

Nonlinear Pricing with Underutilization: A Theory of Multi-Part Tariffs[†]

By ROBERTO CORRAO, JOEL P. FLYNN, AND KARTHIK A. SASTRY*

We study the nonlinear pricing of goods whose usage generates revenue for the seller and of which buyers can freely dispose. The optimal price schedule is a multi-part tariff, featuring tiers within which buyers pay a marginal price of zero. We apply our model to digital goods, for which advertising, data generation, and network effects make usage valuable, but monitoring legitimate usage is infeasible. Our results rationalize common pricing schemes including free products, free trials, and unlimited subscriptions. The possibility of free disposal harms producer and consumer welfare and makes both less sensitive to changes in usage-based revenue and demand. (JEL D11, D21, D42, L86, M37)

Digital goods often feature tiers of service within which marginal prices are zero. That is, they are sold according to multi-part tariffs. For example, making Google searches or browsing Facebook posts is always free; streaming certain movies from Amazon Prime’s library is free after paying for a subscription, while additional movies can be streamed for a price; and reading *Wall Street Journal* articles is free on the margin for anyone up to a trial limit, and free for paid subscribers in unlimited quantities. A common theme in each of these cases is that the seller can monetize the buyer’s time and attention via other channels, like serving advertisements, collecting valuable data, generating network effects, or addicting users. This indirect revenue is big business for the internet’s largest players. For example, Google, Facebook, and Amazon respectively made 98, 108, and 60 percent of their net profit in 2020 from advertisement.¹

*Corrao: MIT Department of Economics (email: rcorrao@mit.edu); Flynn: MIT Department of Economics (email: jpflynn@mit.edu); Sastry: Harvard Department of Economics (email: ksastry@fas.harvard.edu). Jeffrey Ely was the coeditor for this article. This paper previously circulated as “Screening with Under-Utilization: Tiers, Add-ons, and Welfare Implications.” We are grateful to Daron Acemoglu, George-Marios Angeletos, Charles Angelucci, Ian Ball, Dirk Bergemann, Alessandro Bonatti, Oğuzhan Çelebi, Glenn Ellison, Drew Fudenberg, Bob Gibbons, Giacomo Lanzani, Lucas Mayer, Stephen Morris, Marco Ottaviani, Alessandro Pavan, Nicola Rosaia, Jean Tirole, Michael Whinston, Alexander Wolitzky, Kai Hao Yang, four anonymous referees, as well as participants in the 2022 North American Summer Meeting of the Econometric Society, the 2022 Australasian Meeting of the Econometric Society, the 2022 European Summer Meeting of the Econometric Society, the MIT Theory Lunch, the MIT Industrial Organization Lunch, and the MIT Sloan Organizational Economics Lunch for helpful comments. Roberto Corrao gratefully acknowledges the financial support of the Marco Fanno scholarship.

[†]Go to <https://doi.org/10.1257/aer.20220199> to visit the article page for additional materials and author disclosure statements.

¹Net advertisement revenue figures are analyst estimates by eMarketer (eMarketer Insider Intelligence 2020), and net income is from financial statements as collected by the *Wall Street Journal*.

But the fact that indirect revenue *can* make zero-marginal-price units profitable does not explain why they are optimal. The classical theory of nonlinear pricing as screening (Mussa and Rosen 1978; Maskin and Riley 1984; Wilson 1993) predicts that sellers should use smoothly varying marginal prices, instead of tiers of zero marginal prices, to extract maximal profits.² Understanding multi-part tariffs requires an alternative economic mechanism.

In this paper, we introduce a nonlinear pricing model with the following form of noncontractibility: once the buyer purchases the right to use a good, the seller cannot fully enforce the good's utilization. The scope for underutilizing digital goods—and its potential influence on digital markets—is perhaps best illustrated by the historical failure of “pay-to-click” businesses that try to incentivize valuable usage (e.g., to pay people to see advertisements), only to be defrauded by users' simple cheating strategies (e.g., having a computer script click through a website).³ Our model captures this key issue: providers cannot contract upon legitimate, valuable usage, or otherwise prevent fraudulent, valueless usage.

Our main result is that the optimal price schedule in the presence of underutilization and usage-derived revenue is a multi-part tariff. Intuitively, noncontractibility of usage prevents sellers from charging negative marginal prices, which they would like to do to encourage valuable usage. Thus, they instead charge a price of zero on the margin. The remainder of our analysis explores the structure of multi-part tariffs and their welfare implications.

Model.—As in the classical nonlinear pricing framework, buyers differ in their demand for the product, represented by a scalar, privately known type, and have quasilinear utility in money. Higher types correspond to higher demand for the product, embedded in the familiar assumptions that preferences are convex and satisfy strict single crossing. The seller values transfers as well as usage-derived revenues from advertisement, data generation, network effects, and/or user addiction. To model noncontractibility of usage, we give buyers the ability to use less than they are allocated—for instance, if a buyer purchases the right to spend y hours on an online platform, they may choose to spend $x \leq y$ hours.

The seller chooses an arbitrary price schedule that assigns a price to each level of purchases. Buyers first decide how much to purchase, and then what to use, within the scope of what is permitted by the contract. We study the problem of how to design the price schedule optimally, taking both of the buyers' decisions into account.

Optimal Pricing.—We characterize the seller's optimal pricing when buyers can freely dispose of what they purchase. The induced levels of buyer usage in the optimum cap the seller's preferred usage (i.e., what the seller would sell were usage fully contractible) with buyers' bliss points (i.e., what buyers would use were the

²As we clarify in Section II, pooling of many buyer types, as is standard under “ironed solutions,” is unrelated to the issue of zero marginal pricing: under pooling, many buyers purchase the same amount of the good for the same price, but additional units of the good still have a strictly positive marginal price. Models with discrete buyer types or bang-bang solutions can also generate *weakly optimal* multi-part tariffs, but make the counterfactual prediction that no buyer ever consumes in the region with zero marginal prices.

³In Section IB, we discuss the case study of the pay-to-click AllAdvantage.com, and how modern legal infrastructure is designed to prevent platforms from compelling users to engage with advertisements.

product free).⁴ The corresponding price schedule is flat whenever buyers consume their bliss points, since the marginal value of additional usage to the buyer is zero. Thus, sellers price according to *multi-part tariffs*.

We next explore the structure of multi-part tariffs. We show zero marginal pricing applies to more units of the good when there are greater marginal revenues from usage (e.g., from advertising) and smaller marginal information rents, or costs of screening. Intuitively, usage-derived revenue makes higher usage more attractive, while information rents distort down the amount of usage for all but the highest type.

The *shape* of these competing effects determines where in the price schedule zero-marginal-price regions emerge. Figure 1 previews four of the pricing schemes that our model can generate. *Free pricing*, in which all units have zero marginal price (e.g., Google search or Facebook), occurs when marginal revenues from usage globally dominate marginal information rents. *Freemium pricing* or *free trial pricing*, in which initial units of the good have zero marginal price and all subsequent “add-on” units have positive marginal price (e.g., the mobile game “Candy Crush Saga”), occurs when usage-derived revenues are highly concave (e.g., because unique users generate valuable data) and overwhelm information rents for only low-usage buyers. *Premium pricing*, in which initial units of the good are sold for a strictly positive marginal price and subsequent units are sold for zero marginal price (e.g., Amazon Prime Video), occurs when usage-derived revenue dominates information rents for only high-usage buyers. Indeed, as information rents vanish for the highest-usage buyers, globally positive marginal revenue from usage always generates a “premium tier.” Thus, *hybrid pricing* schemes like the combination of a free trial and premium plan (e.g., *Wall Street Journal*) can occur with sufficiently concave and increasing revenues from usage.

Welfare Implications.—We finally use our results to understand the effects on buyer welfare of both contractible usage and changes in the structure of demand and revenue. First, while noncontractibility improves buyer welfare under a fixed price schedule, all buyers would obtain greater welfare with perfectly contractible usage under the corresponding optimal price schedule. The intuition for this result is that the lack of contractibility of usage, which enables buyers to escape being forced to use the good, also prevents them from being paid compensating differentials for usage, which leads to forgone gains from trade. Second, in the absence of contractible usage, buyer welfare is less sensitive to changes in marginal usage-derived revenue. These results echo popular claims, made especially about social media products, that users are not fairly remunerated for “being the product, not the consumer.”⁵

Related Literature.—The closest theoretical analysis to ours is by Grubb (2009), who demonstrates the optimality of three-part tariffs when selling to overconfident

⁴Formally, we show this under the technical conditions that virtual surplus is strictly single crossing in usage and agents’ types, and that virtual surplus is strictly quasiconcave in usage. In online Appendix B.1, we relax the first of these assumptions.

⁵As one example, Apple cofounder Steve Wozniak had the following to say about why he deleted his personal Facebook account: “(Facebook’s) profits are all based on the user’s info, but the users get none of the profits back (...) As they say, with Facebook, you are the product” (Gynn and McCoy 2018).

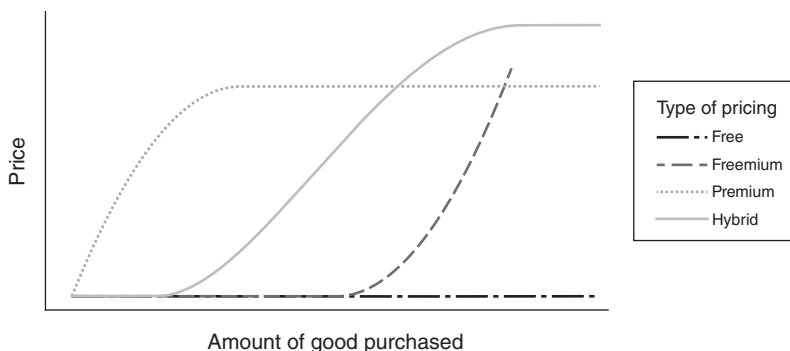


FIGURE 1. EXAMPLE MULTI-PART TARIFFS

consumers who can freely dispose of the purchased good.⁶ In online Appendix C.3, we show how the setting of Grubb (2009) can be mapped to our framework, with overconfidence mapping to a particular kind of external revenue function. By considering a richer class of external revenue functions, our analysis endogenizes a richer set of pricing schemes and is applicable to a greater number of settings.

Our analysis relates to the literature on pricing under various constraints. Sundararajan (2004) studies nonlinear pricing of information goods when the seller faces a “transaction cost” of measuring provision of the good, and derives an optimal tiered pricing schedule. Our analysis, by contrast, endogenizes multi-part tariffs as an optimal strategy in light of underutilization. Amelio and Jullien (2012) and Choi and Jeon (2021) study markets with constraints for nonnegative linear pricing and how bundling products across markets can effectively subvert such constraints. We complement this line of research by studying the *nonlinear* pricing problem, albeit in a single market. Sartori (2021) studies the provision of goods that can be freely duplicated and damaged by the seller (as opposed to being freely disposed by the buyers). The optimal allocation in this setting exhibits inefficient bunching of types, without generating multi-part tariffs.

Our results fit into a theoretical literature on mechanism design with ex post moral hazard (e.g., Laffont and Tirole 1986; Carbajal and Ely 2013; Strausz 2017; Gershkov et al. 2021; Yang 2022). However, given its focus on the possibility of underutilization, our model of ex post moral hazard has a specific structure that admits tractable analysis and, at the same time, has previously not been analyzed.

Outline.—The rest of the paper proceeds as follows. Section I introduces our model. Section II solves for optimal contracts with free disposal. Section III studies the occurrence and structure of optimal multi-part tariffs. Section IV studies the welfare implications of noncontractible usage. Section V concludes.

⁶Free disposal is also relevant for the sale of information goods, where agents may choose to optimally disregard information (see, e.g., Bergemann, Bonatti, and Smolin 2018).

I. Model

A. Consumer Demand

There is a single good that can be bought and consumed in amounts $x \in X = [0, \bar{x}]$. There is a unit measure of consumers with privately known type $\theta \in \Theta = [0, 1]$ that parameterizes their demand. The type distribution $F \in \Delta(\Theta)$ admits a density f that is bounded away from zero on Θ . For example, x might be the time that an agent spends on an online platform, and θ shifts how much they enjoy this activity.

Consumers' type-specific preferences over consumption are represented by a twice continuously differentiable utility function $u: X \times \Theta \rightarrow \mathbb{R}$. We assume that higher types value consumption more and that all types have single-peaked preferences over consumption with the following three conditions: (i) u satisfies strict single crossing in (x, θ) ;⁷ (ii) for each $x \in X$, $u(x, \cdot)$ is monotone increasing over Θ ; and (iii) for each $\theta \in \Theta$, $u(\cdot, \theta)$ is strictly quasiconcave over X . All consumer types value zero consumption the same as their outside option payoff, which we normalize to zero, or $u(0, \theta) = 0$ for all types $\theta \in \Theta$. Agents have quasilinear preferences over consumption and money $t \in \mathbb{R}$, so their total payoff is $u(x, \theta) - t$.

B. Underutilization

The primary departure of our analysis from traditional nonlinear pricing is the ability of consumers to underutilize what they buy at zero cost. A consumer buying y can consume any $x \in [0, y]$. That is, they can *freely dispose* of purchased goods.

We argue that free disposal describes feasible contracting in digital goods markets, our primary application. For example, the *Wall Street Journal* can measure if a given consumer loads an article, but not if they actually read it. Likewise, Google can register that a search has been made, but not that this is done by a human as opposed to a bot.

Perhaps the most direct evidence for the noncontractibility of consumption is provided by the failure of “pay-to-click” internet businesses. One case study is the rise and fall of AllAdvantage.com, a venture that paid users to view a permanent banner ad when browsing the internet. The *New York Times*, who interviewed the company's founder as well as eager customers, asked “Can It Pay to Surf the Web?” in a July 1, 1999, headline (Guernsey 1999). But AllAdvantage.com was quickly bogged down by users' finding simple ways to automate web surfing. This was concisely summarized in the headline of a *Wired* magazine article from July 10, 2000, which (unintentionally) answers the *Times*' original question: “It Pays to Cheat, Not to Surf” (Kang 2000). In our model's language, consumers could purchase y hours of time browsing the internet with the AllAdvantage banner, but then underutilize to consume $x \leq y$ hours, with the residual $y - x$ hours handled by bots.

⁷For avoidance of ambiguity, we mean that $u_{x\theta} > 0$, as per, for example, Nöldeke and Samuelson (2007).

C. Production and Revenues

The seller's revenue derives from two sources. The first is the total transfer from all buyers to the seller. The second is revenue that derives from consumers' usage of the good, net of production costs. This is represented by a continuously differentiable $\pi: X \times \Theta \rightarrow \mathbb{R}$.⁸ The seller values zero consumption the same as their outside option revenue from not selling the product at all, which we normalize to zero, or $\pi(0, \theta) = 0$ for all $\theta \in \Theta$.

We have four primary justifications for usage-dependent revenue π for digital goods.

Advertisement.—Digital goods are commonly bundled with revenue-generating advertisements. For example, Google search results, Facebook social feeds, and *Wall Street Journal* articles all include advertisements. In these and other online settings, advertisers can directly measure both the number of times an advertisement is loaded (*impressions*) and the number of times an advertisement is clicked. Payments from the advertiser to the platform commonly depend ex post on both impressions and clicks per impression (*click-through rate*).

We model these payments via our π as functions of platform consumption x , rather than purchases y ; a *Wall Street Journal* user must load the article containing an advertisement, and perhaps click on it, to register a payment. We argue that the metaphor also applies when human and automated usage of online platforms may be substituted for one another. The aforementioned AllAdvantage failed because advertisers refused to pay out for inauthentic, bot-derived clicks. Modern advertisement contracts, having internalized the mistakes of the AllAdvantage era, tie payouts explicitly to “valid,” human-derived clicks and impressions. As one example, the terms and conditions of Google AdSense, a popular service for adding advertisements to a web page, define invalid activity as that “solicited or generated by payment of money, false representation, or requests for end-users to click on Ads or take other actions” (Google AdSense 2020). For the erstwhile AllAdvantage.com, or any modern website monetized via AdSense, only human consumption x translates into revenue, while any bot-derived residual $y - x$ does not.

The function π is a possibly type-dependent mapping from usage to advertisement revenue. This subsumes details of consumer behavior and the advertisement contract, such as the rate with which consumers click advertisements and the payment per click.

Data Collection.—A trend in digital advertising over the last decade is the rise of *targeted advertisements* tuned toward individuals' interests as revealed by their online activity.⁹ This phenomenon has helped open up a “data economy” in which producers profit from collecting information about consumers, either directly via selling data to marketing intermediaries or indirectly via applying data toward

⁸It may be reasonable to assume that production costs depend on the *allocation* rather than consumption of the good. But if costs are monotone, it is straightforward to argue that the seller will never produce more than is consumed and “waste” the product, leading to a representation of costs in terms of usage.

⁹A report by IHS Markit estimated that, in Europe, advertising that used behavioral data comprised 86 percent of all programmatic digital advertising (IHS Markit 2017).

internal advertisement. As in the advertisement case, we use the function π to model the reduced-form “usage to revenue” schedule subsuming the translation of usage into data and the valuation of that data.

Network Effects.—Social media platforms, matching and networking services (e.g., Tinder and LinkedIn), online games (e.g., Fortnite and Candy Crush Saga), and content-streaming platforms with social rating systems (e.g., Netflix and Hulu) rely on active use to boost the appeal of their product. In online Appendix C.1, we describe how a simple model in which platform externalities generate network effects can microfound an external revenue function π by affecting all agents’ willingness to pay to participate in the platform. The framework can accommodate locally positive social externalities, as suggested by the previous examples, as well as negative externalities, due for instance to crowd-out or congestion.

Addiction.—Conventional wisdom and recent empirical evidence (Allcott, Gentzkow, and Song 2022) suggest that addicted users are a major source of demand for phone apps and social media services. In particular, assume that the quantity $x' \in X$ purchased tomorrow by each consumer is increasing in their “addictive” consumption today $x \in X$, establishing an indirect link between current consumption x and future payments t' . Our revenue function π captures this effect in reduced form. In online Appendix C.2, we illustrate how selling to myopic consumers with habit formation gives rise to an identical nonlinear pricing problem to the one we study, where π is the future revenue obtained by addicting agents today.¹⁰

D. The Nonlinear Pricing Problem

The seller’s problem is to design a total-revenue-maximizing price schedule $T: X \rightarrow \bar{\mathbb{R}}$,¹¹ where $T(y)$ is the payment from a consumer purchasing $y \in X$. Following the choice of T , each type $\theta \in \Theta$ chooses whether to buy anything, how much to purchase $\xi(\theta) \in X$, and how much to ultimately consume $\phi(\theta) \in [0, \xi(\theta)]$. As is standard, we assume that the purchase and consumption functions $\xi: \Theta \rightarrow X$ and $\phi: \Theta \rightarrow X$ are the revenue-maximizing selections from the buyers’ demand correspondence. Hence, the seller’s problem can be formulated as

$$\begin{aligned}
 (1) \quad & \sup_{\phi, \xi, T} \int_{\Theta} [\pi(\phi(\theta), \theta) + T(\xi(\theta))] dF(\theta) \\
 (O) \quad & \text{s.t. } \phi(\theta) \in \arg \max_{x \in [0, \xi(\theta)]} u(x, \theta), \text{ for all } \theta \in \Theta, \\
 (IC) \quad & \xi(\theta) \in \arg \max_{y \in X} \left\{ \max_{x \in [0, y]} u(x, \theta) - T(y) \right\}, \text{ for all } \theta \in \Theta, \\
 (IR) \quad & u(\phi(\theta), \theta) - T(\xi(\theta)) \geq 0, \text{ for all } \theta \in \Theta.
 \end{aligned}$$

¹⁰ Our interpretations of our welfare results do not hold in this case, as π captures future consumption.

¹¹ We use $\bar{\mathbb{R}}$ to mean $\mathbb{R} \cup \{-\infty, \infty\}$.

The first constraint (O), or obedience, establishes that each consumer θ chooses their optimal level of consumption $\phi(\theta)$ by optimally underutilizing their initial purchase $\xi(\theta)$. The second constraint (IC), or incentive compatibility, embodies the consumers' optimal purchase $\xi(\theta)$, taking into account their subsequent ability to underutilize. The final constraint (IR), or individual rationality, ensures that all consumers are willing to participate.¹²

II. Optimal Pricing

We first characterize the solutions of the seller's pricing problem (equation (1)). We then use a simple closed-form example, inspired by our digital goods applications, to both illustrate this result and show that it can imply the optimality of a multi-part tariff, a price schedule featuring units sold for a marginal price of zero.

A. Characterization of Optimal Pricing

We first define some important objects in which the optimal pricing schedule will be expressed. We define the consumer-optimal consumption of type θ as the function $\phi^A: \Theta \rightarrow X$:

$$(2) \quad \phi^A(\theta) = \arg \max_{x \in X} u(x, \theta).$$

This is unique and increasing because of the strict quasiconcavity of $u(\cdot, \theta)$ in x for all $\theta \in \Theta$ and strict single crossing of u in (x, θ) . In a digital platform example, $\phi^A(\theta)$ is the amount of time that type θ would optimally spend on the platform were it freely available.

Second, we define the virtual surplus function $J: X \times \Theta \rightarrow \mathbb{R}$:

$$(3) \quad J(x, \theta) = \pi(x, \theta) + u(x, \theta) - \frac{1 - F(\theta)}{f(\theta)} u_\theta(x, \theta).$$

This is the total surplus $\pi + u$, net of the information rents required to ensure local incentive compatibility. For the remaining analysis, we will assume that the function J satisfies strict single crossing in (x, θ) and is strictly quasiconcave in x . These standard technical assumptions guarantee that virtual surplus has a unique maximum and is maximized pointwise under the optimal contract, thereby ruling out cases with bunching, or multiple agent types' consuming the same bundle.¹³ Therefore, there is a unique, increasing producer-optimal consumption level $\phi^P: \Theta \rightarrow X$:

$$(4) \quad \phi^P(\theta) = \arg \max_{x \in X} J(x, \theta).$$

¹²Ensuring participation of all types is without loss of optimality for the seller owing to their ability to sell nothing and charge a price of zero, given the outside option of zero.

¹³In online Appendix B.1, we relax the single crossing assumption on J , strengthen strict quasiconcavity to strict concavity, and show how to adapt our analysis to settings which feature canonical bunching.

Continuing our example, a revenue-maximizing seller would induce a type θ consumer to spend $\phi^P(\theta)$ time on the platform were usage perfectly contractible.

With these definitions in hand, we state our result describing optimal pricing:

PROPOSITION 1 (Optimal Pricing): *In any optimal contract, consumption is given by*

$$(5) \quad \phi^* = \min\{\phi^P, \phi^A\}.$$

The optimal price schedule for all $x \in [\phi^(0), \phi^*(1)]$ is uniquely given by¹⁴*

$$(6) \quad T^*(x) = u(\phi^*(0), 0) + \int_{\phi^*(0)}^x u_x(z, \phi^{*-1}(z)) dz.$$

Moreover, purchases are part of an optimal contract if and only if they are a selection from the correspondence $\Xi_{\phi^}: \Theta \rightrightarrows X$:¹⁵*

$$(7) \quad \Xi_{\phi^*}(\theta) = \begin{cases} \{\phi^*(\theta)\}, & \text{if } \phi^*(\theta) < \phi^A(\theta); \\ [\phi^A(\theta), \inf_{\theta' \in [\theta, 1]} \{\phi^*(\theta') : \phi^*(\theta') < \phi^A(\theta')\}], & \text{if } \phi^*(\theta) = \phi^A(\theta). \end{cases}$$

The proofs of this and all other results are in online Appendix A. The first part of this result says that consumption in any optimal contract is the producer-optimal consumption level capped by the consumer-optimal consumption level. To understand this result, first observe that forcing a buyer to consume beyond their bliss point level violates obedience—they would avail of free disposal to consume less and reach their bliss point—and is therefore impossible. It is therefore *necessary* that $\phi \leq \phi^A$. We moreover show that $\phi \leq \phi^A$, combined with monotonicity of eventual consumption (also necessary for incentive compatibility by standard arguments), is *sufficient* for obedience and incentive compatibility. Strict quasiconcavity of virtual surplus then implies that simply capping the optimum in the absence of free disposal with the bliss point is optimal.

Our argument generalizes and formalizes the following scenario. A digital platform like Twitter might want users to spend all 24 hours of the day reading content and seeing advertisements. But no contract could support this outcome—users would always prefer to use cheating software to mimic 24-hour usage, while reducing their actual consumption to their bliss point. This was exactly the problem that led to the downfall of AllAdvantage.com. A natural next choice is to design incentives to induce a feasible level of Twitter consumption (i.e., $\phi \leq \phi^A$) that maximizes revenue as well as possible given this constraint. When revenue is hump-shaped (i.e., strictly quasiconcave), this is achieved by enveloping over the bliss points ϕ^A and the revenue-maximizing points ϕ^P . This is what the optimal contract specifies.

¹⁴We define the inverse of a continuous consumption function $\phi: \Theta \rightarrow X$ on $[\phi^*(0), \phi^*(1)]$ as the standard inverse on $(\phi^*(0), \phi^*(1))$, extended to the boundaries with $\phi^{-1}(\phi^*(0)) = \lim_{x \rightarrow +\phi^*(0)} \phi^{-1}(x)$ and $\phi^{-1}(\phi^*(1)) = \lim_{x \rightarrow -\phi^*(1)} \phi^{-1}(x)$. When $\phi^*(0) = \phi^*(1)$, we set $\phi^{-1}(\phi^*(0)) = 0$.

¹⁵With the convention that if the set over which the infimum is taken is empty, then the infimum is equal to the supremum of the codomain of the relevant objective function. For example, ∞ for \mathbb{R} , and 1 for $[0, 1]$.

The second part of the result derives an optimal price schedule that supports the optimum, using standard arguments that invoke the necessity of local incentive compatibility constraints. This price schedule is uniquely pinned down on the space of consumed amounts, $X^* = [\phi^*(0), \phi^*(1)]$, the image of Θ under the optimal consumption function. Away from this interval, there are many available options. For example, the seller can always offer a monotone price schedule by charging $T(\phi^*(0))$ for all $x < \phi^*(0)$ and setting an infinite price for all $x > \phi^*(1)$.¹⁶ For the remainder of the analysis, we restrict attention to $T^*: X^* \rightarrow \mathbb{R}$ and study its properties in detail.

The third part of the result characterizes the set of type-specific purchases that support the optimum. This always includes simply setting purchases equal to consumption. However, when the constraint of free disposal binds, there are many possible levels of purchases which support the optimal allocation and yield the same total revenue for the seller. More precisely, for any interval of types for whom the constraint of underutilization is binding, it is possible to have each type purchase any amount between their own bliss point and the bliss point of the highest type in that interval. An important implication of this multiplicity is that the seller may find it optimal not to offer every level of purchases. Concretely, a possible optimal solution for the seller is to offer only the maximum level of purchases for all types. Under this solution, the seller offers a tiered menu, which features discrete levels of provision, a property we explore in more detail in Section III.

B. A Digital Pricing Example

We provide a more specialized intuition for the form of the optimal contract and foreshadow its pricing implications in a closed-form example.

EXAMPLE 1 (Digital Platform with Advertisement): *A digital platform sells access time $x \in [0, 1]$. Consumers have quadratic payoffs*

$$(8) \quad u(x, \theta) = \theta x - \frac{x^2}{2},$$

where θ is uniformly distributed on $[0, 1]$. A consumer who spends time x on the platform clicks on $x(k - hx)$ advertisements, for $k > 0$ and $h \in [0, k]$. The assumption $h \geq 0$ embodies user fatigue from seeing the same advertisements repeatedly, and $h \leq k$ ensures that total clicks remain positive. Each click yields revenue $p > 0$ for the seller. Moreover, serving time x to the consumer has a constant marginal cost c . The seller's revenue function is therefore

$$(9) \quad \pi(x, \theta) = \left(\frac{\text{Revenue}}{\text{Click}} \cdot \frac{\text{Clicks}}{\text{Time}} - \frac{\text{Cost}}{\text{Time}} \right) \cdot \text{Time} = \alpha x - \frac{\beta}{2} x^2,$$

where we define $\alpha = pk - c$ and $\beta = 2ph$.

¹⁶In Section IIIC, we also give sufficient conditions for optimal price schedules to be flat for $x \geq \phi^*(1)$, capturing the idea of unlimited subscriptions.

We can use Proposition 1 to solve for the seller's optimal pricing in this setting. The consumer-optimal and seller-optimal consumption levels, which respectively maximize payoffs u and virtual surplus J , are

$$(10) \quad \phi^A(\theta) = \theta \quad \text{and} \quad \phi^P(\theta) = \max \left\{ 0, \min \left\{ 1, \frac{\alpha + 2\theta - 1}{\beta + 1} \right\} \right\}.$$

There are two possibilities for the shape of $\phi^* = \min\{\phi^A, \phi^P\}$. If $\beta \geq 1$, corresponding to high concavity of the revenue function (e.g., high user fatigue h), ϕ^A is optimal for sufficiently low types (if any). If $\beta < 1$, corresponding to low concavity of the revenue function (e.g., low user fatigue h), ϕ^A is optimal for sufficiently high types (if any). To most simply illustrate the result and its implications, we restrict attention to when $\alpha \leq 1$ and $\beta < 1$, in which case optimal consumption is

$$(11) \quad \phi^*(\theta) = \min\{\phi^A(\theta), \phi^P(\theta)\} = \begin{cases} \phi^P(\theta), & \text{if } \theta < \frac{1-\alpha}{1-\beta}, \\ \phi^A(\theta), & \text{if } \theta \geq \frac{1-\alpha}{1-\beta}. \end{cases}$$

To derive the optimal price schedule, we use the integral expression for prices (equation (6)) and simplify to obtain

$$(12) \quad T^*(x) = \begin{cases} \frac{1-\alpha}{2}x - \frac{1-\beta}{4}x^2, & \text{if } x < \frac{1-\alpha}{1-\beta}, \\ \frac{(1-\alpha)^2}{4(1-\beta)}, & \text{if } x \geq \frac{1-\alpha}{1-\beta}. \end{cases}$$

We illustrate the properties of this optimal price schedule in Figure 2 when $\alpha = 1/2$ and $\beta = 0$. We show $\phi^*(\theta)$ in the leftmost panel and $T^*(x)$ in the rightmost panel with a solid line. In the middle panel, we illustrate the purchase correspondence defined in equation (7) as the shaded region. Under the optimal contract, the good is sold for a strictly positive marginal price until $x = 1/2$, after which it is sold at a marginal price of zero. As a result, it is as if the seller offers buyers an unlimited subscription for the good at a fixed price, but allows them to buy less than this unlimited level for a discount.

We now contrast these predictions with those in a variant setting with the same demand and external revenues, but no potential for free disposal (i.e., the legitimate time spent on the platform is perfectly contractible). The seller optimally sets $\phi^*(\theta) = \phi^P(\theta)$. We illustrate, in the same case of $\alpha = 1/2$ and $\beta = 0$, the consumption, purchases, and prices in the dashed lines of Figure 2. In this case, optimal prices are nonmonotone. Starkly, $x = 0$ and $x = 1$ are both sold for free. This demonstrates the incentives for sellers in digital markets to use negative prices to induce valuable usage, which would be enforceable absent free disposal.

III. The Occurrence and Structure of Multi-part Tariffs

Having shown that multi-part tariffs can be optimal in an example, we now provide general conditions under which the optimal price schedule is a multi-part tariff. These conditions imply that the price schedule must be a multi-part tariff whenever

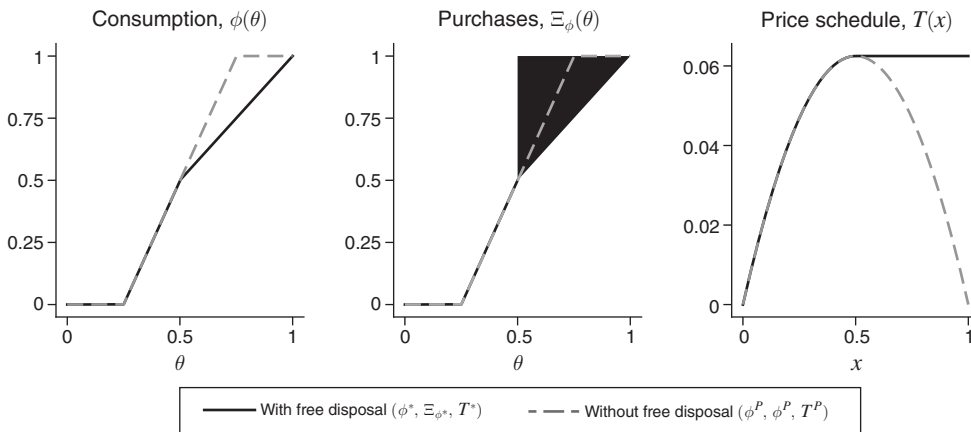


FIGURE 2. OPTIMAL CONTRACTS WITH AND WITHOUT FREE DISPOSAL IN EXAMPLE 1

marginal revenues from usage are strictly positive. We further provide sufficient conditions for optimal multi-part pricing to reduce to the freemium, premium and fixed cost pricing plans often observed in practice, and argue that these conditions realistically describe products with these pricing schemes.

A. The Occurrence of Multi-part Tariffs

We first generally characterize when the optimal price schedule is a multi-part tariff. To this end, we now formally define flatness of price schedules and multi-part tariffs. A price schedule T is *flat* at x if there exists a neighborhood $O(x)$ such that T is constant on $O(x)$.¹⁷ A price schedule T is a *multi-part tariff* if it is flat at some x . This definition builds on the conventional definition of a three-part tariff in which there is a fixed cost, an initial allotment of zero-marginal-price goods, and then additional units with positive marginal price.¹⁸ Our more general definition can account for zero-marginal-pricing in multiple separate tiers and/or away from the “bottom.”

To characterize the optimality of multi-part tariffs, we define the constrained marginal revenue function $H: X^* \rightarrow \mathbb{R}$ that maps outcomes to the marginal revenue of the seller for the type who most prefers that outcome:

$$(13) \quad H(x) = J_x(x, (\phi^A)^{-1}(x)).$$

¹⁷Flatness of prices is more demanding than zero slope, or $T'(x) = 0$ wherever T' is defined, as it must apply on a neighborhood around x . Thus, zero marginal transfers or tariffs “at the top,” meaning a specific maximal point in the type or action space, do not imply flatness by our definition—no unit of size $\varepsilon > 0$ is sold for zero price.

¹⁸A two-part tariff, via the conventional definition, combines fixed costs with positive marginal costs. This of course can be accommodated in the conventional nonlinear pricing framework, with zero tiers, as the “intercept” of the tariff is a free parameter.

To interpret H , imagine that the seller had to offer the product for free. The sign of $H(x)$ determines whether the seller would profit from more ($H(x) > 0$) or less ($H(x) < 0$) consumption from the type whose favorite consumption is x . The following result links flat pricing with this trade-off.¹⁹

PROPOSITION 2 (Multi-part Tariffs): *If $H(x) > 0$, then the optimal price schedule T^* is flat at x and, therefore, a multi-part tariff. Conversely, if the optimal price schedule T^* is a multi-part tariff that is flat at x , then $H(x) \geq 0$.*

To understand the intuition for this result, it is useful to rewrite the sufficient condition for flat pricing ($H(x) > 0$) as a comparison of two dueling economic forces:

$$(14) \quad \underbrace{f(\theta)\pi_x(x, \theta)}_{\text{marginal revenue}} > \underbrace{(1 - F(\theta))u_{x\theta}(x, \theta)}_{\text{marginal information rent}}$$

where $\theta = (\phi^A)^{-1}(x)$. The first force (on the left-hand side) is the total marginal revenue from additional usage. Increasing π_x (e.g., from more intensive advertising) boosts this force. The second force (on the right-hand side) is the total marginal information rent paid to all higher-type consumers. Increasing complementarity or increasing F in the hazard rate order enlarges these information rents.

When marginal revenue dominates marginal information rents, the seller would like to charge a negative marginal price to induce valuable usage. However, the constraint of free disposal makes this impossible. Thus, they do the next best thing and charge a price of zero on the margin. Conversely, under the opposite inequality, or $H(x) < 0$, marginal information rents dominate marginal revenues and the seller wishes for the buyer to consume less than their satiation point. Thus, free disposal poses no constraint, and they charge a positive marginal price to extract the buyer's (strictly positive) willingness to pay.

This result has two immediate implications for the possibility of multi-part tariffs. First, if marginal revenues from usage are everywhere negative ($\pi_x \leq 0$), multi-part tariffs are impossible. Thus, in situations with large production costs (e.g., physical goods), multi-part tariffs are unlikely to arise. Second, if marginal revenues from usage are everywhere strictly positive ($\pi_x > 0$), then the optimal price schedule *must* be a multi-part tariff. This claim follows because marginal information rents necessarily vanish for the highest types of agents, while marginal revenues (by assumption) are always strictly positive. Thus, in settings with low (marginal) production costs and valuable revenue from usage (e.g., digital goods), multi-part tariffs are likely to be optimal.

In the next section, we investigate the structure of multi-part tariffs. Before so doing, we first concretely illustrate why zero marginal pricing obtained in the case

¹⁹Recall the definition of the inverse of a continuous consumption function from footnote 14 in defining $(\phi^A)^{-1}$ on X^* . Moreover, it is not possible to strengthen the claim that if optimal price schedule T is flat at $x \in X^*$, then $H(x) \geq 0$ and make the inequality strict. The proof of Proposition 2 provides an explicit counterexample.

of Example 1 plotted in Figure 2. We also highlight how standard models cannot generate multi-part tariffs in two remarks.

EXAMPLE 1 (Continued): *The constrained marginal revenue function is $H(x) = (\alpha - \beta x) - (1 - x)$, where the first term is the marginal revenue from additional usage and the second is the marginal increase of information rents. The presence of zero marginal pricing therefore relies on there being sufficiently high marginal advertising revenues (high α and low β), which derives in the underlying advertising model from a higher revenue per click or lower user fatigue. In the case plotted in Figure 2, we set $\alpha = 1/2$ and $\beta = 0$. Thus, zero marginal pricing held when $H(x) > 0$ or $x > 1/2$.*

REMARK 1 (Standard Models Do Not Generate Multi-part Tariffs): *Strictly optimal multi-part tariffs are not possible in the canonical Mussa and Rosen (1978) and Wilson (1993) screening models with a (convex) continuum of agent types.²⁰ A weakly optimal multi-part tariff is possible, for instance, in a specialization with a discrete number of types, by extending the domain of the offered menu with a constant price schedule. But we argue this is not economically meaningful, or relevant for our applications, for three reasons. First, no type would ever consume any outcome in the extended menu that does not lie in the initial menu, therefore ruling out variability of consumption within a pricing tier. This is clearly counterfactual—within the single tier of a zero-price product like Facebook, there is large variability in time spent on the platform. Second, there would be as many parts to the price schedule as unique consumer types, which is an arbitrary choice of the modeler. This prevents meaningful comparative statics for the number of observed tiers as a function of primitives. Third, there is no principled reason for arguing that the multi-part tariff is the “right” selection from the set of optimal price schedules which are extended off-menu.*

REMARK 2 (Bunching Is Unrelated to Multi-part Tariffs): *Optimal bunching in standard screening models is a different phenomenon from optimal multi-part tariffs. Intuitively, bunching is a feature that occurs in the type space Θ , where many different types buy the same amount. However, all units of the good are still sold at a strictly positive marginal price and so the optimal price schedule is never flat. In online Appendix B.1, we solve for the optimal contract with bunching when J does not satisfy single crossing in (x, θ) and is strictly concave in x by adapting the method of Nöldeke and Samuelson (2007). Example 1 provides an explicit illustration of how bunching is unrelated to the issue of zero marginal pricing: multiple buyers bunch on buying nothing, but there are still strictly positive marginal prices for the first marginal units of the good (see Figure 2).*

²⁰ Even if the virtual surplus is not strictly concave and a multi-part tariff is weakly optimal, the same argument of this remark would apply.

B. Rationalizing Simple Pricing Schemes

We now leverage our characterization of multi-part tariffs to make theoretical predictions about the structure of pricing in various applications. We moreover argue that these predictions line up with various forms of pricing that we observe in practice. We first define four common pricing schemes:

1. *Regular pricing*, in which all units of the good have strictly positive marginal price.
2. *Fixed pricing*, in which all units have the same price.
3. *Premium-tier pricing*, in which initial units have positive price until some point, after which all subsequent units have zero marginal price.
4. *Introductory-offer pricing*, in which initial units have zero price and subsequent units have positive price.

The last three of these are difficult to understand through the classical lens of non-linear pricing (see Remark 1). We now provide sufficient conditions, based solely on the case of our model where H crosses zero at most once, that delineate these pricing schemes in our model. They all immediately follow from Proposition 2.

COROLLARY 1 (Simple Pricing Schemes): *The following statements are true:*

1. *If $H(x) < 0$ for all $x \in X^*$, then optimal pricing is regular.*
2. *If $H(x) \geq 0$ for all $x \in X^*$, then optimal pricing is fixed.*
3. *If $H(x) < 0$ if and only if $x < \hat{x} \in \text{int}(X^*)$, then optimal pricing is of the premium-tier form, with threshold for strictly positive marginal prices given by \hat{x} .*
4. *If $H(x) \geq 0$ if and only if $x \leq \hat{x} \in \text{int}(X^*)$, then optimal pricing is of the introductory-offer form, with threshold for zero marginal prices given by \hat{x} .*

In several applications of this result, we now describe the model's theoretical predictions for pricing schemes based on plausible structures of external revenues. In each case, we argue that the model's predicted pricing scheme lines up with that which we observe in practice.

Case 1 (Regular Pricing for Physical Goods): Physical goods typically have significant production and/or transportation costs that swamp potential usage-based revenues (e.g., from word-of-mouth advertising). Seen through the lens of our model, this corresponds to a case with $\pi_x < 0$ and implies that $H < 0$. As per Corollary 1, our model therefore predicts that physical goods should feature regular pricing. This case is studied in the classical nonlinear pricing literature (Mussa and Rosen 1978; Wilson 1993) and of course matches the reality of how the majority of physical goods are priced.

Case 2 (Fixed Pricing for Search Engines and Social Media): Search engines and social media platforms, as we have motivated, derive large revenues from advertisement and data collection. Advertisement, in particular, often appears at a uniform rate during regular usage (e.g., sponsored search results in Google) and derives revenue per impression or click (i.e., marginal usage). Mapped to our model, $\pi_x > 0$ for all levels of usage and user types. Our model predicts that, if these revenues globally dominate information rents, then $H > 0$ and optimal pricing is fixed. Moreover, under the assumption that the lowest-type consumers would most prefer to consume nothing (e.g., make zero Google searches or spend zero minutes on Facebook), this fixed price is zero. This matches, of course, the observed pricing scheme of products like Google, Facebook, Twitter, and Instagram.

Case 3 (Premium-Tier Pricing for Content-Streaming): Content-streaming platforms, such as Amazon Prime, do not run external advertisements but instead use consumer data to fine-tune within-platform content recommendations. We might conjecture that the indirect revenues derived from this model are positive, but smaller than the sum of direct revenues earned by search engines and social media platforms from serving advertisements and selling data. Under our model, this corresponds to a case where $\pi_x > 0$, so $H(x) > 0$ for high levels of usage (where information rents are necessarily small, as we will formally clarify in Corollary 2), but $H(x) < 0$ for low levels of usage (where standard screening concerns dominate). Our model therefore predicts that marginal prices should initially be positive and free, in unlimited quantities, for high enough usage. This prediction matches Amazon Prime's streaming library pricing, in which users can buy individual movies or shows for fixed prices or purchase the unlimited Prime subscription.

Case 4 (Introductory-Offer Pricing for Cloud Storage, Mobile Games, and Networking Services): A diverse class of goods derive a disproportionate amount of revenue from initial usage. For example, cloud storage platforms such as Dropbox and iCloud may benefit from a network externality, whereby all users find the product more valuable when they can share files with a larger number of unique people. A similar force is present for networking services platforms, such as LinkedIn. As a further example, mobile games (e.g., "Candy Crush Saga") can directly generate revenue from merely being opened, as they can scrape data points like the user's location and sell them to advertising firms. They are also likely to be addictive and feature diminishing marginal effects of past usage on future willingness to pay (i.e., the first minute of the game hooks the user more than the hundredth).²¹ These settings can be represented in our model with marginal revenues that are diminishing ($\pi_{xx} < 0$), and potentially negative for high levels of usage. Because of this, $H(x) > 0$ is likely for low levels of usage (when marginal revenues are especially high) and $H(x) < 0$ is likely for high levels of usage (when marginal revenues are diminished). Thus, our model predicts an introductory offer of marginally free units. Moreover, if the lowest-type buyers would most prefer to consume nothing, then the

²¹ Allcott, Gentzkow, and Song (2022) show that a model of total demand for popular phone apps, calibrated to empirical evidence, is consistent with habit formation, myopia, and diminishing returns.

introductory offer has no fixed cost either (i.e., it is a free trial). Dropbox features a free trial with limited space allotment, which can be increased with paid upgrades; LinkedIn has a free base version, which can be upgraded to a premium mode with more features; and Candy Crush Saga gives the player a free allotment of lives, which can be replenished at a cost.

A Numerical Example.—We finally use Example 1 to illustrate these four pricing cases numerically, under different assumptions for the shape of the revenue function.

EXAMPLE 1 (Continued): *The constrained marginal revenue function is $H(x) = (\alpha - \beta x) - (1 - x)$, which crosses zero at most once. The four cases of Corollary 1 therefore summarize the possibilities for pricing. In Figure 3, we illustrate each of these cases. We plot the profit functions $\pi(x)$, constrained marginal revenue functions $H(x)$, and tariffs $T(x)$. Case 1 arises when marginal revenues from usage are sufficiently low or negative ($\alpha \leq 1$ and $\alpha \leq \beta$) as in our physical goods application. Case 2 occurs when marginal revenues from usage are always sufficiently high ($\alpha > 1$ and $\alpha > \beta$), as in our search engine and social media platforms applications. Case 3 obtains when marginal revenues from usage are intermediate and are dominated by information rents for low levels of usage ($\alpha \leq 1$ and $\alpha > \beta$), as in our content-streaming platforms applications. Case 4 occurs when there are initial high revenues from usage that then diminish sharply ($\alpha > 1$ and $\alpha \leq \beta$), as in our cloud storage, mobile games, and networking services applications.*

C. More General Pricing Schemes: Unlimited Subscriptions, Trials, and Arbitrary-Part Tariffs

Our model also allows for richer pricing schemes that feature multiple regions of zero-marginal-pricing for various levels of usage. We first establish general sufficient conditions for the occurrence of zero-marginal-pricing at the bottom and top of the pricing schedule and link this to the co-occurrence of free trials and unlimited subscriptions for many goods. Finally, we illustrate in a stylized example how the model can be used to rationalize any number of regions of zero marginal pricing.

Unlimited Subscriptions and Trials.—A price schedule T features an *unlimited subscription* if it is flat at $\phi^*(1)$, the highest level of consumption under the optimal contract, and a *trial* if it is flat at $\phi^*(0)$, the lowest level of consumption under the optimal contract. These features were respectively present under the premium-tier pricing and introductory-offer pricing introduced in the previous subsection. We now derive general conditions under which unlimited subscriptions and trials are obtained in more general, possibly multitier pricing schedules.

COROLLARY 2 (Unlimited Subscriptions and Trials): *The price schedule features an unlimited subscription if*

$$(15) \quad \pi_x(\phi^A(1), 1) > 0.$$

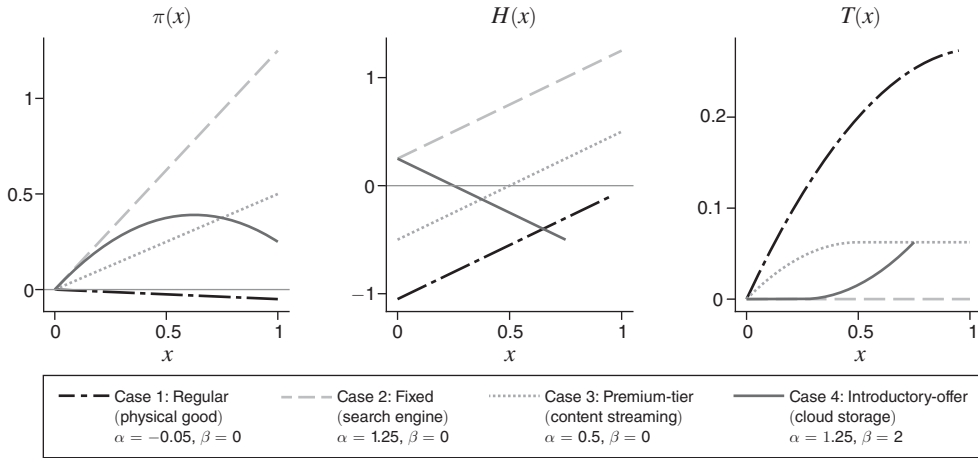


FIGURE 3. CUTOFF PRICE SCHEDULES IN EXAMPLE 1

The price schedule features a trial if

$$(16) \quad f(0)\pi_x(\phi^A(0), 0) > u_{x\theta}(\phi^A(0), 0).$$

The first part shows that an unlimited subscription is optimal whenever marginal revenues from usage are positive when the highest type uses at their bliss point level. As we previously discussed, a more demanding sufficient condition for this is that marginal revenues from usage are globally positive, as in our search engines, social media platforms, and content streaming applications. The lack of a countervailing force from marginal information rents reflects the fact that the seller does not distort allocations of the highest-type agents away from the first-best, surplus-maximizing allocation.

The second part shows that trial tiers are optimal whenever total marginal revenues from usage when the lowest type uses at their bliss point level exceed the marginal information rent paid to *all* higher types. As we have highlighted, with zero or negative marginal revenues, this condition would never hold. Otherwise, it is more likely to hold in environments with high marginal revenues stemming from *low* levels of usage, as in our previous application to cloud storage, mobile games, and networking services.

These two conditions are mutually compatible. Thus, Corollary 2 opens the door to pricing schemes that feature both trials and unlimited subscriptions. We now provide a concrete example of this possibility and apply it to understand the pricing of online newspapers, such as the *Wall Street Journal* (*wsj.com*).

EXAMPLE 2 (Optimal Pricing with Unlimited Subscriptions and Trials): *Consumer preferences, the outcome space, and the type distribution are identical to those in Example 1. Advertisements are served at a constant rate, normalized to one. Consumers notice an advertisement according to an exponential process with hazard rate $\lambda > 0$, click on the first advertisement they notice, and ignore all subsequent*

advertisements. The seller earns α per click. The seller therefore receives the following payoff in expectation:

$$(17) \quad \pi(x, \theta) = \frac{\text{Revenue}}{\text{Click}} \cdot \text{Expected Clicks} = \alpha(1 - e^{-\lambda x}).$$

Thus, $H(x) = \lambda\alpha e^{-\lambda x} - (1 - x)$, which has, depending on the values of (λ, α) , either zero, one, or two tiers. The last case features both a trial and an unlimited subscription. In Figure 4, we illustrate all three possibilities.

The two-tier example, plotted with $\lambda = 2.5$ and $\alpha = 0.5$, satisfies both conditions in Corollary 2. First, the fact that marginal revenues remain positive for all levels of consumption guarantees an unlimited subscription, as this force dominates the vanishing information rents. Second, marginal revenues are high for low consumption because these agents are least likely to have already clicked an advertisement. In the two-tier calibration, this force dominates the marginal information rents that need to be paid to all higher-type agents. Moreover, because the lowest type most prefers zero consumption, the trial tier has zero price, and is therefore a “free trial.”

This combination of a trial and an unlimited subscription is characteristic of online newspaper and content streaming platforms. For example, the *Wall Street Journal* features an initial allowance of free articles before charging a subscription fee for access to unlimited numbers of articles. This can be understood through the lens of our example as the optimal selling response to a situation in which readers’ clicks generate revenue, but readers are unlikely to continue to click on advertisements multiple times.

Arbitrary-Part Tariffs.—In general, our results can be used to rationalize tariffs with an arbitrarily large number of flat regions. That is, if we define an $N + 2$ -part tariff as a price schedule with N flat-pricing intervals (in analogy with the definition of a three-part tariff as having one tier), our model can generate any $N + 2$ -part tariff for $N \in \mathbb{N}$.²² To do this, one may simply construct a constrained marginal revenue function H that crosses zero the desired number of times. We show this constructively in the following (intentionally contrived) example.

EXAMPLE 3 (Multi-part Tariffs Can Have Arbitrarily Many Parts): *Consumer preferences, the outcome space, and the type distribution are identical to those in Example 1. The seller has a revenue function that increases in the usage of each type in proportion to how low that agent’s type is (e.g., they derive particular value from targeting an advertisement at the most marginal users). Moreover, the value of attention paid by all types slightly oscillates over time. We capture this with the profit function:*

$$(18) \quad \pi(x, \theta) = x(1 - \theta) - \frac{k}{2\pi\omega} [\cos(2\pi\omega x) - 1],$$

²²Formally speaking, we refer to a flat-pricing interval as a maximal interval $I \subset X^*$ such that the price schedule is flat at all $x \in I$.

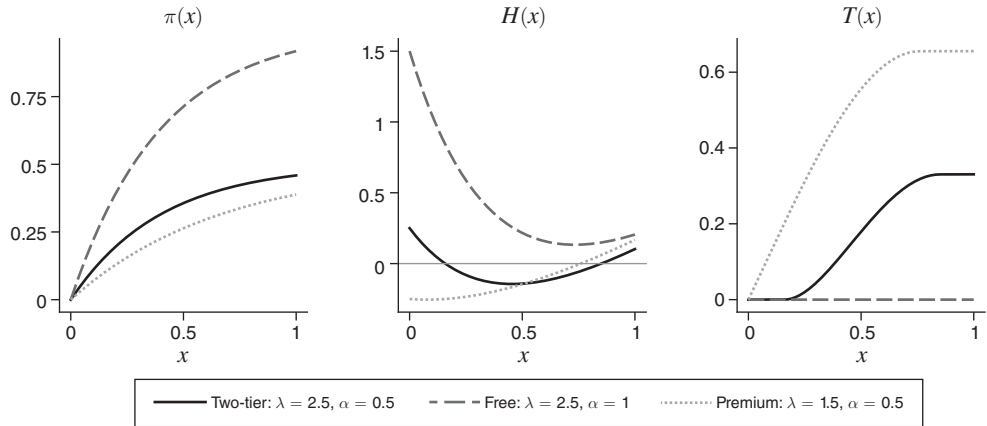


FIGURE 4. HYBRID PRICE SCHEDULES IN EXAMPLE 4

for some $\omega \in \mathbb{N}$ and $0 < k < 1/(2\pi\omega)$.²³ While the functional form of the cosine is intentionally ad hoc, observe in the leftmost panel of Figure 5 that this revenue function only slightly deviates (for small k) from the linear baseline in inducing mild oscillations in revenue over time. The constrained marginal revenue function is $H(x) = k\sin(2\pi\omega x)$, which crosses zero from above ω times, generating an $\omega + 2$ -part tariff. In Figure 5, we plot H and the optimal tariff T for $\omega \in \{1, 2, 3\}$ and $k = 0.04 < 1/(6\pi)$. The tariffs respectively have one, two, and three flat-pricing intervals; these correspond to three-, four-, and five-part tariffs.

IV. Welfare under Multi-part Tariffs

Having studied the positive implications of underutilization for pricing, we now study its welfare consequences. First, we study the effect of introducing underutilization, or removing perfect contractibility, on consumer and producer welfare. Second, we study how the possibility for underutilization mediates the welfare effects of changes in the structure of revenue. We apply these results to understand welfare in digital goods markets.

A. The Impact of Underutilization on Welfare

For an arbitrary price schedule T , we define *consumer welfare* under free disposal for each type θ as

$$(19) \quad V(\theta; T) = \sup_{y \in X, x \in [0, y]} \{u(x, \theta) - T(y)\},$$

²³ Under this condition, it is immediate to verify that the induced virtual surplus function J satisfies our running assumptions: it is strictly concave and satisfies strict single crossing.

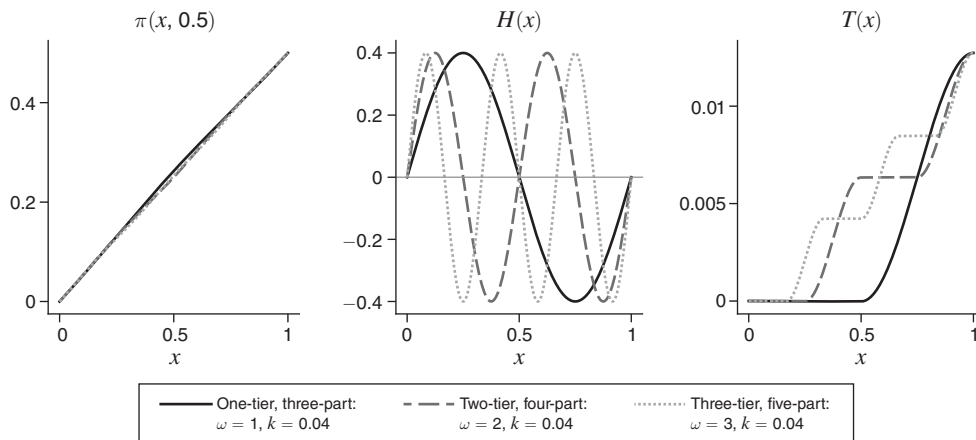


FIGURE 5. ARBITRARY MULTI-PART TARIFFS IN EXAMPLE 5

which is the payoff corresponding to an optimal purchase. We define *producer welfare with free disposal* for each purchasing type θ , $\Pi(\theta; T)$, as

$$(20) \quad \Pi(\theta, T) = \pi(\phi(\theta; T), \theta) + T(\xi(\theta; T)),$$

which is simply total revenue (i.e., the integrand in the objective in equation 1), evaluated at a fixed T and the buyers’ optimally chosen consumption $\phi(\theta; T)$ and purchases $\xi(\theta; T)$. We define $V^*(\theta)$ and $\Pi^*(\theta)$ as the corresponding equilibrium welfare quantities evaluated at the optimal price schedule, T^* . We finally define analogous quantities under no free disposal, or perfect contractibility, as $V_N(\theta; T)$, $\Pi_N(\theta; T)$, $V_N^*(\theta)$, and $\Pi_N^*(\theta)$.

The following result summarizes how consumer and producer surplus are affected by the possibility of underutilization of consumption.

PROPOSITION 3 (Contractibility and Welfare): *For any price schedule T , $V(\theta; T) \geq V_N(\theta; T)$ for all $\theta \in \Theta$. However, under the optimal price schedules, $V^*(\theta) \leq V_N^*(\theta)$ and $\Pi^*(\theta) \leq \Pi_N^*(\theta)$ for all $\theta \in \Theta$.*

All consumers lose out from increased contractibility for any given pricing policy, as the scope for payoff-increasing actions declines; but all consumers *gain* from increased contractibility under the seller’s reoptimized price schedule. Intuitively, consumers would like to commit *ex ante* to avoid the possibility of moral hazard and have the seller pay them to take certain actions; increasing contractibility provides exactly this commitment device. Sellers gain from increased contractibility (from selling to all types and, therefore, in total across types) for the simple reason that it increases the set of implementable allocations.

This result contextualizes the claim that users, particularly of social media platforms, are not fairly remunerated, given the fact that their time and data are the sources of these platforms’ revenues. This idea is exemplified by Apple cofounder

Steve Wozniak’s stated rationale for why he deleted his personal Facebook account (Guynn and McCoy 2018):

Users provide every detail of their life to Facebook and (...) Facebook makes a lot of advertising money off this. (...) The profits are all based on the user’s info, but the users get none of the profits back. (...) As they say, with Facebook, you are the product.

This quote can be understood through the lens of our model. In light of the inherent noncontractibility of usage of social media platforms, consumers cannot be paid on the margin. In this context, Proposition 3 implies that consumers would be better off were their usage perfectly contractible and revenue-generative usage remunerated. However, our rationalization highlights that this issue is fundamentally technological: even if Facebook wished to pay users, they would not be able to do so without being exploited in the manner that led to the downfall of AllAdvantage.

B. Comparative Statics for Welfare

We next study how the presence of free disposal mediates the welfare effects of changes in usage-based revenues π and demand F . Changes in revenue may be driven by underlying changes in advertisement and data-collection technology while changes in demand may be driven by changes in demographics and the quality of goods.

We say that $\tilde{\pi}$ exhibits more profitability of usage than π if every unit of usage generates more revenue, that is, $\tilde{\pi}_x \geq \pi_x$. We say that \tilde{F} exhibits weaker demand than F if F dominates \tilde{F} in the hazard-rate order.²⁴ We write consumer and producer equilibrium welfare as functions of these arguments. We now show how these two comparative statics affect welfare, with and without free disposal.

PROPOSITION 4 (Comparative Statics for Welfare): *If $\tilde{\pi}$ exhibits more profitability of usage than π and \tilde{F} exhibits weaker demand than F , then, for all $\theta \in \Theta$:*

$$(21) \quad 0 \leq V^*(\theta; \tilde{\pi}, \tilde{F}) - V^*(\theta; \pi, F) \leq V_N^*(\theta; \tilde{\pi}, \tilde{F}) - V_N^*(\theta; \pi, F),$$

$$(22) \quad 0 \leq \Pi^*(\theta; \tilde{\pi}, \tilde{F}) - \Pi^*(\theta; \pi, F) \leq \Pi_N^*(\theta; \tilde{\pi}, \tilde{F}) - \Pi_N^*(\theta; \pi, F).$$

This shows that both consumer and producer welfare increase, but less so than under perfect contractibility. In this sense, free disposal erodes the potential welfare gains for both buyer and seller relative to a counterfactual perfect-contractibility world.

The intuition for this result is that, with higher marginal revenue from consumption for the seller or lower demand types (implying smaller distortions from information rents), the seller-optimal consumption is larger. This increases consumer and producer welfare, with and without free disposal. However, with free disposal, there are more types who consume their bliss point. Thus, both buyers and sellers have greater welfare gains under perfect contractibility of usage than

²⁴Our nomenclature reflects the fact the hazard-rate order implies first-order stochastic dominance (see Theorem 1.B.1 in Shaked and Shanthikumar 2007) and is therefore a (strong) notion of decreasing demand.

under free disposal. This further underscores the sense in which perfect contractibility simulates commitment. Proposition 3 shows how this commitment increases welfare in *levels*; Proposition 4 shows how this commitment increases the *sensitivity* of welfare gains from changes in external revenue and demand.

This result allows us to shed light on the welfare effects of changes in advertising technology, such as the past decade's advent of more valuable targeted advertisements based on user data. We can capture this phenomenon as an increase in the profitability of usage for both platforms serving advertisements and platforms profiting from data sales. Through the lens of our model, any marginal increase in profitability, including the introduction of advertisement- or data-based revenue to a previously zero-revenue ($\pi = 0$) service, increases welfare for all consumers. These increases could be larger, however, if digital goods' usage were fully contractible. The reason is precisely the fact that users cannot be paid on the margin for their clicks and data. Moreover, even if only a subset of users are paid for their clicks and data, all users benefit due to changes in the overall price schedule.

This analysis has two potentially relevant implications for policy discussion about digital-goods regulation. First, government regulations that may diminish marginal advertising revenues, like the European Union's General Data Privacy Regulation, may weakly reduce consumer surplus. But the extent of these reductions, and their distributional consequences, depends on whether sellers are far on the interior of the zero-marginal-pricing constraint. That is, if marginal units are already free (e.g., for all units of Google or Facebook), then consumer surplus is unaffected and any (here, unmodeled) gains from the instrumental value of privacy lead to net consumer benefits.

Second, our focus on contractibility contrasts with an influential perspective in the literature that focuses instead on the lack of collective bargaining for users of online products (Posner and Weyl 2018). Taken to the extreme, our results imply a thorny "privacy paradox" for consumers—the only way to properly reap the benefits of the surplus generated by targeted advertising is to surrender *additional* privacy by enabling tools to more precisely monitor usage and attention.

V. Conclusion and Summary of Extensions

We study optimal nonlinear pricing in environments with feasible underutilization and usage-derived revenue, features that are ubiquitous in the digital goods context. We show how the combination of these two forces rationalizes the occurrence of multi-part tariffs, or price schedules that include at least one tier of zero marginal prices. The key mechanism is that sellers have an incentive to pay users on the margin, due to usage-derived revenue (e.g., from advertisement or data collection), but noncontractibility prevents such arrangements from being enforceable. More succinctly, zero marginal pricing is the sellers' constrained optimum in a world in which "pay to click" is impossible.

We apply these results to study positive and normative features of digital goods markets. We show how different structures of external revenue translate into different familiar pricing schemes like free trials, unlimited subscriptions, and free products. We moreover show that the scope for underutilization reduces buyers' welfare and dampens the ability of buyers to reap the rewards from the revenue they generate.

We finally discuss additional analyses contained within the online Appendix and ongoing work.

Optimal Bunching.—We assumed that virtual surplus was strictly single crossing to rule out the possibility that multiple buyer types optimally bunch on the same level of consumption. In online Appendix B.1, we relax this assumption, adapt the assignment approach of Nöldeke and Samuelson (2007), and characterize the optimal contract when bunching is a possibility. Analogously to Proposition 1, the optimal consumption function maximizes the suitable transformation of virtual surplus in this setting, subject to the constraint that no agent consumes more than their bliss point. Under this solution and as per our main analysis (Proposition 2), the price schedule is flat whenever this constraint binds and the optimal price schedule is a multi-part tariff.²⁵

Underutilization with (Perfect) Competition.—It is of course natural to ask how our analysis in a monopoly setting would extend to markets with competition. In online Appendix B.2, we solve for the equilibrium outcome of our screening model when the monopolist faces a perfectly competitive fringe of potential entrants. Specifically, we model this free entry by adding an additional constraint imposing zero profits for the monopolist (i.e., total external revenues plus transfers). The equilibrium price schedule under perfect competition features zero marginal pricing more often than under monopoly pricing (Corollary 3 in online Appendix B.2). The reason is that the perfect competition, modeled this way, leads firms to maximize total surplus instead of virtual surplus—put differently, competition ensures that all profits from screening and distorting down consumption are competed away. Total surplus is maximized by a higher level of consumption owing to the absence of information rents, and so the constraint of underutilization is more often binding. As a result, our model implies that multi-part tariffs are likely to be more prevalent in scenarios with fierce competition between sellers than under monopoly.²⁶

Pricing with Partially Contractible Usage.—In ongoing work, we study the effects of *partial* contractibility of usage on optimal nonlinear pricing. This allows us to capture further settings in which some levels of utilization are contractible (e.g., usage by high-end users) while others are not. In this model, agents switch endogenously from being consumers, who either pay for a good or receive it for free, to being workers, who are paid for their usage of a platform. This model can be used, for example, to capture the pricing of YouTube and TikTok, where low-end content creators use the platforms for free, while high-end content creators are paid.

²⁵In this setting, we cannot develop a similar characterization to that of Proposition 2 by comparing the marginal benefits and costs of additional consumption for the seller as the objective is no longer quasiconcave and global properties determine whether the constraint binds.

²⁶Another setting that features the same pricing implications as perfect competition, yet with a single monopolistic seller, is one in which buyers commit to participate before learning their types. This structure corresponds to the ex ante contracting setting of Grubb (2009).

REFERENCES

- Allcott, Hunt, Matthew Gentzkow, and Lena Song. 2022. "Digital Addiction." *American Economic Review* 112 (7): 2424–63.
- Amelio, Andrea, and Bruno Jullien. 2012. "Tying and Freebies in Two-Sided Markets." *International Journal of Industrial Organization* 30 (5): 436–46.
- Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin. 2018. "The Design and Price of Information." *American Economic Review* 108 (1): 1–48.
- Carbajal, Juan Carlos, and Jeffrey C. Ely. 2013. "Mechanism Design without Revenue Equivalence." *Journal of Economic Theory* 148 (1): 104–33.
- Choi, Jay Pil, and Doh-Shin Jeon. 2021. "A Leverage Theory of Tying in Two-Sided Markets with Nonnegative Price Constraints." *American Economic Journal: Microeconomics* 13 (1): 283–337.
- eMarketer Insider Intelligence. 2020. "Google's US Ad Revenues to Drop for the First Time." *eMarketer Insider Intelligence*, June 22. <https://www.emarketer.com/newsroom/index.php/google-ad-revenues-to-drop-for-the-first-time/>.
- Gershkov, Alex, Benny Moldovanu, Philipp Strack, and Mengxi Zhang. 2021. "A Theory of Auctions with Endogenous Valuations." *Journal of Political Economy* 129 (4): 1011–51.
- Google AdSense. 2020. "Online Terms of Service." Google AdSense. <https://www.google.com/adsense/localized-terms>.
- Grubb, Michael D. 2009. "Selling to Overconfident Consumers." *American Economic Review* 99 (5): 1770–1807.
- Guernsey, Lisa. 1999. "Can It Pay to Surf the Web?" *New York Times*, November 22. <https://www.nytimes.com/1999/07/01/technology/can-it-pay-to-surf-the-web.html>.
- Guynn, Jessica, and Kevin McCoy. 2018. "Apple Co-founder Steve Wozniak Says He's Left Facebook over Data Collection." *USA Today*, April 8. <https://www.usatoday.com/story/tech/2018/04/08/apple-co-founder-steve-wozniak-says-hes-leaving-facebook/497392002/>.
- IHS Markit. 2017. "The Economic Value of Behavioural Targeting in Digital Advertising." IHS Markit, September 7. <https://iabeurope.eu/knowledgehub/policy/the-economic-value-of-data-driven-advertising/>.
- Kang, Y. Peter. 2000. "It Pays to Cheat, Not Surf." *Wired*. July 10. <https://www.wired.com/2000/07/it-pays-to-cheat-not-surf/>.
- Laffont, Jean-Jacques, and Jean Tirole. 1986. "Using Cost Observation to Regulate Firms." *Journal of Political Economy* 94 (3, Part 1): 614–41.
- Maskin, Eric, and John Riley. 1984. "Monopoly with Incomplete Information." *RAND Journal of Economics* 15 (2): 171–96.
- Mussa, Michael, and Sherwin Rosen. 1978. "Monopoly and Product Quality." *Journal of Economic Theory* 18 (2): 301–17.
- Nöldeke, Georg, and Larry Samuelson. 2007. "Optimal Bunching without Optimal Control." *Journal of Economic Theory* 134 (1): 405–20.
- Posner, Eric A., and E. Glen Weyl. 2018. *Radical Markets: Uprooting Capitalism and Democracy for a Just Society*. Princeton: Princeton University Press.
- Sartori, Elia. 2021. "Competitive Provision of Digital Goods." Unpublished.
- Shaked, Moshe, and J. George Shanthikumar. 2007. *Stochastic Orders*. New York: Springer.
- Strausz, Roland. 2017. "A Theory of Crowdfunding: A Mechanism Design Approach with Demand Uncertainty and Moral Hazard." *American Economic Review* 107 (6): 1430–76.
- Sundararajan, Arun. 2004. "Nonlinear Pricing of Information Goods." *Management Science* 50 (12): 1660–73.
- Wilson, Robert B. 1993. *Nonlinear Pricing*. New York: Oxford University Press.
- Yang, Kai Hao. 2022. "Selling Consumer Data for Profit: Optimal Market-Segmentation Design and its Consequences." *American Economic Review* 112 (4): 1364–93.