures all consumer surplus, an efficient outcome that violates any intuition about fairness. In total surplus, perfect competition and perfect discrimination tie and both exceed standard monopoly. In consumer surplus the order reverses and we get perfect competition, standard monopoly, and perfect discrimination. This paper offers a classroom-ready treatment: a five-consumer example, a comparison table, a data-driven pricing application, and a 60-minute lesson plan.

Carlos Liard-Muriente[†]

[†]Central Connecticut State University

1. Introduction: The Paradox That Teaches

Ask a room of intermediate microeconomics students whether a monopolist is bad for society. Nearly every hand goes up. Now tell them that under certain conditions, a monopolist produces exactly the same total output as a perfectly competitive market and is therefore just as allocatively efficient. The room goes quiet.

This is the paradox of the perfectly discriminating monopolist, and it is one of the most instructive moments in all of microeconomic theory. Throughout this paper, the terms perfectly discriminating monopolist, first-degree price discrimination, and perfect price discrimination refer to the same thing. That is, a firm that charges each buyer exactly their reservation price. Students entering intermediate microeconomics already know the standard monopoly story: price is set above marginal cost leads to reduced output, which leads to deadweight loss, which makes consumers worse off. The perfectly discriminating monopolist disrupts every part of that story. The result is not wrong (it is exactly right), but its conditions are so demanding, and its distributional implications so stark, that understanding precisely where the equivalence with perfect competition holds and where it breaks is a lesson in careful economic reasoning that transfers well beyond the classroom. This paper's central pedagogical claim is that the gap between what perfect discrimination requires in theory and what firms can achieve in practice is not a flaw in the theory, but its greatest teaching asset. This is because it forces students to distinguish efficiency from feasibility, and technical optimality from institutional reality.

The paradox has acquired new urgency as digital platforms increasingly use large-scale behavioral data to set prices that vary across individual customers. Machine learning is the tool that converts that data into individualized predictions of willingness to pay, but the economic force at work is the data itself, not the algorithm. Even so, these firms remain far from true first-degree price discrimination because of informational, strategic, legal, and arbitrage constraints. Students who have just learned the model naturally ask: how close are these firms to the theoretical benchmark, and does the remaining gap matter for consumer welfare? Answering that question requires exactly the disciplined welfare analysis this module develops. This is why data-driven pricing is a pedagogical opportunity, and not merely a current-event reference.

This module draws on and contributes to existing economics education literature. This literature documents persistent student difficulty with surplus analysis and with the relationship between efficiency and equity (Walstad & Allgood, 1999), and it finds that the most durable learning comes from working through a small number of concepts deeply rather than covering many topics superficially (Hansen, Salemi, & Siegfried, 2002). It also establishes design principles this paper applies directly: that applied, real-world connections improve retention of economic principles (Becker & Watts, 2001), that having students generate answers before instruction produces more durable understanding (Kapur, 2016), and that the conflation of efficiency with desirability is among the most persistent errors in student reasoning (Feldman & Serrano, 2006). The comparison developed here is precisely the kind of rich conceptual node these authors have in mind: it integrates consumer and producer surplus, deadweight loss, market structure, and distributional considerations in a single teachable example. The theoretical foundations of price discrimination, from the classical treatments (Pigou, 1920; Robinson, 1933) to recent work on personalized pricing (Shiller, 2020; Dubé & Misra, 2023; Rhodes & Zhou, 2024), are introduced where the module uses them: in Sections 2 and 5.

This paper integrates four elements that instructors usually assemble separately (a welfare-comparison framework, a five-consumer example, a data-driven pricing application, and a 60-minute lesson plan), into one self-contained module, and it makes the underlying design principle explicit. The comparison is most durable when students meet the efficiency-

equity tension as a genuine surprise rather than a theorem to memorize. The remainder of the paper proceeds as follows: Section 2 reviews the three market structures and the general requirements for price discrimination; Section 3 presents a structured comparison table; Section 4 develops the five-consumer numerical example with explicit teaching points; Section 5 addresses the information problem and the data-driven pricing application; Section 6 examines the welfare rankings and what they teach about efficiency versus equity; and Section 7 offers discussion questions, a classroom experiment, and the lesson plan.

A. Learning Outcomes

After completing this module, students should be able to:

- Explain why the perfectly discriminating monopolist and perfect competition produce the same total output and welfare, despite completely different distributional outcomes.
- Distinguish allocative efficiency from distributional fairness and apply both criteria to evaluate a market outcome.
- Construct and compare total welfare, consumer surplus, and deadweight loss across the three market structures using a numerical example.
- Evaluate why data-driven personalized pricing approaches but do not achieve true first-degree discrimination, identifying which institutional and informational barriers remain.
- Apply the efficiency-equity distinction to a real policy question, such as pharmaceutical pricing or platform data practices.

B. Classroom Implementation

Throughout several semesters, students consistently identify the efficiency-equity distinction as the module's most surprising element. An outcome can be simultaneously efficient and profoundly unfair. Two moments give them the most trouble: accepting that the perfectly discriminating outcome is efficient despite the complete extraction of consumer surplus, and understanding why consumers fare worse under a smarter monopolist than under a standard one. Both are best addressed through the numerical example in Section 4, where students compute the figures themselves before seeing the comparison table. A useful formative assessment after Section 4 is one sentence answering, "This outcome is efficient. Is it good? Why or why not?" The answers reveal at once whether students have internalized the distinction or are conflating the two.

2. Three Market Structures: A Concise Review

Students typically arrive with a working knowledge of perfect competition and standard monopoly but little or no exposure to first-degree price discrimination. This section provides the minimum review needed to make the differences precise before the comparison table and numerical example.

Price discrimination of any degree requires three conditions. First, the firm must have market power; a price-taking firm cannot charge any buyer more than the market price. Second, the firm must be able to identify, or sort, buyers by their willingness to pay. Third, the firm must be able to prevent resale, so that a buyer charged a low price cannot undercut the firm by reselling to a buyer charged a high price.

First-degree price discrimination is the limiting case in which the second condition holds perfectly. Here, the firm knows each individual buyer's reservation price. Keeping these three requirements in view clarifies why first-degree discrimination is so demanding in practice, a point Section 5 develops.

This paper compares first-degree discrimination against perfect competition and uniform monopoly, and it deliberately brackets second-degree discrimination, the menu-based instruments such as quantity discounts, two-part tariffs, and versioning. In practice, firms often reach for those instruments first, but including them would blur the sharp efficiency-equity contrast that is this module's teaching payoff, and the 60-minute lesson plan is already full. Instructors who wish to situate the material can present second-degree discrimination as an intermediate case lying between uniform pricing and the first-degree benchmark.

A. Perfect Competition

In a perfectly competitive market, price is driven to marginal cost by the presence of many buyers and many sellers, none of whom can influence the market price; each firm takes the price as given. Free entry and exit is a separate force, driving economic profit to zero in the long run. Every consumer whose willingness to pay exceeds marginal cost is served and retains the surplus above the market price. Firms earn normal returns. Total welfare, the sum of consumer and producer surplus, is maximized with no deadweight loss. Perfect competition is the welfare benchmark against which other market structures are evaluated.

B. Standard (Uniform-Price) Monopoly

A standard monopolist sets a single price above marginal cost to maximize profit. Buyers whose willingness to pay exceeds marginal cost but falls below the monopoly price are excluded; they would generate positive net value but are priced out. This exclusion creates deadweight loss; that is, economic value that is neither captured by the firm nor enjoyed by consumers. The standard welfare case against monopoly rests on this deadweight loss, not on the firm's profitability as such.

C. The Perfectly Discriminating Monopolist

Now suppose a monopolist knows exactly what each consumer is willing to pay and can charge each one that exact amount. This is first-degree or perfect price discrimination (Pigou, 1920; Robinson, 1933). Because the firm earns a positive margin on every consumer whose willingness to pay exceeds marginal cost, it will serve all such consumers, the same set served under perfect competition. Total output is therefore identical to the competitive outcome, and there is no deadweight loss. The allocation is efficient.

The crucial difference is distributional. Here, the firm captures all consumer surplus. Every buyer pays exactly their maximum willingness to pay and retains nothing. The efficiency gain relative to standard monopoly is real, but it is achieved through a complete redistribution of surplus from consumers to the firm. This is the central teaching tension, and it is the engine of the paper.

3. Structured Comparison of the Three Market Structures

Table 1 summarizes the three market structures across eight dimensions. It is designed to be used directly in the classroom, either distributed as a handout or displayed on screen, and works best when introduced after students have worked through the numerical example in Section 4. A productive active-learning technique is to give students a blank version of the

table and ask them to complete it individually before the class reviews each row together.

Table 1. Perfect Competition, the Perfectly Discriminating Monopolist, and Standard Monopoly Compared

Dimension	Perfect Competition	Perfectly Discriminating Monopolist	Standard Monopoly
Output / Efficiency	Efficient quantity; price equals marginal cost for all buyers	Efficient allocation; serves all consumers whose willingness to pay is at least marginal cost; identical total output to perfect competition	Inefficient; buyers whose willingness to pay exceeds marginal cost are excluded at the monopoly price
Prices	One price, equal to marginal cost, for all buyers	A different price for every buyer, equal to that buyer's exact willingness to pay	One price above marginal cost for all buyers
Consumer Surplus	High; consumers retain all surplus above price	Zero; the firm extracts the full surplus from every buyer	Positive but reduced; some buyers pay well below their maximum willingness to pay
Producer Surplus	Normal returns (zero economic profit under free entry)	Maximized; captures the entire social surplus	Positive economic profit, but less than under perfect discrimination
Total Welfare	Maximum (consumer surplus plus producer surplus, no deadweight loss)	Maximum, but only under the conditions noted in the footnote to Section 4.C	Below maximum; deadweight loss from excluded buyers
Information and Identification	No individual identification; price emerges from aggregate market behavior	Complete individual identification; the firm must know each buyer's exact willingness to pay and prevent arbitrage between buyers	No individual identification; only aggregate demand matters
Real-World Feasibility	Approximated in commodity, financial, and spot markets	Partially approximated through data-driven personalized pricing in digital markets	The most common market structure observed in practice
Key Teaching Point	Benchmark for efficiency and consumer welfare	Efficiency does not equal equity: identical output to perfect competition, but the opposite distributional outcome	Benchmark for deadweight loss; familiar from introductory courses

Note: Total welfare is defined as consumer surplus plus producer surplus. Under free entry, producer surplus in perfect competition corresponds to zero economic profit.

The table is organized so that the first rows establish the basic welfare results, and the final rows prepare the ground for Section 5. Instructors should spend additional time on those final rows.

The “Total Welfare” row makes the efficiency equivalence explicit: perfect competition equals perfect discrimination in total surplus, both strictly greater than standard monopoly. Students who have been taught that monopoly is bad for welfare often resist this result. The table helps by showing that the welfare comparison depends on which dimension of welfare is being examined, a point that leads naturally into Section 6.

The “Information and Identification” row is the pivot of the entire paper. All three models are conventionally taught under an assumption of perfect information. The perfectly competitive and monopoly models are no exception. What distinguishes first-degree price discrimination is therefore not information in the abstract but the firm’s ability to identify each individual buyer’s reservation price and to prevent that buyer from reselling to others. Perfect competition requires no such individual identification, since prices emerge from aggregate market behavior, and standard monopoly requires only knowledge of aggregate demand. The perfectly discriminating monopolist alone must identify and price each buyer individually. This asymmetry in what firms must identify, not in whether information exists, is what makes first-degree discrimination theoretically elegant but practically demanding, and it is what makes data-driven personalized pricing interesting to analyze through this framework.

4. A Numerical Example

The numerical example is the paper’s primary teaching tool. It can be completed in a single class session in approximately 15 minutes as an instructor-led demonstration. Instructors who prefer to have students work through it in pairs should assign it as pre-class preparation, since the paired version takes roughly 25 minutes and would otherwise crowd the 60-minute lesson plan in Section 7.D. The numbers are deliberately simple so that students can focus on the economic logic rather than the arithmetic.

Consider a market with five consumers whose willingness to pay is \$10, \$8, \$6, \$4, and \$2. The firm has constant marginal cost of \$3. We work through three market structures in sequence.

A. Perfect Competition

At a price equal to marginal cost of \$3, consumers with willingness to pay of at least \$3 buy. That is the \$10, \$8, \$6, and \$4 consumers, so four units are sold. Consumer surplus is $(\$10 - \$3) + (\$8 - \$3) + (\$6 - \$3) + (\$4 - \$3) = \$7 + \$5 + \$3 + \$1 = \$16$. Firm economic profit is \$0, because price equals marginal cost. Total welfare is \$16.

B. Standard Monopoly

The monopolist selects the price that maximizes profit. At a price of \$6, three consumers buy (willingness to pay of \$10, \$8, and \$6). Revenue = $3 \times \$6 = \$18$. Cost = $3 \times \$3 = \$9$. Profit = $\$18 - \$9 = \$9$. Consumer surplus = $(\$10 - \$6) + (\$8 - \$6) + (\$6 - \$6) = \$4 + \$2 + \$0 = \$6$. Total welfare = $\$9 + \$6 = \$15$. Deadweight loss = $\$16 - \$15 = \$1$, representing the \$4 consumer who values the good above marginal cost but is priced out¹.

¹ In the familiar supply-and-demand diagram, deadweight loss appears as a triangle. Here, with discrete unit-sized consumers, that triangle is the single unit of surplus lost from the excluded \$4 consumer; the discrete calculation and the graphical area are the same object viewed at different resolutions.

C. Perfect Price Discrimination

The perfectly discriminating monopolist charges each consumer their exact willingness to pay. The \$2 consumer does not buy, because that willingness to pay is below marginal cost. All others buy at their willingness to pay: \$10, \$8, \$6, and \$4. Firm revenue = $\$10 + \$8 + \$6 + \$4 = \$28$. Total cost = $4 \times \$3 = \$12$. Firm profit = $\$28 - \$12 = \$16$. Consumer surplus = \$0. Total welfare = $\$16 + \$0 = \$16$, identical to perfect competition, with zero deadweight loss².

D. Teaching Point 1: The Efficiency-Equity Distinction

Perfect discrimination and perfect competition produce the same total welfare (\$16) and the same deadweight loss (\$0). But the distribution is completely reversed. Under perfect competition, consumers keep \$16 and the firm earns \$0 of economic profit; under perfect discrimination, the firm takes \$16 and consumers keep \$0. This is not a distributional footnote, it is a complete inversion of who benefits from market activity. Instructors should pause here and ask: 'Is this outcome desirable?' The question has no clear answer, which is exactly the point. Students who have been taught that efficiency is the criterion for evaluating markets are forced to confront the fact that an efficient allocation can be profoundly unfair.

This moment connects directly to a core learning goal in intermediate microeconomics: Students should leave the course understanding that economic efficiency is a well-defined technical concept, not a normative judgment. Feldman and Serrano (2006) argue that distinguishing positive from normative analysis is among the most durable lessons students can take from intermediate economics. The example makes this distinction concrete and memorable.

E. Teaching Point 2: Perfect Discrimination versus Standard Monopoly

Comparing the standard monopoly and perfect discrimination rows produces a second surprise: perfect discrimination is better for total welfare (no deadweight loss) but worse for consumers (zero surplus versus \$6 under standard monopoly). Students often find this counterintuitive. The firm that discriminates more aggressively produces a better aggregate outcome but leaves individual buyers worse off than the less sophisticated monopolist. This is a productive entry point for discussing why consumer welfare and total welfare can diverge, and why antitrust policy typically cares about both.

5. The Information Problem and the Data-Driven Pricing Application

A. Why Perfect Discrimination Is Nearly Impossible, and Whether Data Can Change This

For a firm to practice first-degree price discrimination, it must satisfy three conditions: know each consumer's exact reservation price, prevent resale between buyers, and commit credibly to individualized offers. These requirements historically made perfect discrimination almost entirely theoretical, and they remain demanding even in the digital era. This is what makes data-driven pricing a productive teaching application rather than a settled policy case. The goal of this section is to give students tools for evaluating how close any real firm actually comes to the benchmark.

² This welfare equivalence between perfect competition and perfect discrimination holds under the conditions the example satisfies by construction: constant marginal cost, no fixed costs, and a marginal consumer whose willingness to pay equals marginal cost, so that no surplus is lost at the margin. Instructors who adapt the example with increasing marginal cost, capacity constraints, or fixed costs will find that the equivalence no longer holds exactly.

Large-scale behavioral data and machine learning have partially changed the first condition. A platform that tracks browsing history, purchase patterns, and demographic information can build predictive models of individual willingness to pay without ever asking directly. Shiller (2020) demonstrated this using web-browsing data alone. Predictive models of individual willingness to pay improved substantially on uniform pricing for streaming subscriptions, with simulated profit gains becoming an order of magnitude larger when using behavioral data rather than using demographics alone. Dubé and Misra (2023) implemented machine-learning-based personalized pricing in a field experiment and documented a measurable transfer of surplus from consumers to the firm. This is consistent with moving partway toward the perfectly discriminating outcome while remaining short of it. The key teaching point is not the magnitude of any profit gain, but which specific barriers (informational, strategic, legal, and institutional), prevent firms from closing the remaining gap.

Two familiar platforms make the point concrete. Ride-hailing services, such as Uber, and food-delivery services, such as DoorDash, set prices that vary by time, location, and individual app history, and they are frequently cited as firms approaching individualized pricing. Yet scored against the conditions below, they remain far from first-degree discrimination. They can identify users and predict demand reasonably well, but they cannot prevent arbitrage perfectly, cannot always commit to their personalized prices, and operate under tightening legal constraints. Running these everyday examples through the checklist is more instructive than asserting that such firms have arrived at perfect discrimination because the exercise shows students exactly which conditions bind.

These findings make data-driven pricing a natural classroom application of the theory. Students see the theory they have just learned operating in markets they use every day, a connection that Becker and Watts (2001) identify as central to effective economics pedagogy.

B. A Checklist for Near-Perfect Discrimination: A Classroom Exercise

A productive classroom activity is to evaluate a real firm against the conditions required for perfect discrimination. Ask student groups to assess a specific digital platform using the following five conditions:

1. Does the firm track individual identity reliably across transactions?
2. Can it predict your willingness to pay from behavioral data?
3. Can it prevent resale, that is, arbitrage between buyers?
4. Can it commit not to undercut its own personalized prices later?
5. Is personalized pricing legally permissible in its market?

Condition 4 deserves a brief word, because it is the subtlest and the one which students most often overlook. A firm that charges high-value consumers their full reservation price today gives itself a strong incentive to cut prices tomorrow to reach the buyers it did not serve. Sophisticated buyers anticipate this and wait, which unravels the discrimination before it begins. This is the Coasian commitment problem, and it explains why the ability to commit, not merely the ability to measure, is central to first-degree pricing.

Most real firms score two or three out of five, placing them considerably closer to uniform monopoly than to true perfect discrimination, even with sophisticated data analysis. Conditions 3, 4, and 5 are structural barriers that better data does not remove. This observation helps students see that the gap between data-driven pricing and true perfect discrimination

is not narrowing uniformly. Emerging data-protection frameworks, such as the European Union's General Data Protection Regulation and the United States Federal Trade Commission's commercial-surveillance rulemaking, are making conditions 2 and 5 more demanding, not less. This is a useful anchor for classroom discussion of how institutional context shapes the feasibility of personalized pricing.

This checklist reinforces one of the paper's central lessons: The gap between theoretical perfect discrimination and real-world pricing is not primarily a data problem, it is an institutional one. That distinction is the transferable insight students carry beyond this particular topic.

C. A Note on Competition

The analysis so far, like most intermediate treatments, examines first-degree price discrimination as a single-firm phenomenon, and the welfare calculus can change under competition. Rhodes and Zhou (2024) study personalized pricing in a general oligopoly and show that its effect depends on the degree of market coverage. Here, when coverage is high, because costs are low or many firms compete, personalized pricing intensifies competition and can leave consumers better off than uniform pricing would, reversing the surplus-extraction intuition that holds for a single seller; when coverage is low, the opposite holds. The lesson for students is that the monopoly benchmark developed here is exactly that, a benchmark. Thus, introducing competing sellers can reduce, or even reverse, the harm to consumers.

6. Teaching the Welfare Rankings

A. The Three-Way Hierarchy

Drawing together the analysis, we can state two welfare rankings clearly:

In total surplus: Perfect Competition = Perfect Discrimination > Standard Monopoly
In consumer surplus: Perfect Competition > Standard Monopoly > Perfect Discrimination

These rankings are worth writing on the board and asking students to derive before presenting them. Students who can reproduce both rankings and explain why they differ have understood the core lesson of the paper.

The rankings reveal the fundamental tension in price-discrimination policy. A regulator focused on allocative efficiency might view near-perfect discrimination favorably relative to standard monopoly, while one focused on consumer welfare would view it as the worst of the three. Most antitrust frameworks weigh both, which is why the policy question stays contested.

B. A Teaching Note on Equity and Income Distribution

An important caveat for instructors teaching more advanced students. The efficiency case for perfect discrimination assumes that a dollar of surplus is equally valuable regardless of who holds it. If consumers tend to have lower incomes than shareholders, a social welfare function that weights poorer households more heavily may rank perfect discrimination below even standard monopoly, despite its higher total surplus. Dworczak, Kominers, and Akbarpour (2021) formalize this argument.

Instructors need not present the formal model, but the underlying intuition is valuable for discussion. Ask students, "Suppose the consumers in our example are workers and the firm's shareholders are wealthy. Does the efficiency argument still feel convincing?" This question bridges the gap between technical welfare analysis and the normative questions students

actually want to debate, directly instantiating the positive-normative distinction from Teaching Point 1 in Section 4.

It also helps to give students two words for what they are debating. Procedural fairness asks whether the process was legitimate: Did the firm deceive anyone? Was participation voluntary? Distributive fairness asks whether the resulting distribution is acceptable, regardless of the process. Perfect discrimination can be procedurally unobjectionable and distributively troubling at the same time, and naming the two ideas lets students say precisely which one is driving their reaction.

7. Discussion Questions, Classroom Experiment, and Lesson Plan

A. Discussion Questions

The following questions are designed to deepen engagement with the material. They are suitable for in-class discussion, short written responses, or problem set assignments.

1. In the numerical example, is the outcome under perfect discrimination fair? What conception of fairness are you applying, procedural or distributive, and does economics have anything to say about it?
2. A streaming service uses your viewing history to offer you a personalized subscription price. Score it against the five-condition checklist in Section 5.B. What does your score imply about how close it is to true perfect discrimination?
3. A pharmaceutical company practices near-perfect-discrimination pricing for a life-saving medication, charging wealthy patients much more than poor patients. Is this good policy? Does your answer change if the company uses its profits to fund research and development?
4. If consumers respond to data-driven personalization by investing in privacy tools, who benefits and who loses? Sketch a diagram illustrating the welfare effects.
5. Can you identify a market where near-perfect-discrimination pricing would unambiguously improve welfare relative to any realistic alternative? What features of that market are essential?

B. Assessment

The following prompt has proven effective as a short written assessment or exit ticket after completing the module:

"In the five-consumer example, compute consumer surplus, producer surplus, and deadweight loss under (a) perfect competition, (b) standard monopoly, and (c) the perfectly discriminating monopolist. Then answer: in which scenario are consumers better off? In which is total welfare highest? Are these the same scenario? Explain why or why not."

Students who can correctly produce all three welfare calculations and articulate the efficiency-equity distinction in the final sentence have met the core learning objectives of the module. A common error is to answer that total welfare is highest under perfect competition; this is a sign that the student has not yet fully separated consumer surplus from total welfare in the perfect-discrimination case.

C. The Willingness-to-Pay Simulation: A Classroom Experiment

This experiment, which takes approximately 20 minutes, vividly illustrates how information and market structure jointly determine welfare outcomes. It requires no materials beyond index cards. Because it runs longer than the 60-minute lesson plan allows, it is best used in a 75- or 90-minute session or as a follow-up class.

Divide the class into buyers and one seller. Give each buyer a private willingness-to-pay card with a value drawn from a distribution written on the board. Run three rounds in sequence. First, the seller posts a single price, which approximates standard monopoly. Second, the seller collects and observes every buyer's card and charges each buyer exactly the value on it, which approximates perfect discrimination directly. The defining feature of the perfectly discriminating monopolist is that the seller knows each buyer's willingness to pay, so having the seller read the cards models that knowledge faithfully. Third, introduce multiple competing sellers, which approximates perfect competition. After each round, calculate consumer surplus, producer surplus, and total welfare publicly.

A note on why this is not run as a sealed-bid auction: it is tempting to elicit valuations through an auction, but no simple sealed-bid auction does what this exercise needs. In a second-price auction, truthful bidding is optimal, but the winner pays the second-highest bid and keeps surplus, so the seller does not extract each buyer's full value. In a first-price, pay-your-bid auction, the seller could collect each buyer's own bid, but truthful bidding is no longer optimal, because buyers shade their bids below their true values. Since perfect discrimination requires charging each buyer their actual reservation price, the cleanest classroom model is direct observation of the cards. Instructors can turn this into a teaching point by asking students why an auction fails to reproduce first-degree pricing; the answer previews the information and commitment problems in Section 5.

The experiment consistently produces the pattern predicted by theory. The perfect-discrimination round concentrates surplus with the seller, while the competitive round distributes it to buyers. Students who have participated in all three rounds find the welfare rankings in Section 6.A immediately intuitive. Bergstrom and Miller (2000) document similar results from classroom market simulations, finding that active participation substantially improves students' ability to relate abstract welfare concepts to observed outcomes.

D. A 60-Minute Lesson Plan

Table 2 presents a structured lesson plan for teaching the comparison in a single 60-minute class session. It is designed for a standard intermediate microeconomics course, assuming students have already covered perfect competition and standard monopoly. The numerical example appears in the plan as a 15-minute instructor-led demonstration; instructors who prefer the paired version should assign it as pre-class preparation, as noted in Section 4. Instructors with 75 or 90 minutes can add the willingness-to-pay simulation from Section 7.C.

Table 2. A 60-Minute Lesson Plan for the Perfect-Competition and Perfect-Discrimination Comparison

Time	Phase	Instructor Activity	Student Activity
0-5 min	Hook	Pose the hand-raising question: 'Is a monopolist bad for society?' Then reveal the output equivalence between perfect discrimination and perfect competition. Ask: 'What just happened to your intuition?'	Respond; experience cognitive dissonance
5-15 min	Concept Review	Briefly revisit perfect competition and standard monopoly. Introduce perfect discrimination: the firm charges each consumer exactly their willingness to pay.	Take notes; answer quick comprehension checks
15-30 min	Numerical Example (Section 4)	Work through the five-consumer example on the board in three rounds: perfect competition, standard monopoly, then perfect discrimination. Pause after each to ask students to compute surplus.	Calculate consumer surplus, producer surplus, and deadweight loss for each market structure; compare results
30-40 min	Structured Comparison (Table 1)	Distribute or display Table 1. Guide students row by row, emphasizing the efficiency-equity distinction and the information-and-identification row.	Complete a blank version of Table 1 from memory; discuss with a partner
40-50 min	Data-Driven Pricing Application	Present the five-condition checklist (Section 5.B). Assign one platform per group (Amazon, Spotify, Uber, Netflix). Groups evaluate their platform against the five conditions.	Group work: score assigned platform; report back to class
50-60 min	Discussion and Wrap-Up	Pose Discussion Question 1 (fairness) or Question 3 (pharmaceutical pricing). Summarize: efficiency and equity are separable, which is the lesson.	Open discussion; exit ticket: 'In one sentence, what is the paradox?'

The plan operationalizes the three design principles set out in the introduction. The hook creates cognitive conflict by activating the misconception that monopoly is always welfare-reducing and then confronting it with the output equivalence (Kapur, 2016). The numerical example and the checklist put application before general statement, reflecting Becker and Watts's (2001) finding on applied connection. And the whole sequence targets the efficiency-equity misconception directly, the conflation of efficiency with desirability that Feldman and Serrano (2006) identify as especially persistent. The example is run before the table by design, following the principle of generation before instruction. Here, students who compute the welfare numbers themselves arrive at Table 1 with intuitions already formed, so the table consolidates rather than introduces.

8. Conclusion

The comparison between perfect competition and the perfectly discriminating monopolist is one of intermediate microeconomics' richest teaching opportunities. It forces students to distinguish efficiency from equity, information from market structure, and theoretical results from their conditions of validity. In doing so, it develops the kind of careful analytical reasoning that economics education researchers identify as the discipline's most durable contribution to student thinking (Siegfried et al., 1991).

The result that perfect discrimination replicates perfect competition's allocative efficiency while inverting its distributional properties is not a paradox to be resolved but a genuine feature of market theory to be understood. In an era of data-driven pricing, understanding exactly why the equivalence holds in theory and why it fails in practice equips students to reason clearly about personalized pricing, consumer protection, and the limits of efficiency arguments in policy debate.

The material requires nothing more than supply-and-demand diagrams and a willingness to think carefully about who gets what. The five-consumer example can anchor a single lecture, the data-driven application a seminar discussion, and the lesson plan a coherent 60-minute class. Instructors who wish to go further can pair the module with Dubé and Misra (2023) for empirical evidence, and with Bergemann, Brooks, and Morris (2015) for the formal welfare theory underlying the efficiency claims.

References

- Becker, W. E., & Watts, M. (2001). Teaching economics at the start of the 21st century: Still chalk-and-talk. *American Economic Review*, 91(2), 446-451. DOI: [10.1257/aer.91.2.446](https://doi.org/10.1257/aer.91.2.446)
- Bergemann, D., Brooks, B., & Morris, S. (2015). The limits of price discrimination. *American Economic Review*, 105(3), 921-957. DOI: [10.1257/aer.20130848](https://doi.org/10.1257/aer.20130848)
- Bergstrom, T. C., & Miller, J. H. (2000). *Experiments with economic principles: Microeconomics* (2nd ed.). McGraw-Hill.
- Dubé, J.-P., & Misra, S. (2023). Personalized pricing and consumer welfare. *Journal of Political Economy*, 131(1), 131-189. DOI: [10.1086/720793](https://doi.org/10.1086/720793)
- Dworczak, P., Kominers, S. D., & Akbarpour, M. (2021). Redistribution through markets. *Econometrica*, 89(4), 1665-1698. DOI: [10.3982/ECTA16671](https://doi.org/10.3982/ECTA16671)
- Feldman, A. M., & Serrano, R. (2006). *Welfare economics and social choice theory* (2nd ed.). Springer. DOI: [10.1007/0-387-29368-X](https://doi.org/10.1007/0-387-29368-X)
- Hansen, W. L., Salemi, M. K., & Siegfried, J. J. (2002). Use it or lose it: Teaching literacy in the economics principles course. *American Economic Review*, 92(2), 463-472. DOI: [10.1257/000282802320191813](https://doi.org/10.1257/000282802320191813)
- Kapur, M. (2016). Examining productive failure, productive success, unproductive failure, and unproductive success in learning. *Educational Psychologist*, 51(2), 289-299. DOI: [10.1080/00461520.2016.1155457](https://doi.org/10.1080/00461520.2016.1155457)
- Pigou, A. C. (1920). *The economics of welfare*. Macmillan.
- Rhodes, A., & Zhou, J. (2024). Personalized pricing and competition. *American Economic Review*, 114(7), 2141-2170. DOI: [10.1257/aer.20221524](https://doi.org/10.1257/aer.20221524)
- Robinson, J. (1933). *The economics of imperfect competition*. Macmillan.
- Shiller, B. R. (2020). Approximating purchase propensities and reservation prices from broad consumer tracking. *International Economic Review*, 61(2), 847-870. DOI: [10.1111/iere.12442](https://doi.org/10.1111/iere.12442)
- Siegfried, J. J., Bartlett, R. L., Hansen, W. L., Kelley, A. C., McCloskey, D. N., & Tietenberg, T. H. (1991). The status and prospects of the economics major. *Journal of Economic Education*, 22(3), 197-224. DOI: [10.1080/00220485.1991.10844710](https://doi.org/10.1080/00220485.1991.10844710)
- Walstad, W. B., & Allgood, S. (1999). What do college seniors know about economics? *American Economic Review*, 89(2), 350-354. DOI: [10.1257/aer.89.2.350](https://doi.org/10.1257/aer.89.2.350)