Simple and Robust Quality Disclosure: The Power of Quantile Partition

Shipra Agrawal[1] Yiding Feng[2] Wei Tang[3]

sa3305@columbia.edu ydfeng@ust.hk weitang@cuhk.edu.hk

[1]Columbia University [2]Hong Kong University of Science and [3]Chinese University of Hong Kong

Online platforms—from Upwork and Airbnb to Amazon and Etsy—routinely convey product or seller quality through simple, percentile-based badges such as “Top 10%” or “Guest Favorite.” These signals are remarkably stable across diverse markets, despite vast heterogeneity in quality distributions and consumer preferences. This raises a fundamental question: *Why do such simple rules work so well?* In this work, we provide a sharp theoretical justification by studying *robust quality disclosure* in a canonical marketplace setting.

We consider a platform that commits to a public disclosure policy mapping a seller's private product quality (drawn from an unknown prior) into a signal. After observing the signal, the seller sets a signal-contingent monopoly price. Buyers have heterogeneous private types and linear valuations in quality. Critically, the platform designs its disclosure rule *without knowledge* of the quality distribution, buyer type distribution, or exact valuation function. We evaluate policies via a *minimax competitive ratio*: the worst-case revenue relative to the Bayesian-optimal benchmark that knows everything.

Our main result focuses on K**-quantile partition** policies, which pool qualities whose prior quantiles fall within fixed intervals (e.g., top 1%, next 4%, etc.). We fully characterize the robustly optimal policy:

- The optimal worst-case competitive ratio Γ_K^* is the unique solution to a scalar fixed-point equation involving a K-fold composition of a simple function.
- The optimal thresholds follow a backward recursion and allocate *finer resolution to upper quantiles*—reflecting that high-quality items generate more surplus and are more sensitive to information precision.
- The ratio converges as $\Gamma_K^* = 1+\Theta(1/K)$. With just $K = 5$ signals, the platform guarantees at least **91.7%** of the optimal revenue *uniformly over all market primitives.*

In contrast, we show that alternative “simple” policies based on *fixed quality thresholds* (e.g., “quality > 0.8”) are fundamentally limited: no such policy can beat a factor-2 approximation, regardless of K. This stark separation highlights that *rank-based signaling—not absolute quality—is key to robustness* in heterogeneous markets.

Technically, we introduce a **robust disclosure design** framework that reduces the economic problem to optimizing over a convex functional space of indirect revenue functions. We prove a tight extremal representation: every feasible indirect utility is a mixture of hinge functions $\max\{c, x\}$, and the worst case is always attained at an extreme point. This enables exact analysis via convex optimization and Jensen-gap arguments.

Our results offer clear managerial guidance: adopt an *asymmetric quantile-partition disclosure with 3–5 tiers concentrated at the top*, and avoid fixed quality thresholds. This balances simplicity, interpretability, and worst-case performance—precisely the desiderata for real-world platform governance.

A full version of this paper can be found at https://arxiv.org/abs/2602.01066.