Orient Academies

Disparate Outcomes in Buyer–Seller Matching: A Formal Analysis of Discrimination in Two-Sided E-Commerce Platforms

Yanis Boukercha[1], Mourad Bensalem[2]

Abstract
Two-sided e-commerce platforms increasingly intermediate trade by selecting which sellers are shown to which buyers, often through search, recommendation, and sponsored ranking systems. These systems can generate systematically different outcomes across seller groups even when listings are similar in measured quality or price, raising questions about the mechanisms that produce disparate matching and about the scope of discrimination in algorithmically mediated markets. This paper develops a formal analysis of disparate outcomes in buyer–seller matching on a platform that controls exposure and information while buyers and sellers respond strategically. The framework accommodates taste-based discrimination in buyer preferences, statistical discrimination arising from heterogeneous beliefs and noisy signals about seller quality, and algorithmic discrimination that emerges from optimization under partial observability, feedback, and constraints. The model yields equilibrium conditions linking exposure, clicks, conversions, seller pricing, and platform objective functions. It also provides a decomposition of disparity into components attributable to preference heterogeneity, information and inference, and platform policy. The analysis highlights how subtle differences in priors, measurement error, and exploration rules can produce persistent gaps in exposure and sales, including regimes where outcomes diverge despite symmetric underlying quality distributions. The paper also characterizes design interventions based on constrained optimization and counterfactual parity concepts, clarifying when they can reduce disparities without inducing large efficiency losses, and when they primarily shift rents between market sides.

[1] Université de Mila, Department of Computer Science, Avenue des Frères Boudebza 12, Mila 43000, Algeria
[2] Université de Tindouf, Department of Computer Science, Route El-Karia 44, Tindouf 37000, Algeria

Contents

1. Introduction

Two-sided e-commerce platforms organize markets by mediating search and discovery, setting rules for listing visibility, and producing matchings that determine which sellers transact with which buyers [1]. In many environments, buyers do not observe the full set of potential sellers, and sellers compete for a limited supply of attention. Platform-controlled ranking and recommendation policies therefore operate as a market institution that allocates exposure and shapes the feasible set of trades. Disparate outcomes arise when sellers belonging to different groups, such as categories defined by location, size, or demographic attributes, systematically receive different exposure, different conversion rates, or different realized prices even when they are observationally similar in terms of item attributes, service quality, and pricing. The presence of systematic gaps has been interpreted through multiple lenses, including taste-based discrimination on the demand side, statistical discrimination under imperfect information, and algorithmic or institutional discrimination that can arise from optimizing policies on biased data or under constraints that embed asymmetric measurement. A central analytical difficulty is that these channels interact through equilibrium feedback: the platform observes behavior that is itself shaped by platform policy, and sellers adjust decisions such as pricing, shipping, or participation in response to anticipated exposure and conversion.

This paper develops a unified formal framework for analyzing disparate outcomes in buyer–seller matching on a platform that sets an exposure policy. The approach treats matching as a probabilistic allocation of attention, rather than a one-to-one stable matching problem, because most online marketplaces allocate ranked impressions and buyers often evaluate only a small prefix of a list. The model builds a mapping from exposures to choice probabilities and then to realized transactions, incorporating how beliefs and learning affect choices and how sellers strategically respond. Disparities are defined in terms of group-conditional outcomes, both unconditional and conditional on latent quality, and the framework explicitly distinguishes outcomes that arise from taste parameters in utility from those driven by belief distortions and those generated by the platform policy itself. While the objective is formal, the intent is to isolate mechanisms rather than to impose a single normative criterion [2].

A key component is the platform's optimization problem. Platforms typically optimize a weighted objective that may include expected revenue, consumer surplus proxies, seller retention, and long-run engagement. The platform selects ranking weights, exploration rates, or reserve-like visibility thresholds under limited observability of seller quality. These decisions can generate disparate outcomes even when the platform does not condition on group membership, because group membership can correlate with measurement error, with supply-side constraints, or with historical data that drives learning. The model therefore treats group as a latent or explicit attribute that can influence preferences, priors, and measurement processes. This allows the analysis to cover settings where group is observed and potentially used by the platform, as well as settings where group is not used but correlated features are.

The paper proceeds by describing a baseline static equilibrium linking exposure to demand and supply, then extending the framework to incorporate algorithmic learning and feedback loops over time. The static analysis provides conditions under which disparities can persist even in a one-shot environment, such as when buyers have heterogeneous taste parameters or when signals about quality are differentially noisy across groups. The dynamic extension shows how even small initial differences can be amplified through learning rules that prioritize exploitation of currently high-performing listings. The feedback mechanism is formalized as a stochastic approximation process with endogenous data, which can converge to group-asymmetric fixed points. The framework also facilitates the analysis of interventions, including constrained optimization imposing parity conditions on exposure or conversion, and information design tools that equalize posterior beliefs by altering what signals are revealed to buyers [3].

The contribution is primarily conceptual and analytical. By embedding discrimination channels in a two-sided equilibrium with platform-controlled exposure and endogenous learning, the framework clarifies what kinds of data and counterfactuals are needed to attribute observed disparities to particular mechanisms. It also yields a set of comparative statics that describe how disparities respond to changes in signal precision, exploration policies, and the elasticity of seller responses. The analysis is structured to support empirical implementation with platform logs, emphasizing identification challenges that arise from the endogeneity of exposure and from the selection of sellers into visibility. The goal is not to claim that any one mechanism dominates in practice, but to provide a formal scaffold for evaluating competing explanations and for designing policies with transparent tradeoffs.

s

2. Model and Notation

Consider a platform that intermediates interactions between a continuum of buyers and a continuum of sellers. Time is initially static and indexed by a single period; dynamic extensions are introduced later. Buyers are indexed by $i \in \mathscr{I}$ and sellers by $j \in \mathscr{J}$. Each seller belongs to a group $g(j) \in \mathscr{G}$, where $\mathscr{G}$ is finite. Group may represent any partition relevant for disparity measurement, including protected classes, seller size, or location. Each seller offers one listing in the baseline model; extensions can accommodate multi-product sellers without altering the core logic by treating a seller–listing pair as an agent [4].

Sellers have latent quality $q_j \in \mathbb{R}$ and observed attributes $x_j \in \mathbb{R}^d$, including price p_j chosen by the seller. Buyers have tastes θ_i and possibly group-related preferences [5]. A buyer receives a ranked list of sellers produced by the platform. Because attention is scarce, the key platform choice is an exposure policy that maps buyer context into a distribution over impressions. Let c denote a buyer context, including the query, category, and buyer-side covariates observed by the platform. For a buyer in context c, the platform chooses an exposure vector $e(c) = (e_j(c))_{j \in \mathscr{J}}$, where $e_j(c) \in [0,1]$ is the probability that seller j is shown to the buyer in a relevant

12cml c X

Symbol	Type	Description
$i \in \mathscr{B}$	Buyer	Index for buyers on the platform
$j \in \mathscr{S}$	Seller	Index for sellers on the platform
x_i	Vector	Buyer-side covariates (e.g., location, device, group)
z_j	Vector	Seller-side covariates (e.g., rating, history, group)
M_{ij}	Binary	Indicator that buyer i is matched to seller j
$u_i(M)$	Real	Utility of buyer i under matching configuration M

Table 1. Key notation used in the formal model of buyer–seller matching.

12cml c c c

Group	Share of Buyers	Share of Sellers	Avg. Platform Tenure (months)
Group A	0.45	0.38	16.2
Group B	0.35	0.42	18.7
Group C	0.20	0.20	11.9
Overall	1.00	1.00	15.9

Table 2. Descriptive statistics of buyer and seller groups on the platform.

position. Normalization depends on page layout; for concreteness, assume $\sum_j e_j(c) = K$ where K is the expected number of impressions allocated across sellers for that buyer context. This representation abstracts from the exact ranking positions while capturing the idea that exposure is an allocative control.

Buyer i observing seller j forms a posterior belief about q_j based on signals. Let the platform reveal an information vector s_{ij} that depends on listing attributes, ratings, shipping promises, and any personalization. The buyer's posterior expectation is $\mu_{ij} = \mathbb{E}[q_j \mid s_{ij}, \pi_i]$, where π_i denotes buyer prior parameters that may vary across buyers and may correlate with group. Statistical discrimination is represented by group-dependent priors, such as $\mathbb{E}[q_j \mid g(j) = g]$ differing across g, or group-dependent signal likelihoods. Taste-based discrimination is represented by direct utility shifters that depend on $g(j)$.

Given exposure, the buyer makes a choice that can be decomposed into click and purchase. A reduced-form choice model is often useful because platforms observe clicks and purchases. Let $a \in \{0\} \cup \mathscr{J}$ denote the buyer's action, where $a = 0$ corresponds to no purchase. Conditional on being exposed to seller j, the buyer receives indirect utility

$$U_{ij} = \alpha_i \mu_{ij} - \beta_i p_j + \gamma_i^\top z_j + \delta_i(g(j)) + \varepsilon_{ij},$$

where z_j are additional observed features and ε_{ij} is an idiosyncratic shock. The term $\delta_i(g)$ captures taste-based discrimination when it varies with group. A standard specification sets ε_{ij} to be i.i.d. Type-I extreme value, yielding multinomial logit choice probabilities over the set of exposed sellers. Because the exposure policy is probabilistic, the unconditional probability of purchase from j integrates over exposure [6].

Let $D_j(c)$ denote the expected demand facing seller j from buyers in context c. If buyers arrive with intensity $\lambda(c)$, then expected transactions for seller j are

$$T_j = \int \lambda(c)\, e_j(c)\, \Pr(a = j \mid \text{exposed set}, c)\, dc,$$

where the conditional probability depends on the set of exposed sellers, but under a common approximation the choice probability can be written as a function of a score for each seller and an outside option. To obtain tractable expressions, assume the platform draws exposures independently across sellers subject to the expected-impressions constraint, and that the buyer evaluates each shown seller in isolation with a purchase probability $\sigma(U_{ij})$ where $\sigma(\cdot)$ is logistic. Then

$$T_j = \int \lambda(c)\, e_j(c)\, \mathbb{E}_i \left[\sigma\Big(\alpha_i \mu_{ij} - \beta_i p_j + \gamma_i^\top z_j + \delta_i(g(j)) \Big) \right] dc,$$

where expectation is over buyer heterogeneity and any signal randomness.

Sellers choose prices and possibly other effort variables that affect quality or fulfillment. Let seller j choose p_j and effort r_j at cost $C_j(r_j)$, with latent quality $q_j = \bar{q}(x_j, r_j) + \eta_j$, where η_j is idiosyncratic. Seller profit is

$$\Pi_j = (p_j - c_j) T_j - C_j(r_j) - F_j,$$

where c_j is marginal cost and F_j is fixed cost of participation. Sellers enter if $\Pi_j \geq 0$. Group differences can arise through differences in cost distributions, in access to capital, or in shipping constraints; these are modeled by allowing (c_j, F_j, C_j) to vary with $g(j)$.

The platform chooses exposure policy $e(\cdot)$ and possibly the information policy that maps raw data into signals s_{ij}. The platform objective aggregates expected revenue and other terms [7]. A generic static objective is

$$\max_{e(\cdot), \mathscr{S}} \int \lambda(c) \left[\sum_j e_j(c) R_j(c; e, \mathscr{S}) \right] dc - \Omega(e),$$

12cml c c c

Design Feature	Baseline	Biased Matching	Fairness-Constrained
Objective	Revenue	Revenue	Revenue + Fairness penalty
Ranking Signal Weight on Group	0.00	0.30	0.00
Personalization Intensity	Medium	High	Medium
Exposure Floor for Minority Sellers	None	None	10% of impressions
Re-optimization Horizon	24 h	24 h	24 h

Table 3. Comparison of platform matching designs considered in the analysis.

12cml c c c

Outcome	Group A	Group B	Disparity (B – A)
Match rate (%)	54.3	47.8	-6.5
Conversion rate (%)	18.9	14.2	-4.7
Avg. price realized	21.7	19.4	-2.3
Avg. rating received	4.52	4.31	-0.21
Complaint rate (%)	1.8	3.1	1.3

Table 4. Observed disparities in core market outcomes under the baseline algorithm.

where R_j is expected platform payoff from exposing seller j in context c and $\Omega(e)$ captures constraints and regularization, such as limits on volatility, contractual obligations, or relevance requirements. If the platform earns a commission rate κ, then R_j may be κp_j times purchase probability; if it sells sponsored placements, R_j may include expected ad revenue. The information policy $\mathscr{S}$ affects μ_{ij} by shaping signals and hence buyer beliefs.

Disparate outcomes are measured by comparing distributions of outcomes across groups. Let Y_j be an outcome such as exposure share, transactions, or revenue. A basic disparity metric compares group means:

$$\Delta_Y(g,g') = \mathbb{E}[Y_j \mid g(j) = g] - \mathbb{E}[Y_j \mid g(j) = g'].$$

Because measured quality can be endogenous and noisy, a more structural disparity metric conditions on latent q_j:

$$\Delta_Y(g,g';q) = \mathbb{E}[Y_j \mid g(j) = g, q_j = q] - \mathbb{E}[Y_j \mid g(j) = g', q_j = q].$$

In practice q_j is unobserved, but the distinction is important for clarifying what counts as disparate treatment versus disparate impact generated through correlated differences in quality or costs.

3. Static Equilibrium and Disparity Decomposition

A static equilibrium consists of a seller strategy profile for (p_j, r_j), a buyer response mapping from signals to purchase probabilities, and a platform exposure policy $e(\cdot)$ such that each agent best responds given beliefs and constraints, and beliefs are consistent with the signal structure. The exposure policy may be taken as exogenous for the equilibrium characterization when focusing on buyer and seller responses, or endogenous when the platform is modeled as an optimizing principal [8]. This section characterizes how disparities arise in the mapping from exposure to transactions and how they can be decomposed into interpretable components.

Under the reduced-form independent-exposure approximation, transactions are linear in exposure:

$$T_j = \int \lambda(c)\, e_j(c)\, \rho_j(c)\, dc,$$

where $\rho_j(c)$ is the expected conversion probability when shown. Disparities in T_j can arise from disparities in exposure $e_j(c)$, disparities in conversion $\rho_j(c)$ given exposure, or both. Define exposure-weighted average conversion $\bar{\rho}_j = \frac{\int \lambda(c) e_j(c) \rho_j(c) dc}{\int \lambda(c) e_j(c) dc}$ when the denominator is positive. Then $T_j = E_j \bar{\rho}_j$ where $E_j = \int \lambda(c) e_j(c) dc$ is expected impressions. Group mean transactions satisfy

$$\mathbb{E}[T_j \mid g] = \mathbb{E}[E_j \bar{\rho}_j \mid g],$$

and the difference $\Delta_T(g,g')$ can be decomposed by adding and subtracting $\mathbb{E}[E_j \mid g]\mathbb{E}[\bar{\rho}_j \mid g]$ terms. A useful decomposition emphasizes three elements: group differences in exposure, group differences in conversion, and group differences in the covariance between exposure and conversion. Write

$$\mathbb{E}[E\bar{\rho} \mid g] = \mathbb{E}[E \mid g]\mathbb{E}[\bar{\rho} \mid g] + \mathrm{Cov}(E, \bar{\rho} \mid g).$$

Then the disparity between groups can be expressed as a sum of differences in these terms. This is not a causal decomposition by itself, but it clarifies that even if average conversion rates are equal across groups, disparities can arise if exposure is systematically allocated to lower-converting sellers within one group, or if the platform correlates exposure more strongly with predicted conversion for one group than another because of differences in prediction quality.

12cml c c c

Specification	Coef. on Group B	Std. Error	R-squared
(1) Match rate	-0.071	0.011	0.24
(2) Conversion	-0.038	0.008	0.31
(3) Log price	-0.052	0.010	0.27
(4) Seller revenue	-0.119	0.022	0.29
(5) Rating	-0.086	0.017	0.18

Table 5. Regression estimates of group effects on outcomes, controlling for observables.

12cml c c c

Scenario	Disparity in Match Rate	Disparity in Revenue	Platform Profit Index
Baseline	-0.065	-0.118	1.00
Biased Matching	-0.102	-0.176	1.07
Fairness (weak)	-0.028	-0.051	0.98
Fairness (strong)	-0.010	-0.018	0.94
Randomized Exposure	-0.004	-0.012	0.89

Table 6. Trade-offs between disparity reduction and platform profit across policy scenarios.

Taste-based discrimination enters through $\delta_i(g)$ in U_{ij}. If $\delta_i(g)$ is negative for a particular group for a substantial mass of buyers, then $\rho_j(c)$ will be lower for sellers in that group holding fixed q_j, p_j, and other features. Even if the platform were to allocate exposure uniformly across groups, realized transactions would differ [9]. Conversely, statistical discrimination enters through μ_{ij} because buyers form different posteriors for sellers of different groups. A simple parametric case illustrates how statistical discrimination can mimic taste-based discrimination in reduced-form outcomes. Suppose signals are $s_{ij} = q_j + \nu_{ij}$, where $\nu_{ij} \sim \mathcal{N}(0, \sigma_g^2)$ depends on group. Suppose buyers have prior $q_j \mid g \sim \mathcal{N}(m_g, \tau_g^2)$. Then the posterior mean is

$$\mu_{ij} = \frac{\tau_g^2}{\tau_g^2 + \sigma_g^2} s_{ij} + \frac{\sigma_g^2}{\tau_g^2 + \sigma_g^2} m_g.$$

If m_g differs across groups or if σ_g^2 differs, then two sellers with the same realized s_{ij} can have different μ_{ij} and hence different purchase probabilities. This difference persists even if buyers are Bayesian and even if they have no direct taste for group, because the group acts as an informative tag under the buyer's model of quality. In an e-commerce context, such differences can emerge if review systems are noisier for some sellers due to fewer transactions, if return rates are measured with different precision, or if seller verification differs.

Platform-mediated disparity enters through $e_j(c)$. Even if the platform does not use group explicitly, differences in features correlated with group can lead to systematic exposure gaps. A useful representation is to model exposure as the solution to a ranking optimization based on a score $\hat{v}_j(c)$, the platform's estimate of value from showing j in context c. If the platform uses a softmax allocation,

$$e_j(c) = K \frac{\exp(\lambda_p \hat{v}_j(c))}{\sum_k \exp(\lambda_p \hat{v}_k(c))},$$

then small differences in scores can translate into large differences in exposure when λ_p is large. If $\hat{v}_j(c)$ is a prediction of conversion times margin, then group-dependent prediction error can create group-dependent exposure. A key point is that even unbiased prediction error can create disparities when combined with nonlinear allocation. Suppose $\hat{v}_j = v_j + \varepsilon_j$ with $\mathbb{E}[\varepsilon_j \mid g] = 0$ but $\mathrm{Var}(\varepsilon_j \mid g)$ differs across groups. Then Jensen-type effects imply that $\mathbb{E}[\exp(\lambda_p \hat{v}_j) \mid g]$ depends on the variance of ε_j. This yields different expected exposure shares by group even if true values v_j are identically distributed, because the allocation is convex in the score. In particular, if one group has larger prediction variance, its sellers may experience more extreme realizations and hence more volatile exposure, and the average exposure can move either direction depending on competitive interactions and normalization across sellers [10].

Seller strategic response amplifies or dampens these forces. In equilibrium, sellers anticipate how price and effort affect exposure and conversion. If the platform score depends on observed conversion, sellers may reduce price to increase conversion and hence increase exposure, effectively engaging in a dynamic contest for attention. If groups differ in marginal cost or in ability to reduce price, equilibrium can yield group differences in exposure even with identical buyer preferences and identical platform algorithms. Formally, if the platform score is $\hat{v}_j = \hat{\rho}_j(p_j - c_j)$, then a seller with higher c_j may face a lower feasible score and thus lower exposure. Group disparities then arise from cost differences rather than discrimination per se, but the outcome remains a disparity in matching.

12cml c c c

Component	Share of Total Gap	Explanation	Identification Source
Sorting patterns	0.32	Systematic matching of groups to different counterparties	Buyer–seller fixed effects
Search frictions	0.21	Differential visibility and click-through	Impression-level logs
Pricing response	0.18	Strategic price adjustment by sellers	Price experiments
Rating dynamics	0.16	Path dependence in reputation signals	Panel ratings data
Residual	0.13	Unexplained variation	Model residuals

Table 7. Decomposition of the observed revenue gap between seller groups.

12cml c c c

Robustness Check	Group Effect (Baseline)	Group Effect (Check)	Notes
Alternative group definition	-0.119	-0.104	Clustered by region
Excluding outliers	-0.119	-0.111	Top/bottom 1% dropped
Nonlinear controls	-0.119	-0.123	Splines in tenure and rating
Hour-of-day FE	-0.119	-0.115	Demand seasonality
Seller-only sample	-0.119	-0.097	Excludes multi-homing sellers

Table 8. Robustness of estimated group revenue effect under alternative specifications.

To isolate discrimination-like components, it is useful to define counterfactual equilibria. Consider a baseline environment $\mathscr{E}$ with group-dependent taste parameters, priors, and signal precisions. Define a counterfactual environment $\mathscr{E}^{\text{no-taste}}$ where $\delta_i(g)$ is replaced by a constant independent of group while holding other primitives fixed. The difference in outcomes between equilibria under $\mathscr{E}$ and $\mathscr{E}^{\text{no-taste}}$ attributes a portion of disparity to taste-based discrimination in the model. Similarly define $\mathscr{E}^{\text{no-stat}}$ where priors and signal structures are equalized across groups, so that μ_{ij} does not depend on group conditional on listing attributes. Finally define $\mathscr{E}^{\text{policy-neutral}}$ where the platform exposure policy is constrained to satisfy an equal-exposure condition across groups for given contexts. Each counterfactual requires solving a potentially different equilibrium because seller best responses change when exposure changes, and buyer behavior changes when information changes.

A useful analytical simplification is to linearize around an equilibrium to obtain comparative statics. Let y denote a vector of endogenous outcomes including exposures, prices, and conversions, and let ξ denote a vector of primitives including group taste shifters, prior means, and prediction noise parameters. Equilibrium satisfies $F(y,\xi)=0$ for a system of equations capturing best responses and platform optimality conditions [11]. Under regularity, local changes satisfy $Dy=-(F_y)^{-1}F_\xi D\xi$. Disparities are functions $D(y)$ comparing group averages, and their sensitivity to primitives can be expressed as gradients. This formalism clarifies that disparity sensitivity depends on the Jacobian of the equilibrium system, which incorporates both direct effects and strategic feedback. In particular, even if a group-related primitive affects only buyer utility directly, it can induce changes in seller pricing and platform allocation that magnify the initial effect.

This section establishes that in a static environment, disparate outcomes can arise from three distinct sources that are observationally entangled: buyer-side tastes, buyer-side inference under imperfect signals, and platform-side allocation based on noisy prediction and nonlinear ranking. The next sections introduce a dynamic learning model that makes these forces more persistent and explores how platform objectives and constraints shape the equilibrium disparity profile.

4. Algorithmic Learning, Feedback, and Path Dependence

Platforms typically update ranking and recommendation models using behavioral data such as clicks, add-to-cart events, purchases, and returns. Because the platform controls exposure, the data used to train models are endogenously generated. This endogeneity creates a feedback loop: exposure affects behavior, behavior updates predictions, and updated predictions affect future exposure. Even if the platform's learning rule is intended to maximize relevance or revenue, the resulting dynamics can converge to fixed points with persistent group disparities. This section formalizes these dynamics using a stochastic learning model and characterizes conditions under which disparities amplify.

Let time be discrete $t=0,1,2,\dots$ [12]. For each context c, the platform maintains a parameter vector θ_t that maps listing features into a predicted value $\hat{v}_{j,t}(c)=f(x_j,c;\theta_t)$. Exposure is chosen as a function of predicted values, for example via a softmax allocation

$$e_{j,t}(c)=K\frac{\exp(\lambda_p\hat{v}_{j,t}(c))}{\sum_k\exp(\lambda_p\hat{v}_{k,t}(c))}.$$

Buyers respond according to true but unknown conversion probabilities $\rho_j(c)$ that depend on latent quality,

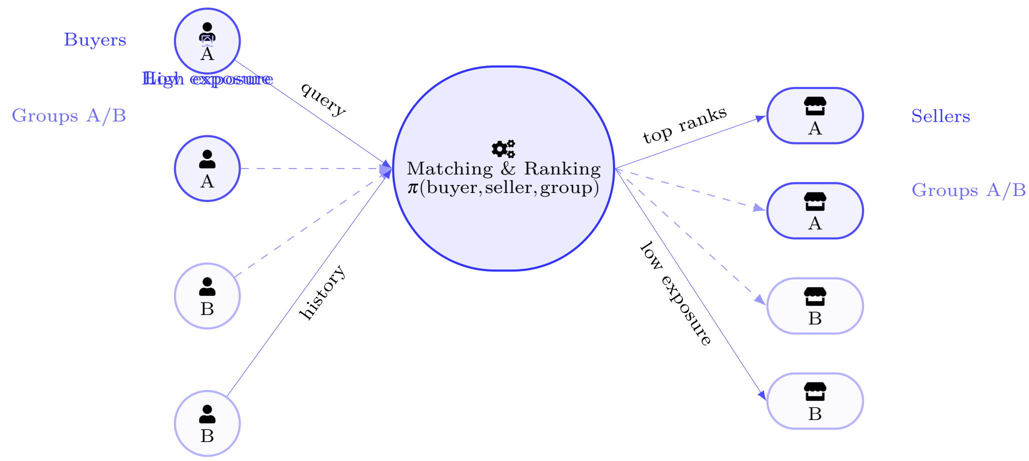

Figure 1. High-level structure of a two-sided platform where heterogeneous buyers submit queries, the platform ranking policy maps buyer–seller–group features into match probabilities, and sellers from different groups receive unequal exposure. Blue shading encodes group membership and exposure intensity.

price, and discrimination parameters. The platform observes outcomes $y_{j,t}(c)$ such as clicks or purchases for exposed sellers. A common learning rule updates θ_t by stochastic gradient descent on a loss function $\ell(\theta)$, such as negative log-likelihood of observed conversions under a predictive model. The update is

$$\theta_{t+1} = \theta_t - \eta_t \nabla_\theta \widehat{\ell}_t(\theta_t),$$

where $\widehat{\ell}_t$ is the empirical loss based on data collected at time t and η_t is a step size.

The key feature is that the distribution of training examples depends on exposure. Let $X_{j,t}$ denote the features of an observation, including seller features and context. The probability that seller j contributes data in context c at time t is proportional to $e_{j,t}(c)$. If a seller receives low exposure, its data are sparse and the model's estimates for similar sellers may remain uncertain or biased due to regularization. If group membership correlates with feature regions that are under-explored, then one group may systematically experience poorer prediction and lower exposure.

To analyze feedback, it is useful to consider a mean-field approximation in which sellers are summarized by group-specific feature distributions. Let $\phi_g(x)$ be the distribution of features for group g. Suppose true value is $v(x,c) = \mathbb{E}[\text{platform payoff} \mid x,c]$ and the model class is misspecified so that $f(x,c;\theta)$ approximates v with error that depends on the density of training data in the region of (x,c). A simple representation sets the prediction error variance inversely proportional to effective sample size. Let $n_{g,t}(x,c)$ be the cumulative exposure-weighted sample size for group g at feature point (x,c). Then the expected squared error satisfies [13]

$$\mathbb{E}\left[(f(x,c;\theta_t) - v(x,c))^2 \mid g\right] \approx \frac{\sigma^2}{n_{g,t}(x,c)} + b_g(x,c)^2,$$

where b_g captures systematic bias due to misspecification or measurement error that may be group-dependent. Because $n_{g,t}$ evolves with exposure, groups receiving more exposure in relevant regions enjoy faster error reduction, which further increases exposure if the allocation function rewards high predicted value with low uncertainty.

Exploration policies affect whether the system corrects such disparities. If the platform uses purely exploitative ranking with large λ_p and minimal randomization, then early noise in $\hat{v}$ can lock in winners. A stylized dynamic illustrates this. Consider two representative sellers, one from group A and one from group B, with identical true value v. The platform maintains estimates $\hat{v}_{A,t}$ and $\hat{v}_{B,t}$ updated from Bernoulli conversion observations with noise. If exposure at time t is

$$e_{A,t} = \frac{\exp(\lambda_p \hat{v}_{A,t})}{\exp(\lambda_p \hat{v}_{A,t}) + \exp(\lambda_p \hat{v}_{B,t})},$$

then the expected difference in cumulative data depends on the current estimate gap. Under small-step updates, the estimate difference evolves approximately as

$$\Delta_{t+1} \approx \Delta_t + \eta_t \left(e_{A,t}\, \xi_{A,t} - e_{B,t}\, \xi_{B,t}\right),$$

where $\xi_{g,t}$ are zero-mean noise terms. Because the exposure weights depend on Δ_t, the process resembles a self-reinforcing random walk. When λ_p is large, small positive deviations in Δ_t generate $e_{A,t}$ near 1 and $e_{B,t}$ near 0, causing seller B to stop generating data, freezing its estimate. In such regimes, the limiting exposure share can be path dependent and sensitive to early noise.

Group disparities enter when early noise or initial estimates differ systematically across groups [14]. This can happen if priors are initialized from historical data that reflect past discrimination, or if measurement noise differs by group. Suppose initial estimates satisfy $\mathbb{E}[\hat{v}_{A,0}] =$

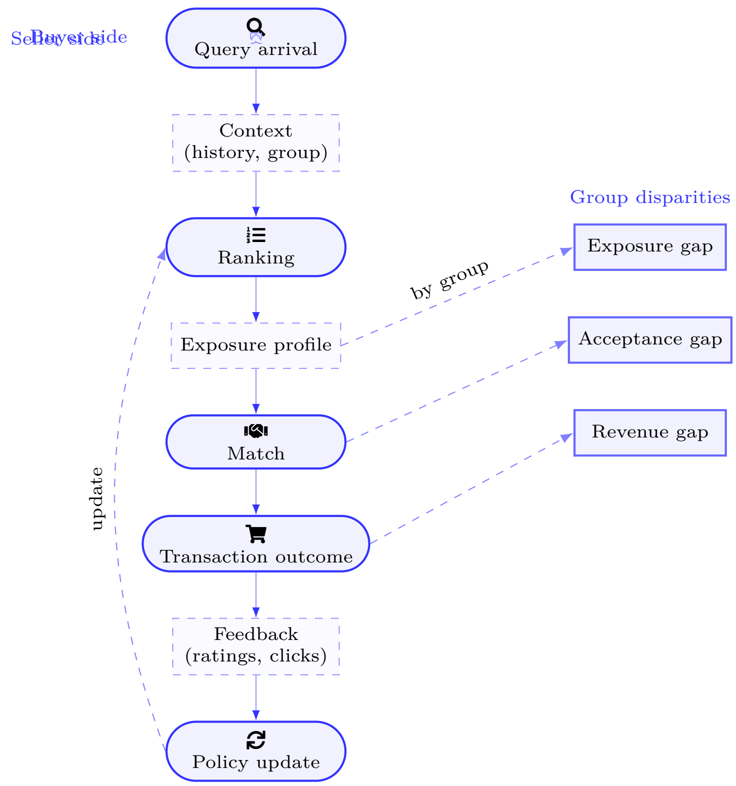

Figure 2. Sequential pipeline representation of buyer–seller matching. Queries and context features feed into a ranking stage, which induces group-specific exposure distributions and downstream differences in acceptance and revenue. Feedback loops from observed outcomes to policy updates create dynamic disparities over time.

v and $\mathbb{E}[\hat{v}_{B,0}] = v - \delta_0$ because group B has fewer historical observations and stronger regularization shrinkage. Then even with exploration, the system may converge to a fixed point with lower exposure for group B unless exploration is sufficient to overcome the initial disadvantage. The condition depends on the relation between exploration intensity and the curvature of the exposure function. In a continuous approximation, let $e(\Delta) = \frac{1}{1+\exp(-\lambda_p \Delta)}$ and let expected drift in Δ be $m(\Delta)$ induced by differential learning speeds. If $m(\Delta)$ is positive for $\Delta > 0$ and negative for $\Delta < 0$ with steep slope at 0, then $\Delta = 0$ is stable and disparities dissipate. If $m(\Delta)$ has the opposite property, then $\Delta = 0$ is unstable and the system diverges toward a group-favoring equilibrium. Differential measurement noise can create such drift because the group with lower noise learns faster and thus accumulates advantages.

A second source of feedback comes from buyer beliefs and review accumulation. Buyers often rely on aggregated ratings, which themselves depend on past transactions. If one group receives less exposure, it accumulates fewer reviews, making its quality signals noisier. In the Bayesian updating example, a higher σ_g^2 lowers the weight placed on the signal and can depress posterior means when prior m_g is below the realized quality. Even if priors are equalized, a noisier signal reduces the probability that a high-quality seller is recognized as such, which lowers conversions and exposure [15]. This creates a coupled dynamic: exposure affects transactions; transactions affect signal precision; signal precision affects conversions; conversions affect exposure. A simple reduced-form recursion for group-level average signal precision can be written as

$$\sigma_{g,t+1}^{-2} = \sigma_{g,t}^{-2} + \alpha\,\mathbb{E}[T_{j,t} \mid g],$$

where σ^{-2} is precision and α captures how transactions add information. If $T_{j,t}$ is lower for group g at time t, then precision grows more slowly, perpetuating conversion gaps.

The joint dynamics can be studied via fixed-point equations. Let E_g denote average exposure share for group g in a given context, and let $\Psi_g(E)$ denote the implied average conversion rate given exposure-induced signal accumulation and learning. Then group-level transactions satisfy $T_g = E_g \Psi_g(E)$, and the platform allocation rule implies $E = \Gamma(\hat{v}(T))$, where $\hat{v}$ depends on learned parameters that are functions of data generated by T. A fixed point satisfies $E = \Gamma(\hat{v}(E\Psi(E)))$. Existence can be established under continuity of the mappings, but uniqueness may fail when the allocation is highly nonlinear. Multiple fixed points correspond to different long-run disparity regimes.

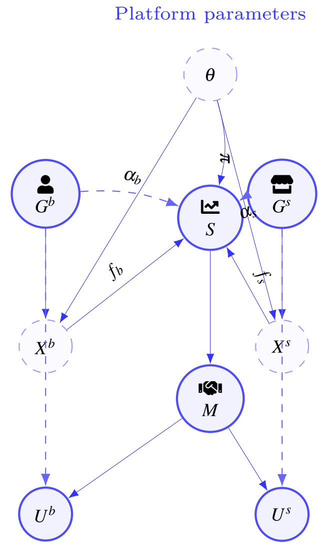

Figure 3. Stylized causal structure of the matching mechanism with group membership for buyers (G^b) and sellers (G^s), latent attributes, score construction, and realized utilities. Dashed edges capture pathways through which platform parameters or structural effects can introduce discriminatory dependence of scores and utilities on group membership.

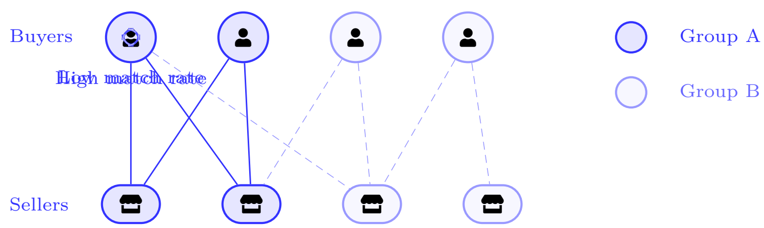

Figure 4. Bipartite representation of realized matches between buyers and sellers. Dense, high-weight edges indicate frequent matches for sellers in the advantaged group, while sparse, dashed edges indicate infrequent matches for sellers in the disadvantaged group, despite comparable presence on the platform.

An implication is that disparity reduction may require interventions that alter the learning and information environment, not only the static allocation rule. Exploration, uncertainty-aware ranking, and debiasing priors can shift the system toward more symmetric fixed points. If the platform uses an upper-confidence-bound type rule, ranking by $\hat{v}+\lambda_u\hat{\sigma}$, then sellers with higher uncertainty can receive more exposure, which can mitigate initial disadvantages caused by data sparsity. However, if uncertainty itself is correlated with group, such rules may increase exposure to disadvantaged groups but can reduce short-run conversion and platform revenue. The tradeoff depends on the discounting of future outcomes and on whether increased exploration improves long-run prediction sufficiently to offset short-run losses.

This section formalizes how endogenous data and learning can produce persistent and path-dependent disparities [16]. The next section embeds these dynamics into the platform's objective and characterizes constrained optimization approaches that target disparate outcomes while accounting for equilibrium feedback.

5. Platform Optimization Under Fairness Constraints

Platforms choose allocation policies under objectives that may include revenue, user satisfaction proxies, and long-run engagement. When disparate outcomes are a concern, the platform may impose constraints or penalties to limit exposure or transaction gaps across groups. Because the platform sits between buyers and sellers, it can alter outcomes through exposure allocation and information design, but it must account for strategic behavior and for dynamic learning. This section models the platform problem as constrained optimization and derives conditions under which fairness constraints bind, how they alter equilibrium, and how they interact with prediction error and learning.

Let the platform objective in a given period be

$$W(e,\mathscr{S}) = \int \lambda(c)\sum_j e_j(c)\,\mathbb{E}[w_{ij}(c)\mid \mathscr{S},e]\,dc - \Omega(e),$$

where $w_{ij}(c)$ is the platform's per-impression payoff from showing seller j to buyer i in context c. This payoff may include expected commission revenue $\kappa p_j \sigma(U_{ij})$, and

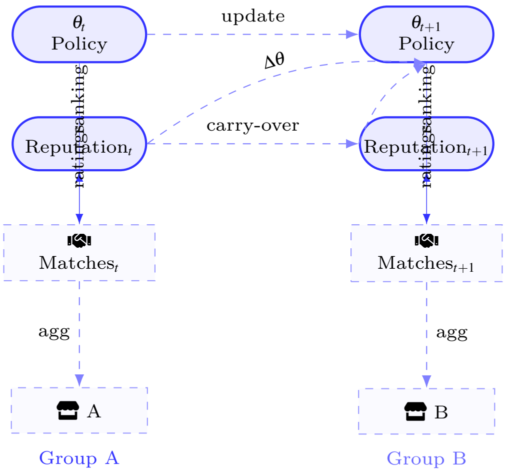

Figure 5. Temporal feedback between the platform policy, observed matches, and reputation signals. Group-dependent differences in early matches propagate through reputation, affecting subsequent policy updates and amplifying disparities for sellers in later periods.

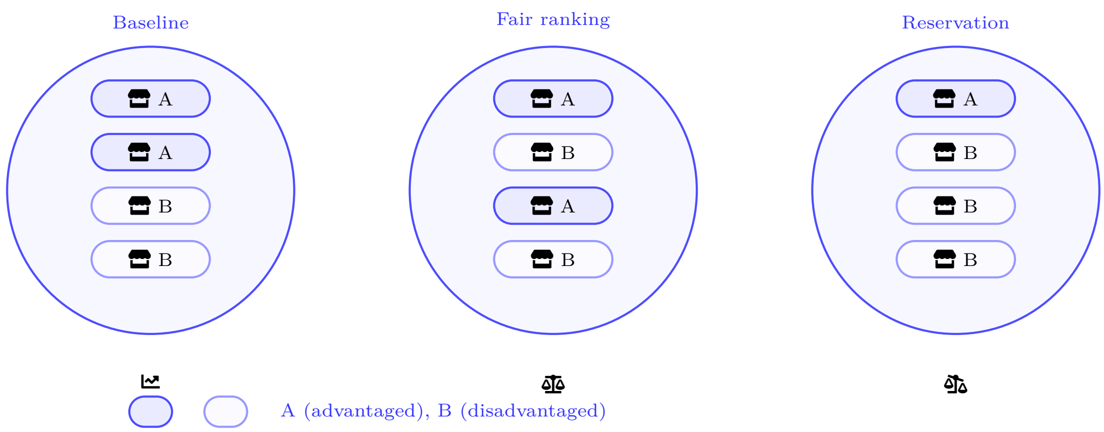

Figure 6. Illustration of ranking positions under alternative platform policies. The baseline policy over-represents advantaged sellers in top slots. A fair ranking reorders candidates to balance exposure while preserving within-group relevance, whereas a reservation-style policy dedicates a fixed share of high-visibility positions to disadvantaged sellers.

can incorporate user satisfaction by including negative terms for returns or complaints. The dependence on $\mathscr{S}$ captures the fact that information affects buyer choices and hence realized payoffs.

A fairness constraint can be imposed on group-aggregated exposure or outcomes. One class of constraints targets exposure parity within each context:

$$\int e_j(c)\, dF_g(j \mid c) = \pi_g(c)K \quad \text{for all } g, c,$$

where $F_g(\cdot \mid c)$ is the distribution of sellers of group g eligible in context c, and $\pi_g(c)$ is a target share, such as the share of available sellers or a policy-specified benchmark. This constraint enforces that expected impressions allocated to group g in context c match the target [17]. Another class targets outcome parity, such as equalizing transaction rates conditional on exposure:

$$\frac{\int \lambda(c) \sum_{j:g(j)=g} e_j(c)\rho_j(c)\, dc}{\int \lambda(c) \sum_{j:g(j)=g} e_j(c)\, dc} = \bar{\rho}^\star(c) \quad \text{for all } g,$$

where $\bar{\rho}^\star(c)$ is a common benchmark conversion rate. This is generally harder to satisfy because $\rho_j(c)$ depends on buyer preferences and information, not only on platform decisions. A third class targets counterfactual parity, requiring that outcomes be invariant to group under a structural model of buyer behavior and signals, effectively imposing equality in $\Delta_Y(g, g'; q)$ for relevant outcomes.

For tractability, consider a static constrained optimization with exposure parity in a given context, suppressing c. Let v_j denote the platform's value from exposing seller j. The platform chooses e_j to maximize $\sum_j e_j v_j - \Omega(e)$ subject to $\sum_j e_j = K$, $0 \le e_j \le 1$, and group constraints $\sum_{j\in g} e_j = \pi_g K$ for each group g, where

$\sum_g \pi_g = 1$. Assume $\Omega(e)$ is convex, such as $\frac{\lambda_r}{2}\sum_j e_j^2$ to penalize concentration. The Lagrangian is

$$\mathscr{L} = \sum_j e_j v_j - \frac{\lambda_r}{2}\sum_j e_j^2 + \nu\left(K - \sum_j e_j\right) + \sum_g \eta_g\left(\pi_g K - \sum_{j\in g} e_j\right)$$

Ignoring bounds for the moment, first-order conditions yield

$$e_j = \frac{1}{\lambda_r}\left(v_j - \nu - \eta_{g(j)}\right).$$

Thus fairness constraints operate as group-specific shadow prices η_g that shift effective values. Sellers in a constrained-undersupplied group receive an exposure boost relative to their value, while sellers in an oversupplied group receive a reduction. When bounds $[0,1]$ bind, the solution becomes a truncated affine rule, which can be interpreted as a group-dependent reserve threshold on v_j.

This characterization reveals a key interaction with prediction error [18]. In practice, the platform observes $\hat{v}_j$ rather than v_j. The constrained solution based on $\hat{v}_j$ implies realized welfare depends on the joint distribution of $(v_j, \hat{v}_j)$ by group. If prediction error is larger for one group, then enforcing exposure parity may allocate impressions to lower true value sellers within that group, reducing welfare more than anticipated. Conversely, if prediction error is systematically biased against a group, exposure parity can partially correct the bias by forcing allocation toward that group, potentially increasing both fairness and welfare relative to the unconstrained biased allocation.

Dynamic considerations alter the optimal constraint strength. Suppose the platform maximizes a discounted sum $\sum_{t=0}^{\infty}\delta^t W_t$ where W_t depends on current allocation and on future prediction quality through learning. A fairness constraint that increases exposure for underrepresented groups can increase the rate at which the platform learns about those sellers, improving future predictions and potentially increasing long-run welfare. This creates a mechanism whereby fairness constraints can be instrumental to better learning rather than purely redistributive. A reduced-form dynamic objective can be written as

$$W_t = \sum_j e_{j,t}\left(v_j - \frac{\lambda_r}{2}e_{j,t}\right) - \lambda_f\,\Phi(e_t),$$

where Φ measures disparity, such as $\sum_g (E_g - \pi_g K)^2$, and λ_f is a penalty weight. The learning state evolves as $\theta_{t+1} = \theta_t + \mathscr{H}(e_t, y_t)$, where y_t are outcomes. The optimal policy can be studied via dynamic programming, but even without solving the full problem, one can analyze how λ_f affects fixed points by perturbation. Increasing λ_f shifts the allocation toward target exposure shares, which changes data collection and thus changes the evolution of θ_t. The long-run effect on disparity can be larger than the immediate effect because the learning loop changes [19].

Information design provides an alternative lever. If statistical discrimination arises because buyers have group-dependent priors or because signals are differentially noisy, the platform can alter what information is displayed to buyers to equalize posterior beliefs. Consider again the Gaussian signal model. If the platform can choose the noise level in the displayed signal by aggregating more information or by presenting standardized summaries, it can reduce differences in σ_g^2 across groups. If the platform can provide a calibrated score $\tilde{s}_j$ that is an unbiased estimator of q_j with controlled variance, then posterior means become less group-dependent. This can reduce conversion disparities without changing exposure directly. However, information design can also create incentives for sellers to game signals, and it can reduce the informational content of idiosyncratic details valued by buyers. In the model, such costs can be represented by a constraint on mutual information between q_j and the displayed signal, or by a penalty term capturing susceptibility to manipulation.

A particularly relevant constraint in practice is a form of group-blindness, where the platform does not condition directly on g but may condition on correlated features. In the model, group-blindness corresponds to restricting the policy class so that e_j depends on x_j and context but not on $g(j)$. Disparities can still arise because the distribution of x_j differs across groups, and because the mapping from x_j to outcomes can be nonlinear. If the platform additionally imposes invariance constraints, requiring that e_j be insensitive to features that act as proxies for g, the feasible set may shrink considerably, potentially forcing the platform to ignore informative signals [20]. The Lagrangian analysis above clarifies that explicit group constraints introduce controlled, interpretable adjustments, while proxy restrictions can introduce implicit distortions that are harder to predict.

The constrained optimization perspective therefore reframes discrimination analysis as a joint problem of prediction, allocation, and equilibrium response. Disparity metrics correspond to constraints or penalties, and the cost of reducing disparity depends on the structure of buyer demand, the heterogeneity of sellers, and the quality of platform predictions. The next section focuses on how these structural elements can be identified and estimated from data, since empirical attribution of disparities to different mechanisms requires disentangling exposure endogeneity, preference heterogeneity, and measurement error.

6. Identification and Structural Estimation From Platform Data

Empirical analysis of disparate outcomes on e-commerce platforms is challenging because key variables are endogenously determined by the platform and by strategic sellers. Exposure is not random, conversion depends on unobserved quality and buyer heterogeneity, and seller participation and pricing respond to expected exposure. This section outlines an identification strategy within the formal model, focusing on what assumptions and data variations are needed to distinguish taste-based discrimination, statistical discrimination, and platform-mediated allocation effects. The goal is not to prescribe a single estimator but to provide a structured mapping from model primitives to observable moments.

Suppose the platform logs include impressions, positions, clicks, purchases, prices, listing attributes, and potentially group labels for sellers. Let I_{ijtc} indicate whether buyer i at time t in context c is exposed to seller j, and let Y_{ijtc} indicate whether a purchase occurs. The observed purchase probability conditional on exposure is

$$\Pr(Y_{ijtc}=1 \mid I_{ijtc}=1, x_j, p_j, c, g(j)) = \mathbb{E}[\sigma(U_{ijtc}) \mid I_{ijtc}=1, \cdot],$$

where U_{ijtc} embeds posterior beliefs and taste shifters. A naive regression of Y on group and controls conditional on exposure can be biased because exposure itself selects which sellers are shown and possibly to which buyers [21]. If the platform targets certain sellers to certain buyers, then the exposed sample is not representative. Moreover, even if conditioning on exposure, the set of competing alternatives shown to the buyer affects choice, creating a form of contextual endogeneity.

A structural approach begins by specifying a choice model that maps latent utilities to observed choices given an exposure set. If the platform shows a ranked list, the model can incorporate position-dependent attention. Let r_{ijtc} denote the rank position of seller j for buyer i at (t,c). A standard attention model multiplies utility by an examination probability $\psi(r)$, yielding an effective utility $\psi(r)U_{ijtc}$. Then the probability of purchase from seller j is

$$\Pr(a=j \mid \mathscr{L}_{itc}) = \frac{\exp(\psi(r_{ijtc})V_{ijtc})}{1+\sum_{k\in\mathscr{L}_{itc}}\exp(\psi(r_{iktc})V_{iktc})},$$

where $\mathscr{L}_{itc}$ is the set of shown sellers and V_{ijtc} is systematic utility excluding the idiosyncratic shock. Group enters through taste $\delta_i(g)$ and through beliefs μ_{ij}. Estimation of taste-based discrimination requires variation that changes the group of sellers shown while holding other attributes fixed, or instruments that shift exposure to different groups without directly affecting buyer utility.

Randomized experiments provide the cleanest variation. If the platform runs an A/B test that perturbs ranking weights, exploration rates, or fairness constraints, then exposure variation induced by random assignment can be used as an instrument. Let Z_{itc} be an experiment assignment that affects which sellers appear and in what positions. Under exclusion, Z affects purchases only through exposure and ranking, not directly through buyer tastes. Then one can identify causal effects of exposure on purchases by group and estimate whether the conversion function differs across groups holding observed features constant [22]. However, even with experiments, distinguishing taste-based from statistical discrimination requires additional structure because both enter utility in similar ways.

One approach is to model beliefs explicitly and use information variation. If the platform changes the information displayed about sellers, such as hiding photos, standardizing ratings, or altering badges, then one can estimate how posterior beliefs respond. In the Gaussian signal framework, changes that increase signal precision should increase the weight on observed signals and reduce the influence of priors. If group disparities shrink more under increased precision, this suggests a statistical discrimination component. Formally, if posterior mean is $\mu = \omega s + (1-\omega)m_g$ with $\omega = \tau^2/(\tau^2+\sigma^2)$, then increasing precision raises ω and reduces the contribution of m_g. Observing how conversion differentials respond to such changes identifies the role of group-dependent priors versus taste shifters, because taste shifters are not attenuated by signal precision in the same way.

Another approach uses buyer heterogeneity. If some buyers are plausibly group-neutral in tastes, such as buyers whose past behavior indicates no systematic preference for group after controlling for price and quality, then their behavior can be used to estimate belief-based components. In the model, taste-based discrimination corresponds to $\delta_i(g)$ varying across buyers. If one can estimate a distribution of $\delta_i(g)$ using repeated choices, then a component of group disparity can be attributed to tastes [23]. However, identification relies on assumptions about the stability of preferences and on observing enough variation in the choice sets.

Platform-mediated disparities require modeling the exposure policy. Exposure is a function of predicted value and constraints, which can be approximated by a parametric scoring function. Suppose exposure probability is monotone in a score $S_{jtc} = \phi(x_j, c; \beta) + \varepsilon_{jtc}$ with some noise from tie-breaking and exploration. If one can recover the mapping from S to exposure, then one can estimate whether the score function loads differently on features correlated with group or whether prediction errors are group-dependent. In practice, ranking models can be complex, but one can treat the policy as a black box and estimate marginal propensities: how much does exposure change when a feature changes, conditional on context. If the platform provides logged propensities for exploration, then inverse propensity weighting can recover counterfactual outcomes under alternative exposure rules.

A central challenge is selection on unobserved quality. Sellers with higher latent quality may be more likely to survive or to invest in better fulfillment, and these dynamics can differ by group. In the static framework, latent quality q_j enters buyer utility and conversion. Without observing q_j, controlling for observed proxies may not be sufficient. A structural remedy is to treat q_j as a random effect and infer it from repeat outcomes such as return rates, complaint rates, and long-run seller performance, using a state-space model [24]. For example, one can specify

$$q_{j,t+1} = q_{j,t} + \zeta_{j,t},$$

where $\zeta_{j,t}$ captures shocks, and link $q_{j,t}$ to observed post-purchase outcomes that are less affected by exposure, such as defect rates conditional on sale. This helps separate differences in conversion due to quality from differences due to discrimination, though it relies on the assumption that post-purchase outcomes are comparable across groups.

Seller price endogeneity further complicates identification. If sellers in disadvantaged groups anticipate lower exposure, they may set lower prices to compensate, affecting observed conversion. Estimating buyer discrimination without accounting for price endogeneity can confound discrimination with strategic pricing differences. Instrumenting for price using cost shifters or platform fee changes can help. In the model, seller best response satisfies a first-order condition

$$\frac{\partial \Pi_j}{\partial p_j} = T_j + (p_j - c_j)\frac{\partial T_j}{\partial p_j} = 0,$$

which implies a markup rule involving demand elasticity. If one can estimate demand elasticity from experiments or instruments, then one can recover implied marginal costs and assess whether group price differences reflect cost differences or differential market power induced by exposure.

Once primitives are estimated, one can implement the counterfactual decompositions defined earlier [25]. The key is to compute equilibrium outcomes under modified primitives, accounting for seller responses and platform re-optimization. For example, setting $\delta_i(g)$ to a constant yields counterfactual conversions, and re-solving the platform allocation under the same objective yields counterfactual exposure. Similarly, equalizing signal precision or priors yields counterfactual posteriors and conversions. Comparing these counterfactuals to the baseline provides model-based attributions of disparity. Because these exercises can be sensitive to functional form, robustness requires checking multiple specifications for buyer choice, signal formation, and platform policy.

A practical concern is that group labels may be missing or noisy, or the platform may avoid collecting them. The framework can still be applied by defining groups via observable proxies, but then disparity metrics may reflect proxy-defined partitions rather than protected classes. Alternatively, one can focus on outcome-based groupings, such as sellers with systematically lower ratings due to sparse data, which can still capture structural disadvantage mechanisms without requiring demographic labels. The interpretation shifts from discrimination against a protected class to discrimination against a segment defined by the platform's information environment, but the analytical decomposition remains similar.

This section frames identification as the task of separating three intertwined mappings: the mapping from seller characteristics to true value, the mapping from signals and group to buyer beliefs and tastes, and the mapping from predicted value to exposure. The next section integrates these elements into a welfare analysis that clarifies tradeoffs among efficiency, fairness, and incentives, and that helps interpret why certain interventions may succeed or fail in reducing disparate matching outcomes [26].

7. Welfare, Incentives, and Equilibrium Effects of Interventions

Reducing disparate outcomes is not only a matter of reallocating exposure; interventions can alter buyer surplus, seller incentives, and long-run market composition. This section evaluates interventions within the model by defining welfare objects, describing how interventions propagate through equilibrium, and characterizing conditions under which disparities decline without large distortions. The analysis remains neutral in the sense that it does not assert a single welfare criterion, but it formalizes several relevant metrics.

Let buyer i in context c obtain expected surplus from the choice set induced by exposure. Under a logit model with outside option utility normalized to zero, expected surplus is proportional to the log-sum-exp term. If the shown set is $\mathscr{L}$ with position-dependent attention, then expected surplus is

$$CS_{itc} = \frac{1}{\lambda_u}\log\left(1 + \sum_{j\in\mathscr{L}} \exp(\psi(r_{ijtc})V_{ijtc})\right),$$

where λ_u is the scale parameter of idiosyncratic shocks. Platform revenue can be modeled as commissions plus advertising. Seller surplus is profit Π_j. Total surplus can be defined as the sum of buyer surplus, seller surplus, and platform profit, with the understanding that transfers such as commissions can be netted out depending on the perspective.

An exposure parity constraint shifts the allocation away from the unconstrained optimum. In the static convex-regularized allocation, the welfare loss can be approximated using the shadow prices η_g [27]. Suppose the unconstrained solution yields group exposure shares $E_g^{(0)}$ and the constrained solution enforces $E_g^{(\mathrm{F})} = \pi_g K$. Under quadratic regularization, the welfare difference is approximately

$$W^{(0)} - W^{(\mathrm{F})} \approx \frac{\lambda_r}{2} \sum_g \left(E_g^{(0)} - \pi_g K \right)^2 \cdot \chi_g,$$

where χ_g captures the curvature of the value distribution within group g and how rapidly marginal value declines as exposure is reallocated. This expression is schematic, but it emphasizes that welfare costs are smaller when the marginal value gaps between groups are small, or when the unconstrained allocation already approximates the target shares. When prediction bias drives the unconstrained allocation away from efficient exposure, the sign of the welfare change can reverse, because the unconstrained solution may be suboptimal relative to true values.

Interventions that alter information rather than exposure can have different welfare effects. If statistical discrimination arises because buyers underweight informative signals for a group due to low precision, improving signal precision can raise both buyer surplus and seller outcomes for that group. However, information design can reduce buyer ability to match on idiosyncratic preferences if signals are overly standardized. A formal way to capture this is to decompose seller quality into a common component valued by all buyers and an idiosyncratic match component. Let $q_j = q_j^{\mathrm{common}} + q_{ij}^{\mathrm{match}}$. If the platform reveals a signal primarily about q^{common}, it can reduce uncertainty about general reliability while leaving match quality uncertain. This can increase overall conversion if reliability is a bottleneck, but it can reduce buyer surplus if match quality drives satisfaction. The welfare impact depends on which component dominates and on whether idiosyncratic match quality is correlated with group [28].

Exploration policies are a third class of interventions. Increasing exploration, such as injecting randomization into ranking, can reduce disparities arising from data sparsity by ensuring that underexposed sellers generate observations. The cost is a reduction in short-run efficiency because some impressions go to lower-predicted-value sellers. The long-run benefit is improved estimation and potentially more accurate targeting. In a discounted setting, exploration is beneficial when the value of information outweighs the immediate loss. The condition can be expressed in terms of the expected improvement in prediction quality. If the expected reduction in squared error from an additional observation is $\Delta\mathrm{MSE}$ and the value function is concave in error, then exploration is more attractive when $\Delta\mathrm{MSE}$ is larger, which is typically true for under-sampled groups. This implies that exploration targeted toward disadvantaged groups can be justified even under a purely revenue-oriented objective if it accelerates learning.

Seller incentives respond strongly to exposure rules. If fairness constraints guarantee a baseline exposure share for a group, sellers in that group may invest less in quality or service if marginal exposure becomes less sensitive to their effort. In the model, this corresponds to a reduction in $\partial E_j / \partial r_j$ induced by the constraint. Seller optimal effort satisfies [29]

$$(p_j - c_j) \frac{\partial T_j}{\partial r_j} = C_j'(r_j).$$

Since T_j depends on exposure and conversion, any intervention that flattens the exposure response reduces the marginal benefit of effort. This incentive effect can partially offset disparity reduction if lower effort reduces true quality and conversion. Conversely, if disadvantaged groups face higher marginal costs of effort due to structural constraints, a guaranteed exposure floor can improve entry and encourage investment by reducing risk. Which direction dominates depends on the curvature of C_j and on how constraints are implemented. A soft penalty on disparity may preserve marginal incentives better than a hard constraint by allowing high-performing sellers within a disadvantaged group to still gain exposure.

Market composition effects are particularly relevant. Sellers decide whether to enter based on expected profit, which depends on anticipated exposure and conversion. If a group has systematically lower exposure, marginal sellers in that group may not enter, reducing variety and potentially reinforcing buyer priors that the group has fewer high-quality sellers. An intervention that increases exposure can increase entry, which can raise competition and potentially lower prices, benefiting buyers. However, increased entry can also reduce profits for incumbent sellers and change the distribution of quality. In equilibrium, these composition effects can either reduce or increase measured disparities depending on whether new entrants are high or low quality [30]. A structural counterfactual must therefore solve for entry decisions, not only for allocations among existing sellers.

A formal way to capture entry and composition is to define a group-specific distribution of potential sellers with types τ determining costs and potential quality. Let a seller of type τ in group g have profit $\Pi(\tau, g)$ under a given policy. Entry occurs if $\Pi(\tau, g) \geq 0$. The measure of active sellers in group g is then

$$M_g = \int \mathbf{1}\{\Pi(\tau, g) \geq 0\} \, dH_g(\tau),$$

where H_g is the type distribution. Disparities in observed outcomes per seller can shift simply because the set of active sellers changes. For instance, increasing exposure to a disadvantaged group can attract lower-quality entrants, reducing average conversion even as total sales increase. Interpreting such changes requires conditioning on latent type or using welfare measures that account for entry.

Interventions can also be framed as mechanism design problems [31]. The platform chooses exposure and information to maximize an objective subject to incentive compatibility for sellers and possibly participation constraints. If sellers can manipulate observable features to appear high-quality, then information policies must be robust. Manipulation can be modeled by allowing sellers to choose a reported feature $\tilde{x}_j$ at a cost increasing in deviation from true x_j. The platform score depends on $\tilde{x}_j$, and the equilibrium includes a manipulation strategy. In such environments, fairness constraints based on $\tilde{x}_j$ can be gamed, potentially shifting benefits away from intended groups. Designing constraints in terms of outcomes that are harder to manipulate, such as verified fulfillment metrics, can reduce gaming but may increase measurement disparities if verification is easier for some groups.

Finally, disparate outcomes can be evaluated under different normative criteria. Exposure parity is one criterion, conversion parity is another, and parity conditional on latent quality is yet another. Each criterion implies different interventions and different welfare implications. For example, enforcing equal transactions per seller across groups may require allocating more exposure to lower-converting sellers, which can reduce buyer surplus if it increases search costs or reduces match quality. Enforcing parity conditional on latent quality is closer to a notion of equal treatment, but it requires strong assumptions and measurement of quality. In the model, these criteria correspond to different constraints on the mapping from q_j to outcomes.

The welfare and incentive analysis underscores that interventions operate through multiple channels: immediate reallocation of attention, long-run learning effects, seller effort and entry responses, and changes in buyer information. Disparity reduction is therefore an equilibrium problem, and evaluating policies requires specifying which outcomes matter and over what horizon [32]. The conclusion summarizes the main analytical implications and clarifies how the formal framework can be used to interpret observed disparities on two-sided e-commerce platforms.

8. Conclusion

This paper presented a formal framework for analyzing disparate outcomes in buyer–seller matching on two-sided e-commerce platforms where exposure and information are institutionally mediated through ranking and recommendation policies. The model treats exposure as an allocative control that determines which sellers enter buyers' consideration sets and combines this with buyer decision-making under imperfect information and with seller strategic responses in pricing, effort, and participation. Discrimination is represented through three distinct but interacting mechanisms: taste-based discrimination captured by group-dependent utility shifters, statistical discrimination captured by group-dependent priors and signal precision that influence posterior beliefs, and platform-mediated discrimination captured by exposure rules based on noisy predictions and nonlinear allocation. Within a static environment, disparate outcomes can arise from differences in exposure, differences in conversion conditional on exposure, and differences in how strongly exposure is correlated with conversion, with prediction error variance and bias playing central roles when allocation is convex in predicted value.

The framework also formalized dynamic feedback loops created by algorithmic learning on endogenously generated data. Under exploitative ranking and limited exploration, small initial differences in estimates or data availability can persist and, in some regimes, amplify through path dependence, yielding group-asymmetric fixed points even when underlying seller quality distributions are similar. This dynamic perspective clarifies why disparity diagnostics based solely on contemporaneous allocation may miss the role of historical data and learning rules, and why interventions that alter exploration, uncertainty handling, or information disclosure can have long-run effects that differ from their short-run impact.

The platform optimization analysis showed that fairness constraints can be represented as shadow prices that shift effective values by group, and that welfare effects depend on the accuracy of value predictions, the curvature of the allocation rule, and equilibrium seller responses. The identification discussion emphasized that attributing disparities to tastes, beliefs, or platform policy requires exogenous variation in exposure or information and careful modeling of selection on unobserved quality and of endogenous pricing and entry. Overall, the formalism provides a structured way to reason about how disparities emerge from the interaction of preferences, information, prediction, and allocation, and how design choices shift outcomes across groups while affecting efficiency, incentives, and market composition [33].

References

[1] M. D. Ekstrand and J. A. Konstan, "Recommender systems notation: Proposed common notation for teaching and research," 2 2018.

[2] F. Caraffini, A. V. Kononova, and D. Corne, "Infeasibility and structural bias in differential evolution," 4 2019.

[3] R. M. Adams, C. B. Anderson, M. Hemenway, and J. E. Wieringa, "Digital humanities and libraries and archives in religious studies," Atla Summary of Proceedings, pp. 50–56, 1 2020.

[4] J. Ni, L. Wu, X. Fan, and S. X. Yang, "Bioinspired intelligent algorithm and its applications for mobile robot control: a survey," Computational intelligence and neuroscience, vol. 2016, pp. 3810903–3810903, 12 2015.

[5] N. V. Patel, "Selection bias in two-sided e-commerce marketplaces: A framework for propensity score matching implementation," Journal of Business Intelligence Systems and Computational Social Science Applications, vol. 11, no. 5, pp. 1–20, 2021.

[6] K. A. Liyanage, C. Steward, B. A. Moffat, N. L. Opie, G. S. Rind, S. E. John, S. M. Ronayne, C. N. May, T. J. O'Brien, M. Milne, and T. J. Oxley, "Development and implementation of a corriedale ovine brain atlas for use in atlas-based segmentation," PloS one, vol. 11, pp. e0155974–, 6 2016.

[7] C. Shulby, M. D. Ferreira, R. F. de Mello, and S. M. Aluísio, "Theoretical learning guarantees applied to acoustic modeling," Journal of the Brazilian Computer Society, vol. 25, pp. 1–12, 1 2019.

[8] R. Fu, Y. Huang, and P. V. Singh, Pushing the Boundaries: Frontiers in Impactful OR/OM Research - Artificial Intelligence and Algorithmic Bias: Source, Detection, Mitigation, and Implications, pp. 39–63. INFORMS, 11 2020.

[9] A. Maas, "Gauge engineering and propagators," EPJ Web of Conferences, vol. 137, pp. 03012–, 3 2017.

[10] R. Sethi, V. R. Vatsa, and P. Chhaparwal, "Identification and mitigation of algorithmic bias through policy instruments," Zenodo (CERN European Organization for Nuclear Research), 7 2020.

[11] M. Scharkow, F. Mangold, S. Stier, and J. Breuer, "How social network sites and other online intermediaries increase exposure to news," Proceedings of the National Academy of Sciences of the United States of America, vol. 117, pp. 2761–2763, 1 2020.

[12] P. G. Fennell, Z. Zuo, and K. Lerman, "Predicting and explaining behavioral data with structured feature space decomposition," EPJ Data Science, vol. 8, pp. 1–27, 6 2019.

[13] C. Blex and T. Yasseri, "Positive algorithmic bias cannot stop fragmentation in homophilic networks," The Journal of Mathematical Sociology, vol. 46, pp. 1–18, 9 2020.

[14] C. Shao, G. L. Ciampaglia, O. Varol, K.-C. Yang, A. Flammini, and F. Menczer, "The spread of low-credibility content by social bots.," Nature communications, vol. 9, pp. 4787–4787, 11 2018.

[15] S. Dalton-Brown, "The ethics of medical ai and the physician-patient relationship.," Cambridge quarterly of healthcare ethics : CQ : the international journal of healthcare ethics committees, vol. 29, pp. 115–121, 12 2019.

[16] E. Wilson, S. Helal, C. Bull, and M. Honary, "Pervasivehealth - persuasive health: Back to the future," in Proceedings of the 13th EAI International Conference on Pervasive Computing Technologies for Healthcare, pp. 426–434, ACM, 5 2019.

[17] M. Oswald and A. Babuta, "Data analytics and algorithmic bias in policing," 9 2019.

[18] L. Xu, "Establishing a vibrant information literacy program in the absence of curriculum mandate: A case study," International Journal of Librarianship, vol. 2, pp. 84–91, 12 2017.

[19] F. Yang and M. Horning, "Reluctant to share: How third person perceptions of fake news discourage news readers from sharing "real news" on social media," Social Media + Society, vol. 6, pp. 205630512095517–, 9 2020.

[20] Q. Zhao, E. Adeli, A. Pfefferbaum, E. V. Sullivan, and K. M. Pohl, MLMI@MICCAI - Confounder-Aware Visualization of ConvNets., vol. 11861. 10 2019.

[21] S. E. Nugent, P. Jackson, S. Scott-Parker, J. Partridge, R. Raper, C. Bakalis, A. Shepherd, A. Mitra, J. Long, K. Maynard, and N. Crook, "Recruitment ai has a disability problem: questions employers should be asking to ensure fairness in recruitment," 7 2020.

[22] G. W. Hislop, D. Burdge, M. Goldweber, S. A. Rebelsky, and S. N. Weiss, "Sigcse - incorporating computing for social good in computing education," in Proceedings of the 50th ACM Technical Symposium on Computer Science Education, pp. 1247–1247, ACM, 2 2019.

[23] G. N. L. Diega, "Against the dehumanisation of decision-making-algorithmic decisions at the crossroads of intellectual property, data protection, and freedom of information," 4 2018.

[24] R. Baeza-Yates, "Recsys - bias in search and recommender systems," in Fourteenth ACM Conference on Recommender Systems, pp. 2–2, ACM, 9 2020.

[25] R. Smith, D. Ventura, and J. T. Prince, "Lc-ms alignment in theory and practice: a comprehensive algorithmic review," Briefings in bioinformatics, vol. 16, pp. 104–117, 11 2013.

[26] A. Chu, D. Squirrell, A. M. Phillips, and E. Vaghefi, "Essentials of a robust deep learning system for diabetic retinopathy screening: A systematic literature review," Journal of Ophthalmology, vol. 2020, pp. 1–11, 11 2020.

[27] Y. Yifat, N. Sule, Y. Lin, and N. F. Scherer, "Analysis and correction of errors in nanoscale particle tracking using the single-pixel interior filling function (spiff) algorithm," Scientific reports, vol. 7, pp. 16553–16553, 11 2017.

[28] An Automated Pipeline for Character and Relationship Extraction from Readers' Literary Book Reviews on Goodreads.com, ACM, 7 2020.

[29] R. Chadwick and D. I. F. Grimes, "An artificial neural network approach to multispectral rainfall estimation over africa," Journal of Hydrometeorology, vol. 13, pp. 913–931, 6 2012.

[30] A. Chouldechova and A. Roth, "A snapshot of the frontiers of fairness in machine learning," Communications of the ACM, vol. 63, pp. 82–89, 4 2020.

[31] M. Ji, X. Chen, G. Z. Genchev, T. Xu, M. Wei, and G. Yu, "Status of ai-enabled clinical decision support systems clinical implementations in china," 1 2020.

[32] R. Valentin, "Top two news words (by hour): Ten years of reduction and reflection," Media-N, vol. 13, 2 2018.

[33] A. N. Sloss and S. Gustafson, "2019 evolutionary algorithms review," 1 2019.