Personalised Pricing and the Price Signal

Technology makes this look like a refinement of markets: same goods, prices calibrated more precisely to each buyer. It captures the mechanics but not the economics. Widespread personalised pricing would change something more fundamental than the distribution of surplus; it would change what prices are for. A posted price is, above all, a coordinating signal: the same number, on the same shelf or webpage, visible to anyone who looks. Buyers can search across them; rivals can undercut them. Market power exercised through posted prices is, at minimum, visible.

Personalised pricing swaps out this public signal for something quite different: a private quote calculated around inferred willingness to pay, and offered to a single individual. There is indeed an efficiency argument for such a practice, and it must be addressed head on. Sometimes, it can increase output by reaching buyers that a single price would exclude. But that's contingent on whether output actually expands, and many conditions must hold for it to do so. The more likely widespread outcome of this practice isn't increased efficiency. It's the loss of meaningful price comparison, increased informational asymmetry between seller and buyer, and surplus extraction via improved profiling rather than better production.

To understand why, we need to understand what personalised pricing is, and what it is not. It is not transparent categorical discounting: student fares, senior rates and off-peak train tickets. These are open and publicly verifiable and roughly the same for all members of the category. Personalised pricing, as enabled today, is unilateral, data-driven pricing based on the seller's inference of an individual buyer's willingness to pay, derived from behavioural, demographic and contextual signals. Buyers are not offered prices that they can accept or reject, but prices that are effectively take-it-or-leave-it based on inferred willingness to pay. They have no idea which characteristics determined their price, and can't know what prices have been offered to others.

This is not the perfect price discrimination that economics textbooks describe. It is imperfect, and the asymmetry it generates will only deepen. That said, it remains distinguishable from posted prices: much more granular, much more individualised, much more opaque. What's more, the information runs one way: the seller profiles the buyer.

The most potent argument for this practice is its efficiency potential. Under a single price, the firm excludes all consumers with a willingness to pay between that price and marginal cost. If better data allowed the firm to sell to those excluded consumers at prices closer to marginal cost, without lowering prices for existing customers, total output would rise and deadweight loss would fall. There is indeed real force here, especially when marginal cost is low, as it often is in digital markets; selling one unit costs close to zero and the price needed to cover it is negligible. Yet the efficiency defence is constrained. As Varian (1985) argues, a necessary condition for efficiency gains from price discrimination is that the strategy actually increase total output. If the firm is simply segmenting existing consumers better, charging each one more precisely what they are willing to pay but selling to no one new, there is no efficiency improvement, only a transfer from consumers to the firm. Any argument for personalised pricing must therefore make a specific claim that total sales have increased, and this requires more than good data. The seller must face costs high enough to have genuinely lost demand at the single price; the data must be accurate enough to identify those consumers; the consumer base must be large enough that they represent an economically meaningful opportunity; and arbitrage frictions must be substantial enough for the firm to profitably maintain this scheme. In short, the efficiency claim is contingent on specific conditions; it is not a structural feature of the practice across all markets.

Does the actual record of real pricing regimes satisfy that condition?

Across the 14 regimes I have classified by their structural capacity to generate output expansion, only a subset satisfies the conditions the efficiency claim requires. Only in three of the 14 does unambiguous output expansion occur, and none of those conditions holds for algorithmic pricing in commercial markets: explicit segmentation along group-based price elasticity (e.g. student and senior discounts), real-time supply expansion due to a binding capacity constraint (e.g. ride-hail surge pricing, where Castillo (2025) finds total welfare rises through increased driver supply and labour allocation), and an externally mandated obligation to serve all potential customers (e.g. university financial aid). In a further three regimes, there is a theoretical case for output expansion, supported by only limited evidence, and in these cases the mechanism is not individualised, algorithmic classification based on consumer characteristics but demand-state classification, with available supply allocated across states of demand. In contrast, the conditions associated with surplus transfer to the seller, namely opaque prices, weakened comparison power and high switching costs, appear in at least 8 of the 14 regimes, and notably in those characterised by personalised algorithmic pricing: customer recognition regimes (in the UK, the Insurance renewal regime, in which the FCA found an estimated 1.2 billion pounds of potential overcharging of loyal customers each year, and the Telecom Teaser renewal regime) and opaque individualisation regimes (online retailers, enterprise software, auto

insurance). The single case where an opaque-individualisation regime produces unambiguous output growth is not in the commercial market but in a nonprofit regime for which it is mandated: university financial aid. Opaque individualisation appears to serve consumers' interests only where an externally imposed constraint on firm behaviour exists, one that commercial algorithmic pricing cannot provide.

The danger from the asymmetry is not the erosion of a pricing convention, but the impairment of a market function. An impersonal posted price serves several critical functions. It allows consumers to compare offers. The price that Buyer A sees is roughly what Buyer B sees, so search is not free, but it is not wholly opaque. It allows rivals to observe the prices they need to displace. It makes the exercise of market power at least visible as a basis for challenge and competition.

Personalised pricing undermines each of these functions. For a consumer comparing offers from multiple suppliers, the question "which is cheapest?" is replaced by "do the offers I received represent different products, or different histories or locations or whatever else I cannot see?" For an entrant, the target becomes invisible: it is forced to compete on terms it cannot observe, for customers it cannot profile as well as the incumbent. This is not a mere risk. The evidence suggests that personalised pricing already undermines price comparison in ways that may escape regulatory remedy.

My own documentary audit of consumer-facing prices for a single concert, the BTS show on 27 May 2026 at Allegiant Stadium, across five ticketing websites, found upper-level seat prices ranging from $105 to $137, a 26 per cent dispersion. The same mid-level seat was priced at $95 on Ticketmaster and $110 on StubHub. When $30 and $17 fees were added respectively, the final prices became nearly identical, $125 and $127, across platforms with different base prices and different fee structures. These differences are not merely a matter of service fees. Even after the FTC's 2025 all-in pricing mandate, Senate investigators found that Ticketmaster's service fees at twenty-six venues exceeded the processing fees they ostensibly replaced. The price signal was not being restored, because the degradation was structural, not simply the concealment of a fee. The consumer does not face a posted price; she faces a conditional quote that depends on platform choice, fee structure and the offer she receives, in a market whose architecture is designed to support personalised extraction. This is the asymmetry at the core of this question.

The efficiency arguments for personalised pricing are local in their application. They depend on a particular confluence of low marginal costs, excluded demand, the capacity to identify and reach those buyers, and limited arbitrage. Where the efficiency argument applies, it applies to the specific case. The losses are general. They arise in any market where prices cease to function as public signals. Once personalisation is available, those losses arise in any market where sellers adopt it. Price comparison is less informative where

offers are not public. Contestability is undermined wherever incumbents entrench their position with retention offers targeted at specific customers. The scope for surplus transfer is greater wherever prices are private. The efficiency story requires specific conditions to hold; the harms arise whenever those conditions do not obtain. It is the pattern that matters, not the individual cases. As the comparative evidence shows, the conditions the efficiency story requires are considerably rarer than those that enable surplus extraction. This is not to say that personalised pricing cannot improve markets. It can, and in some cases has. The question is whether the structural context in which commercial algorithmic pricing operates allows these circumstances to arise. It does not, at least not consistently.

Things get worse once the lens turns from one seller to an equilibrium with many firms personalising. In the benign access story, one seller extends discounts to reach marginal consumers. Once many firms personalise simultaneously, the strategic picture changes. If many firms profile consumers, all face an incentive to attract high-value customers from rivals via discounts in one period and extract rents from locked-in customers in later periods, when switching costs or information asymmetries make departure costly. Pricing becomes history-dependent. The fact that a buyer made a past purchase gives a firm useful information about her demand that it may use against her in the future. This incentivises consumers to manipulate their own behaviour: delaying purchases, concealing preferences, switching strategically to evade profiling.

Although competition among switchable consumers could produce lower prices, the incentive to profile for retention and extraction is likely to outweigh the incentive to expand access. The fact that a buyer would be disappointed to discover that another buyer received an identical product for less is a demand-side constraint on widespread personalising at the margin. But that single-seller story does not characterise the equilibrium. What should we expect, then? Not a world of perfectly exploitative extraction, as algorithmic pricing is far too crude for that. Nor a world of widely extended access and improved allocation, except in those limited markets where output-expansion conditions hold.

The expected result is a shift in the function that price serves. Posted prices, while noisy and imperfect, are an operational public signal: enabling comparison, keeping the exercise of market power visible, and sustaining the preconditions for entry and rivalry. Under personalisation, by contrast, price becomes a private signal, shaped by profiling, selectively disclosed and difficult to compare or contest. The effect is not primarily on allocation. It is primarily on structure. The division of surplus shifts toward those who hold the informational advantage, and that is the seller. Buyers transact against an inference about themselves they didn't supply and couldn't possibly audit. Perhaps the most striking case of opaque individualised pricing that nonetheless expands quantity is

university financial aid. The lesson is that such pricing is compatible with expanding market size only when there's an institutional commitment to access that doesn't exist in commercial algorithmic pricing regimes. The environment that widespread personalised pricing creates is one in which the public signal of posted prices, once capable of enabling comparison, entry and competitive discipline, is quietly and by design transformed to serve the seller's inference rather than the buyer's judgment.

References

- Castillo, J.C. (2025). Matching and pricing in ride hailing: Wild goose chases and how to solve them. *Management Science*.
- Financial Conduct Authority (2020). *General Insurance Pricing Practices: Final Report* (MS18/1.3). Updated December 2020.
- United States Senate, Permanent Subcommittee on Investigations (2026). *So Casually Cruel: How Ticketmaster's Monopoly Supercharges Prices and Fees*.
- Varian, H.R. (1985). Price Discrimination and Social Welfare. *American Economic Review*, 75(4), 870–875.