NBER WORKING PAPER SERIES

MUSIC FOR A SONG:
AN EMPIRICAL LOOK AT UNIFORM SONG PRICING AND ITS ALTERNATIVES

Ben Shiller
Joel Waldfogel

Working Paper 15390
http://www.nber.org/papers/w15390

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2009

We thank the Mack Center at Wharton for financial support. Tom Holmes provided helpful comments on an earlier draft. We are also grateful to seminar participants at Kellogg, Ohio State, Michigan, the NBER 2008 Summer Institute IO group, QME, INFORMS in Washington, and ZEW in Manheim for comments. All errors are our own. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

Music for a Song: An Empirical Look at Uniform Song Pricing and its Alternatives
Ben Shiller and Joel Waldfogel
NBER Working Paper No. 15390
October 2009
JEL No. L12,L82

ABSTRACT

Economists have well-developed theories that challenge the wisdom of the common practice of uniform pricing. With digital music as its context, this paper explores the profit and welfare implications of various alternatives, including song-specific pricing, various forms of bundling, two-part tariffs, nonlinear pricing, and third-degree price discrimination. Using survey-based data on nearly 1000 students' valuations of 100 popular songs in early 2008 and early 2009. We find that various alternatives – including simple schemes such as pure bundling and two-part tariffs – can raise both producer and consumer surplus. Revenue could be raised by between a sixth and a third relative to profit-maximizing uniform pricing. While person-specific uniform pricing can raise revenue by over 50 percent, none of the non-discriminatory schemes raise revenue's share of surplus above 40 percent of total surplus. Even with sophisticated pricing, much of the area under the demand curve for this product cannot be appropriated as revenue.

Ben Shiller
University of Pennsylvania
bshiller@wharton.upenn.edu

Joel Waldfogel
Business and Public Policy
University of Pennsylvania, Wharton School
1400 Steinberg Hall-Dietrich Hall
Philadelphia, PA 19104-6372
and NBER
waldfogj@wharton.upenn.edu

The prominence of the iTunes Music store, where until recently all songs sold for \$0.99, has focused attention on uniform pricing and its alternatives. Economists have well-developed normative theories that raise questions, at least in theory, about the wisdom of uniform pricing. Alternatives to uniform pricing include song-specific pricing, various forms of bundling, two-part tariffs, and nonlinear pricing. Many of these approaches are well understood in theory.[1] But determining the amount of additional profit or consumer surplus available from using alternatives to uniform pricing is an empirical question in each context where it might be applied.

Quantifying the surplus foregone by uniform pricing is a matter of current practical as well as academic interest. Apple has now sold over 6 billion songs at iTunes, over 2 billion in 2008 alone (Christman, 2008; Musgrove, 2009), until recently using only uniform pricing. During the summer of 2007, some record labels declined to renew their contracts with Apple out of a desire for more flexibility in pricing.[2] In September 2007, Amazon launched a music downloading service featuring, among other things, song-specific, or "variable" pricing.[3] Earlier this year Apple moved toward song-specific pricing with a three-tier pricing structure. Nokia implements a different alternative, pure bundling, with unlimited song downloads on phones sold with a "Comes with Music" surcharge.

In general, it is hard to know how much money uniform pricing leaves on the table because the data needed to evaluate this question – the full distribution of reservation prices across buyers and products – are hard to come by. Usually, researchers

---

[1] See Stigler (1963), Adams and Yellen (1976), Schmalensee (1984), Armstrong (1999), Bakos and Brynjolfsson (1999) and others discussed below.
[2] David Carr, "Steve Jobs: iCame, iSaw, iCaved," *New York Times*, section C, page 1, September 10, 2007.
[3] Amazon's service also features music without digital rights management. See Ed Christman, "Amazon: Keeping it on the Down Low," Billboard, October 6, 2007.

estimate some sort of demand system allowing inference about individuals' valuations of various quantities of various products.[4] If one could directly observe buyers' reservation prices for products, some sophisticated forms of pricing would be easily implementable. This paper pursues this goal using survey-based direct elicitation of students' valuations of 100 popular songs at iTunes during January 2008 and January 2009.

The effect of sophisticated pricing on revenue has implications beyond pricing strategy. The welfare economics of imperfect competition depends crucially on the extent to which the social benefit of a product – the area under the demand curve – can be captured as revenue for the seller.[5] It's obvious that the producer surplus under single-price monopoly can fall short of fixed and variable costs even when the joint surplus would exceed costs. Hence, the market can fail to provide goods with benefit in excess of costs. Of course, whether sellers can capture surplus as revenue depends on the effectiveness of price discrimination. Perfect first degree price discrimination eliminates, or at least substantially mitigates, the underprovision problem.[6] But we do not know what share of surplus is appropriable as revenue in practice. One of this paper's contributions is to show the share of surplus that is appropriable in one context.

The paper proceeds in five sections. First, we briefly discuss the pricing problem, along with the relevant theoretical and empirical literatures on various kinds of pricing. In section 2 we discuss our data collected from two surveys in 2008 and 2009. Section 3 discusses our parametric approach to characterizing the valuation distributions needed for our pricing exercises. Section 4 presents our results on revenue and the distribution of

[4] See, for example, Chu, Leslie, and Sorenson (2007) and other studies discussed below.
[5] See Spence (1976ab), Dixit and Stiglitz (1977), and Heal (1980) for elaboration of these arguments.
[6] See Edlin, Epelbaum, and Heller (1998) for a discussion of how price discrimination can bring about efficiency.

surplus available under various pricing schemes, including revenue-maximizing uniform, song-specific pricing, two-part tariffs, nonlinear bundle-size pricing, pure bundling, (simple versions of) mixed bundling, and individual customer-specific pricing. Section 4 also presents bootstrap-based evidence on the precision of our estimates. Section 5 finds Pareto-improving pricing schemes that deliver at least the levels of consumer and producer surplus achieved under profit maximizing uniform pricing. A brief conclusion follows.

Because our songs and respondents are not representative, we cannot prescribe an optimal pricing policy (or pass judgment on the wisdom of, say, a $0.99 uniform price). But we can compare the breakdown of the area under the demand curve into producer surplus, consumer surplus, and deadweight loss under different pricing schemes. We find, for example, that revenue-maximizing uniform pricing collects just over a quarter of surplus as revenue, while leaving almost 45 percent of surplus as consumer surplus and the remainder as deadweight loss. Using revenue-maximizing uniform pricing as the benchmark, we find that various alternative schemes such as pure bundling and two part tariffs could raise revenue by about 17 percent in one sample and by about 30 percent in another. While none of the self-selecting schemes can raise revenue much beyond a third of total surplus, person-specific third degree price discrimination can raise producer surplus by more than 50 percent over its level under uniform pricing, allowing sellers to appropriate over half of surplus as revenue. Third degree price discrimination based on observables, by contrast, has very little effect on producer surplus.

I. Theoretical Setup and Literature Review

1. Setup

Each consumer *i* has a reservation price for each song *s* (V_{is}). These reservation prices, in conjunction with the pricing schemes, determine whether the individual purchases 1 or 0 units of each song (and we assume resale impossible). Consumers evaluate bundles of songs by adding their valuations of each of the songs included in bundles.

The seller's problem is to choose a vector of prices *P* for the songs, or groups of songs, to maximize his profit. For example, if the seller is a uniform single-price monopolist selling 50 different products, then he chooses only one price, and the price of each bundle depends proportionally on the number of songs on the bundle. That is, if bundle *k* contains 5 songs, and *p* is the single price per product, then $P_k=5p$. With other pricing schemes, the *P* vector becomes more complicated. If the seller engages in component pricing, with a potentially different price for each song, then the bundle price is the sum of the prices of each of the elements of the bundle. If the seller employs a two part tariff, the first song purchased has one price, while each additional song has another. If the seller engages in "bundle size pricing," then P_k simply depends – generally nonlinearly – on the number of elements in the bundle purchased, so there are 50 separate prices to set. If the seller engages in mixed bundling, then each combination of products can potentially have its own price. With *n* products, there are (2^n – 1). With 3, this amounts to 7. With 50 products, the number of combinations is about a quadrillion. With pure bundling, the seller sets a single price for the entire bundle of 50 songs. Finally, under person-specific pricing, and extreme form of 3rd degree price discrimination, the seller sets a (potentially different) price for each consumer.

Define V_{ik} as individual *i*'s valuation of song bundle *k*. (The bundle could be an individual song, or any combination of the 50 songs available). Each consumer's problem is to maximize $V_{ik} - P_k$ by his choice of *k*, or which product (or bundle of songs) to consume.

Given a pricing scheme along with consumers' valuations of products, a computer can mechanically determine how much revenue each scheme would collect. For some of these schemes – notably, nonlinear pricing and mixed bundling – the number of possible pricing schemes is large so that even a fast computer can examine only a limited range of schemes (in particular for nonlinear pricing with more than about 5 separate prices, and for mixed bundling just a few products).

2. Related Literature

The textbook theory of single price monopoly (see, for example, Pindyck and Rubinfeld, 2006) provides a guide to profit-maximizing uniform pricing. The simplest alternative to uniform pricing (UP) is song-specific, or component pricing (CP). Analytically, it involves the same apparatus as UP, albeit with a single separate price per song. Song-specific pricing is currently employed by digital music sellers Amazon as well as Aimee Street.com, and many observers consider Apple's recently abandoned practice of uniform pricing a blunt instrument for revenue maximization.[7] Because UP is a constrained special case of CP – with equal prices for all songs – UP cannot produce more revenue than CP, and CP would in general be expected to produce more.

---

[7] The puzzle of uniform prices arises in other product markets, for example in movie theater pricing. See Orbach and Einav (2007).

There is a substantial theoretical literature on sophisticated alternatives to uniform pricing. This body of work provides guidance about how pricing schemes other than uniform pricing can be expected to affect profits and other aspects of surplus. Stigler (1963) presented a two-product example showing that bundling could produce more revenue pricing the products separately when consumers' valuations of the products are negatively correlated. Adams and Yellen (1976) introduced mixed bundling (MB) with examples where mixed bundling produced more profit than either pure bundling (PB) or product-specific pricing. Schmalensee (1984) shows that PB can be more profitable than product-specific pricing even when the correlations of consumer valuations are positive. McAfee, McMillan and Whinston (1989) show that mixed bundling (MB) always beats pure bundling although – again – it results in complex pricing schemes.

A few recent papers outline the situations in which a firm can extract surplus by selling large multiproduct bundles. Bakos & Brynjolffson (1999) show that if the valuation of a large bundle is more predictable than the valuation of individual products, then as the number of products grows large, pure bundling extracts the entire surplus as revenue. Fang and Norman (2006) obtain related results with finite numbers of products. Armstrong (1999) also shows that when tastes are correlated across products – for example because of income differences across consumers – a menu of two-part tariffs is almost optimal.[8]

A few papers use characterizations of demand to create pricing schemes. Chu, Leslie, and Sorenson (2007) estimate demand for eight plays at a Palo Alto theater using data on purchases of individual play tickets as well as bundles. They use their estimated

[8] Sundararajan (2004) develops a model of optimal pricing for large bundles when there are costs of administration.

model to create profit-maximizing pricing schemes under uniform pricing, component pricing, pure bundling, bundle-size pricing, and mixed bundling. Relative to uniform pricing, component pricing raises revenue 1.4 percent, bundle size pricing raises revenue 2.3 percent, and mixed bundling raises revenue 4 percent, but none of the alternatives to uniform pricing raises revenue by more than 5 percent.[9] McMillan (2007) estimates a model of demand for soft-drinks at a grocery store which he uses to calculate that uniform pricing would cost the store $60 in profit (on revenue of roughly $10,000). Crawford (2008) estimates a model of demand for cable television bundles which he uses to simulate effects of adding channels to available bundles. He estimates that bundling an average top-15 special interest cable network gives rise to a 4.7 percent increase in profit, a 4.0 percent decrease in consumer surplus, and a 2.0 percent increase in total surplus.

II. Data

1. *Sources and Validity*

The basic data for this study are drawn from two surveys, each of about 500 Wharton undergraduates, performed in January of 2008 and January of 2009. In the 2008 survey, undergraduates – mostly freshmen – at Wharton were required to fill out an online survey which presented them with the top 50 songs at iTunes (as of January 11, 2008). Students were given instructions and paper worksheets in class on January 16 and 17. For each song, students were told to listen to a clip to remind themselves of the song, then to write down the maximum amount they would be willing to pay to get the song

[9] They emphasize that bundle size pricing achieves 98 percent of the revenue of mixed bundling. Said another way, bundle size pricing achieves 60 percent of mixed bundling's improvement over uniform pricing.

from the sole authorized source. Completion of the assignment was necessary for problem set credit to motivate students to participate carefully.

In particular, students were given the following instructions:

> Imagine that, unlike in current reality, there is only one authorized source for each song. Put aside what you know about prices at existing outlets because for this survey we're pretending that they don't exist.
>
> For each song listed in the survey, indicate the maximum amount you would be willing to pay to obtain it from the sole authorized source. **For this exercise, I'm asking you to report what it is worth to you, not what price you think would be fair or what price you are accustomed to paying. That is, I'm asking you to indicate the maximum amount you would be willing to pay to obtain it from the authorized source.**
>
> For example, if you already purchased it, then at the time you bought it, you were willing to pay at least the price you paid but you might have been willing to pay more. If you would prefer not to have it even if it were free, you would indicate 0.
>
> On the following pages, you will be presented with a list of songs and artists. In the space provided for each song, enter the maximum amount you are willing to pay for the song (for example 1.75, NOT $1.75). You must enter a dollar amount for each song.

The resulting 2008 dataset includes 23,150 observations on individual-song valuations for 463 people, as well as a small amount of information on the respondents: age (mostly 18-20), gender, race, self-reported level of interest in music (not interested, somewhat interested, very interested), and the size of their music library.

The 2009 data, which includes 21,650 observations from 433 respondents, differs in two respects. First, rather than the top 50 songs at iTunes, the 2009 data include the top 25 songs at iTunes, along with 10 songs that were the top 10 songs 6 months earlier, and 15 songs drawn randomly from those ranked between 26 and 100 at the time of the survey. The goal of the different sampling scheme is to explore if the effectiveness of alternative pricing schemes in the 2008 data arises from the narrow choice of songs. A second difference in the 2009 data is that respondents were asked not only their valuation of each song but also whether they possessed the song and, if so, whether they had purchased or obtained it via file sharing.

One concern with valuation data drawn from surveys is that respondents would anchor on prices they know to be charged at actual websites, for example the well-known \$0.99 iTunes price employed when the data were collected. We are not interested in the students' beliefs about current pricing; we are interested in their maximum willingness to pay.

This leads us to visual inspection of the data, and nearly all of the valuations are "reasonable": 98 percent of valuations fall on [0,10], and 86 percent fall on [0,2]. There is some clustering at the familiar price around \$1 but not substantially more than the clustering on other multiples of \$0.25. Of nearly 24,000 responses for 2008, there are 7257 zeroes. Data for 2009 are similar (see Figure 1).[10] The clustering of valuation responses at round and focal numbers (notably, zero, multiples of \$0.25, \$0.99, and \$1.99) gives rises to plateaus on the demand curve that create "sawtooth" spikes in the relationship between revenue and price. If we believed that the reported valuations were literally true, we could use the raw data to find optimal pricing schemes. Instead, as we explain below, we will fit the data to some parametric distribution. (The four panels of Figure 1 cover data for 2008 and 2009, both raw – top – and fit to a parametric distribution – bottom – via a process we describe below at III.)

It is prudent to ask whether surveys can elicit meaningful valuation information. It is well known that question wording affects responses. In earlier work on music valuation, questions asking for willingness to pay tended to elicit much lower valuations

[10] Another test of reasonableness, pursued below, is whether optimal prices charged to these valuations are close to actual prices. Of course, such a test involves the joint hypothesis of "reasonable" data and profit-maximizing pricing decisions.

than questions asking for amounts required to give up music (Rob and Waldfogel, 2006). This is the familiar endowment effect (Knetsch and Sinden, 1984). Sensitivity of results to question wording has led some researchers to be skeptical of survey responses (Diamond and Hausman, 1994). One response to this concern stems from the fact that the profit-maximizing prices implied by buy-based valuations in Rob and Waldfogel (2006) – similar to those used here – tend to be close to observed prices while those implied by sell-based valuations are far off. We find that below as well, providing some assurance that the survey wording gets at the valuations relevant to the pricing decision. A second response is that here, people are valuing familiar items rather than, say, pristine Alaskan wilderness they have never seen.

Our 2009 data allow a more direct test of reasonableness. We need to be able to rely on our reported valuations to make statements about which songs – and bundles of songs – respondents would purchase when facing various prices. Because the 2009 data include information on whether the respondent possesses the song, we can perform a simple check on the data: do respondents actually possess the songs that they report valuing at or above $0.99? Because it is easier for the respondent to accurately report whether he or she has a copy of a song than to accurately report the valuation he or she attaches to the song, this check can validate the valuation responses.

Each observation is a song-respondent pair. Recall that respondents report valuations for each song, along with whether they possess the song. For examining the relationship between valuation and ownership, we discard song-respondent observations that the respondents obtained via file sharing, since obtaining a song without payment is consistent with low or high valuation. With the remaining data we ask whether the probability that the observations in each valuation range are owned relates to the reported valuation.

Figure 2 summarizes the relationship between valuation and purchase, and it is clearly positive. The data are divided into 50 equal-sized groups, ordered by valuation. In each cell, we compute the probability that song-respondent observations are owned. The horizontal axis records the logarithm of the cell's average song valuation, while the vertical axis indicates the probability that the song is owned. As the valuations rise, the probability of purchase rises as well, to about 20 percent for average valuations near – but below - $0.99. When the average valuations reach $0.99 – and the log valuation approaches 0 – the probability of ownership jumps to about 50 percent. As the average valuation continues to increase, so does the probability of ownership. Figure 2 shows that valuation information is predictive of purchase behavior.

A related check on whether the data are reasonable is whether the songs for which respondents frequently report high valuations are also the songs with higher sales among buyers generally. Although we cannot observe song sales directly, we have two indirect measures of sales. The number of weeks a song has been on the "Billboard Hot Digital" chart, along with its peak single chart position, provides indirect measures of its cumulative sales.[11] A regression of the share of respondents reporting valuations of at least $0.99 on these two sales proxies (for the 2008 data) yields:

[11] The Billboard Chart, "Hot Digital Songs," is available at http://www.billboard.com/bbcom/charts, accessed March 14, 2008. The Hot Digital Chart, unlike many other charts, reports sales of singles, regardless of whether a CD version of the single exists.

$$share(V > \$0.99) = \underset{(0.052)}{0.352} + \underset{(0.0017)}{0.0033} * (chart\_weeks) - \underset{(0.0015)}{0.0044} * (chart\_peak), \text{ where}$$

chart_weeks indicates the number of weeks the song had been on the Billboard chart as of March 8, 2008; and *chart_peak* indicates the song's peak chart position. Standard errors are in parentheses, and the R-squared from this regression is 0.34. Results support the notion that the data are reasonable in some rudimentary sense: Songs on the chart longer have a higher simulated sales penetration in this sample, and songs with a lower chart peak (a higher peak rank) have higher penetration.

Our survey approach to eliciting reservation valuations has antecedents in the marketing and operations literatures. Hanson and Martin (1990) employ direct elicitation of reservation values for an exercise reported in their study. Kalish and Nelson (1991) explore direct elicitation along with preference rating and ranking measures and find that reservation values do well in terms of fit while the other measures are superior in predicting choice on a holdout sample. Venkatesh and Mahajan (1993) survey respondents on their willingness to pay for performances of Indian music. The authors calculate revenue under component pricing, pure bundling, and mixed bundling. Jedidi, Jagpal, and Manchanda (2003) estimate an empirical model of survey respondents' reservation values of individual goods and bundles of two. While Jedidi, Jagpal, and Manchanda (2003) observe that self-stated reservation prices are subject to measurement error, especially for infrequently purchased products, the songs in our survey are familiar and commonly purchased.

2. Data Summary

Table 1 reports the average valuations per song in the 2008 sample, as well as the median, 25th, and 75th percentile valuations. The most highly valued songs on the list have average valuations over \$2.00. They include "Stronger" by Kanye West, "Apologize" by Timbaland (feat. One Republic), and "The Way I Are" by Timbaland (feat. Keri Hilson & DOE). The lowest valued songs, with mean valuations below \$0.7, are those targeted at consumers younger than our college-student respondents. Examples include artists such as Alvin and the Chipmunks and Disney artists such as the Jonas Brothers and Miley Cyrus. The 25th percentile valuation for most songs is around \$0.10, and the 75th percentile valuation is typically over \$1.5. To conserve space, we do not report the analogous table for 2009, but like 2008 the 2009 data displays substantial inter-song variation.

Valuations vary substantially both within and across respondents. Figure 3 characterizes the distribution of cumulative valuations (e.g. how highly a respondent values his top 10 songs, for example). This figure shows that the median valuation of (each individual's) top 10 songs among these 50 in the 2008 sample is about \$20, while the 75th percentile valuation is about double that (\$40), and the 25th percentile valuation is around \$15. Valuations in the 2009 sample are lower: for example, the median person's valuation of his or her top 10 is around \$15. The flattening of each of these curves indicates substantial difference between the valuations of the most highly and least highly valued songs. Analyzed a different way, the valuation data indicate that the vast majority of the variation in the reported valuations arise across individuals, as opposed to songs. A regression of 2008 (2009) valuations on song fixed effects yields an R-squared of 4.0

(5.9) percent. The R-squared from a regression on only individual effects is 40.2 (31.5) percent, and the R-squared with both individual and song effects is 44.2 (37.4) percent.

The correlations of song valuations across persons help to determine the extent to which non-uniform pricing schemes can capture additional revenue. For example, a common intuition from bundling theory is that bundling raises revenue more as products' valuations are less positively correlated. Song valuations in our datasets are positively correlated. With 50 songs there are 1,225 pairwise song correlations. The mean correlation is 0.37 for the 2008 sample and 0.25 for the 2009 sample. Figure 4 shows the whole distribution of pairwise correlations for each year.

Because the two samples are made up of different groups of songs (by vintage and ranking), it is interesting to explore the cause of the lower correlation among songs in the 2009 sample. Each of the samples contains the top 25 iTunes songs at the time of the survey. When we restrict attention to these songs, the pairwise correlations average 0.40 for the 2008 sample and 0.26 for the 2009 sample. Hence, the lower correlations in the 2009 sample are not attributable to the latter sample's wider breadth in popularity and vintage. Whatever the cause of the lower correlation, we can expect bundling to raise revenue more for the 2009 songs.

III. Characterizing the Valuation Distributions

In order to implement pricing schemes we need a characterization of the distribution of reservation prices across persons and products. Given the nature of our data – based on direct elicitation of each person's valuation of each product – one approach is obvious: use the observed valuation responses directly. The major shortcoming of this approach is the irregularity of the resulting revenue functions and surfaces arising from the spikes in the top panel of Figure 1. This irregularity strains credulity if we think that underlying utility functions are smooth; such irregularity also makes it hard to efficiently identify revenue-maximizing solutions. An alternative that allows for smoother revenue functions is to fit the valuation data to a parametric distribution.

We explored the multivariate normal, and while the approach has the appeal of simplicity, we are concerned about various aspects of fit. First, the probability of having a positive valuation appears not to be governed by the same underlying index that generates the positive valuations. Second, the positive values appear to be better described by a lognormal than a normal distribution. Accordingly, we employ a "zero-inflated multivariate lognormal" model.

That is, we assume that positive valuations are distributed multivariate lognormal. Thus, individual i's log valuation of song (s) is $v_{is} = \mu_s + \varepsilon_{is}$. The parameter σ_s is the standard deviation of log valuations for song (s). Finally, ρ_{st} is the correlation of valuations between song (s) and song (t). We estimate μ_s and σ_s using the positive valuations of song (s). We estimate ρ_{st} directly using logs of valuations for individuals who report positive valuations for both songs (s) and (t).

We then estimate the probabilities that respondents report positive valuations using the following model: $y_{is} = \theta_s + \epsilon_{is}$where y is binary: it equals 1 if the valuation is positive and 0 if the reported valuation is 0. The standard deviation is unity by construction. By estimating this equation for two songs at a time, we obtain estimates of the correlation between the tendencies for each pair of songs to have positive valuations

ρ^I_{st}. The parameters μ and θ are unconstrained, so the probability that respondents report positive valuations for a song can be correlated with the average value that respondents report, but we treat ε and ϵ as uncorrelated.

We then simulate valuations by first simulating multivariate lognormal draws according to our estimated values of μ, σ, and ρ. We exponentiate the resulting log valuation estimates. We then simulate the process generating zero valuations by simulating the indices of y using our estimates of θ_i and ρ^I_{ij}. We "zero-out" the simulated valuations for observations where the probit predicts a zero valuation. This gives us a simulated dataset generated by underlying distributions that are smooth. We then apply pricing algorithms outlined below to these data. For each dataset, we create valuation data on the 50 songs from 5000 simulated individuals. An appendix reports selected pricing exercises on raw data. Figures 1, 3, and 4 include sub-figures with both simulated and raw data, allowing visual comparison.

For each pricing scheme, the optimal price generated by the process is a statistic subject to sampling variation. For the pricing schemes that allow quick calculation of optimal prices, calculating optimal prices from simulated data resulting from bootstrap re-estimation of the parameters { μ, σ,ρ, θ_i , ρ^I_{ij}} using 500 repetitions and re-sampling on individuals.

1. *Raw Data or Parametric Estimates*

A natural interpretation of the bunching of reported valuations at multiples of \$0.25 is that respondents tend to round their reported valuations, in particular that respondents often report, say, \$1.25, when their true valuation is between \$1.125 and

\$1.375. In such rounded data, the implied quantity sold at a price of \$1.25 overstates the quantity of songs for which respondents actually value at or above \$1.25 because the quantity of songs with reported valuations at or above \$1.25 includes songs with valuations as low as \$1.125. Likewise, the implied quantity sold at \$1.20 is overstated, because some valuations that are truly below \$1.20 are rounded up to \$1.25 in the observed data. In contrast, the implied quantity sold at \$1.26 will be understated, because the quantity of songs reportedly valued at \$1.26 excludes some songs that are truly valued between \$1.26 and \$1.375. This sort of bunching gives rise to plateaus on the demand function and ensuing spikes on the revenue function. Thus the raw data will exaggerate the revenue available from uniform pricing at multiples of \$0.25 and values slightly below these multiples, and understate the revenue available at prices slightly higher than a multiple of \$0.25. If the spikes are sufficiently large (i.e. when rounding is common), the estimated maximum revenue available will exceed the true maximum revenue available even if the true optimal price is slightly higher than a multiple of \$0.25.

Optimal bundle prices depend on individuals' valuations of bundles of songs. Valuations of bundles of songs are sums of the often-rounded-to-the-nearest-quarter valuations of particular songs. The process of summing valuations will tend to average out the rounding error in particular songs, removing misleading spikes from the bundle-pricing revenue functions. Thus, bundle pricing on (sums of) raw data will not overstate the profit available from offering a bundle at any particular price.

Uniform pricing provides the benchmark we use to evaluate other pricing approaches. The use of raw data and its resulting spiky revenue functions overstates the maximum revenue available with uniform pricing and therefore understates the benefit of pricing techniques that involve individuals' valuations of bundles (bundling, two-part tariffs, and nonlinear pricing). Hence, we believe that the raw data provide misleading answers, and we focus on estimates using the parametric approach throughout the paper.

IV. Results Using Different Pricing Schemes

1. The Single-Price Monopoly Baseline

The four panels of Figure 5 show the empirical demand curves, treating all songs as a single good (music), using both raw and parametric approaches, and using 2008 and 2009 data. The raw data, in the top two panels, produce demand curves with noticeable plateaus, while the parametric approach produces smooth demand curves. Given that marginal cost of digital music is zero, the surplus at stake in this pricing problem is the entire area under the raw-data demand curves, or \$60.34 per person in 2008 and \$43.56 per person in 2009. [12]

We can calculate the profit-maximizing price – and breakdown of the available surplus – with a simple manipulation of the valuation data. We order the valuation data from highest to lowest among the n (song x individual) valuations: $V_1,\ldots V_n$. Because n songs are sold when the price is V_n, we can calculate the revenue when we charge V_n per song as $n^* V_n$. Figure 6 shows the empirical revenue functions relating $n^* V_n$ to n for the various approaches.

---

[12] Music selling has three parties: the artists, their labels, and the retailer. At the time of the surveys, Apple reportedly paid a flat rate of \$0.7 per song. Analysis of the retailer's pricing decision could then take the marginal cost to be \$0.7. We will instead treat the problem as if the three parties acted collectively to maximize the pie that they split. To that end we treat the marginal costs as its technical value, zero.

The revenue functions based on raw data have various local maxima, confirming our concerns about the revenue overstatement due to bunching (the highest of which occurs at a price of $1.99 in the 2008 data and at $0.99 in the 2009 data). The data from the parametric estimates give rise to smooth, single peaked revenue functions, and the revenue-maximizing uniform prices from the parametric approaches are $2.30 in the 2008 sample and $1.46 in the 2009 sample. We discuss the precision of these estimates at section IV.8 below.

Table 2 summarizes the breakdown of total surplus into its constituent parts under uniform pricing with the 2008 and 2009 data, respectively. Based on the 2008 parametric estimates, uniform pricing allows producers to captures 27 percent of surplus, while consumers get 44 percent of surplus, and the remaining 29 percent of the area under the demand curve is deadweight loss. Using the 2009 data gives nearly identical results.

We put the uniform pricing results to two uses. First, they are interesting on their own. Because we don't normally have data on the full distribution of valuations, we cannot generally calculate the shares of surplus accruing to different parties. Instead, we typically appeal to functional form assumptions on the demand curve itself. For example, with linear demand and zero marginal cost, under uniform pricing PS is half of total surplus, while CS and DWL are each a quarter. Because the demand curves here (in Figure 5) are far from linear, results are quite different. Second, and more important, the uniform price results provide the benchmark we use for evaluating more sophisticated pricing approaches. For each approach below, we can ask how the components of surplus grow or shrink relative to their values under uniform pricing.

2. Song-Specific "Component" Pricing

A conceptually simple alternative to uniform pricing across all songs is component pricing (uniform pricing within songs). It might be complicated in practice because it requires song-specific valuation information, but putting this practical complication aside, we can implement this with our data by simply calculating maximum revenue for each song as we calculated maximum revenue overall above. Figure 7 shows histograms of resulting optimal song prices from the two years' samples. The median optimal song price for the 2008 sample is \$2.28, and the inter-quartile range runs from \$1.94 to \$3.00. For 2009 the median is \$1.20, and the inter-quartile range extends from \$0.98 to \$1.52.

Interestingly, in light of the clamor to convince Apple to employ song-specific pricing at iTunes, component pricing – using 50 prices rather than one – has little effect on the components of surplus. It raises PS by 3 percent in both the 2008 and 2009 data based on parametric estimates. Component pricing reduces CS by 6 percent in the 2008 sample and raises CS by 5 percent in the 2009 sample. DWL under component pricing is 6 percent higher than the DWL under uniform pricing using the 2008 sample and 11 percent lower using the 2009 sample. It is notable that component pricing's benefit remains small even with the 2009 sample and its inclusion of older and less popular songs. It should be noted that these estimates of the benefit of component pricing, while small, overstate the practical benefit of song-specific pricing since in actuality sellers would need to set prices for songs in advance of knowing each song's realized demand.

3. Pure Bundling

Another simple alternative to uniform pricing is "pure bundling" (PB), in which the entire group of songs is offered, as a group, for a single price. Since Stigler (1963) and Adams and Yellen (1976), economists have understood the intuition that negative correlations in valuations allow a seller to capture more revenue by bundling products. As we have seen, sample song valuations are positively correlated across individuals, and the correlation is higher for the 2008 than the 2009 data. Schmalensee (1984) shows that bundling can increase revenue even with positive correlations of valuations.[15] And as Bakos and Brynjolffson (1999) and Armstrong (1999) argue, as the number of products grows large, PB will approach perfect price discrimination as long as consumers' tastes aren't too highly correlated across products. How does pure bundling affect surpluses in our context? And how does this vary with the size of the bundles?

To calculate the optimal full (50-song) bundle price we sum the song valuations across songs within each individual to arrive at that individual's valuation of the entire bundle. We then calculate the revenue-maximizing bundle price, as we would under single price monopoly. The optimal 50-song bundle price is $74.25 using the data derived from the parametric approach for 2008 and $36.84 for the 2009 data. Pure bundling raises revenue substantially, by 17 percent relative to uniform pricing with the 2008 sample and by 29 percent using the 2009 sample. These gains to consumers come partly at the expense of consumers: CS under bundling is 15 percent below its value under uniform pricing in the 2008 sample and 5 percent below for the 2009 data. While pure bundling improves revenue – and substantially more than component pricing – the resulting revenue falls far short of the perfect price discrimination revenue predicted by some theoretical models. The reason for this failure is that tastes are correlated across

[15] Fang and Norman (2006) demonstrate similar results.

products, which arises, as Armstrong (1999) notes, "because of income or other systematic differences across consumers."

We also explore how the effect of pure bundling on revenue varies with the size of the bundle, by calculating the maximal revenue for random song bundles, with 500 draws for each of the following bundle sizes: 2, 3, 4, 5, 10, and 25. See Table 3. Each column corresponds to one of the sample years and either the raw or parametric approach.

Using the 2008 sample, bundles of 2 raise revenue 5 percent relative to uniform pricing. As bundle size increases, the benefit of bundling increases. More than half of the benefit is achieved with bundles of 5. Similar results hold using the 2009 sample. While larger bundles raise revenue, the rate of increase declines as bundle size increases. If the flattening continues, it appears that there would not be much more revenue benefit available to bundles larger than 50 (so that 50 is a "large" number, in the sense of the theory) and that most of the revenue benefit of large bundles is achieved with 5-song bundles.

4. Two-Part Tariffs

Another scheme we can explore is the two part tariff with a hookup fee (T) independent of the number of songs purchased and a per-song price (p). We have already explored two of its special cases: When p=0, this is pure bundling, and when T=0, this is uniform pricing. We explore this family of schemes as follows. Try a pair (T, p). Given the p, each individual would purchase some number of songs and would have some level of consumer surplus from the songs with valuations at or above p. If the individual's total consumer surplus exceeds T, he would pay the hookup fee, and then the revenue from that individual would equal p*(number of units purchased) + T. If T exceeds his consumer surplus, by contrast, then he would make no purchase.

Our goal is to find values of (T, p) that maximize revenue. If the revenue function were well-behaved, we could enlist hill-climbing algorithms to find the revenue-maximizing price pair. Despite our parametric estimation and its attendant smoothing, resulting functions remain irregular. There is a multitude of substantially different, though functionally related, two-part pricing schemes that yield similar revenue. The hill climbing program is generally unable to find, or come close to finding, which of these many functionally related pricing pairs maximizes revenue. To approximate the best price pair we begin our maximization with a rather fine grid, searching over 161 values of $p \in \{0, 0.05, 0.1, \ldots, 8\}$ in conjunction with 1000 values of $T \in \{0, 0.1, 0.2, \ldots, 100\}$. We then use the per unit price and fixed fee from each of the top ten revenue generating prices from this grid search as starting points for a hill-climbing program.

Using the 2008 data, the best two-part schemes identified involve (T, p) = (\$52.31, \$0.48). Using the 2009 sample, the best tariff is (\$21.19, \$0.37). It is interesting to examine the top 10 tariffs – and the top 1 percent of tariffs – among those identified by the grid search, along with the best identified. Figure 7 shows these tariffs, and the top tariffs lie along a downward-sloping line between a high unit price with a low hookup fee and a much lower unit price with a higher hookup fee. The two-part tariff achieves results very similar to those achieved with pure bundling - PS, CS, and DWL are almost identical under two-part pricing and bundle pricing.

5. Nonlinear "Bundle Size" Pricing

Two part tariffs are a special case of more general nonlinear prices that vary with the number of units purchased, what Chu, Leslie, and Sorenson (2007) call, "bundle size pricing." As outlined in Wilson (1993), calculating a nonlinear tariff is straightforward in principle. Provided that the tariff crosses each individual's demand curve only once, from below, the nonlinear tariff can be calculated as a sequence of optimal prices, for the first, second, and n^{th} units, in this case for n up to 50. When we perform this exercise with our data, however, the resulting tariff has some problems. First, the optimal prices do not decline monotonically. Second, and more important, the number of buyers of successive numbers of units does not decline monotonically. Indeed, the number of persons buying, say, the 3^{rd} unit exceeds the numbers of buyers of the first unit. Single crossing does not hold, and as a result, the simple method cannot be employed for calculating the nonlinear tariff in this context.

Without a workable simple algorithm for determining the exact tariff, we are left with some other alternatives for computing an approximate nonlinear tariff, including a simple parameterization of the tariff and grid search. Given the computational cost of a fine grid search for a pricing scheme with more than a few prices, we turn to a parameterization.

Prices must be non-increasing in the number of songs purchased, and prices must be positive. This suggests some restrictions. If t is an index for a song's position in the sequence, then a simple parameterization is $p(t) = \alpha * e^{\beta t}$, where $\alpha > 0$, $\beta < 0$. We perform a grid search over $\alpha \in \{0.25, 0.5, 0.75, \ldots, 100\}$, $\beta \in \{-2, -1.95, \ldots, 0\}$ to find the revenue-maximizing tariff in this family, using each of the four data sets. The parameters that maximize revenue using the 2008 and 2009 data, respectively, are α=(12.0, 10.75) and β=(-0.15, -0.25). The associated tariff for 2009 is illustrated in a budget constraint in Figure 9. The prices of the first ten songs are, in order: \$8.37, \$6.52, \$5.08, \$3.95, \$3.08, \$2.39, \$1.87, \$1.45, \$1.13, \$0.88, \$0.67. The per-song price reaches about \$0.10 by about the 30th song in the 2008 data, and by about the 20th song in the 2009 data. The total price of all 50 songs in both datasets is slightly higher than the respective bundle price. Not surprisingly, as Table 2 shows, the nonlinear tariff performs very similarly to pure bundling and two part tariffs: it does slightly better than pure bundling and not quite as well as the two part tariff.[16]

6. Comparing Mixed Bundling with Alternatives

One of the goals of this paper is to determine what share of the area under the demand curve can be appropriated as revenue with sophisticated pricing. Each of the schemes we have considered so far is a special case of mixed bundling (MB), which is known to produce the most revenue. To see how much revenue can be obtained, we need to explore mixed bundling. We do this in two parts. We first explore general mixed bundling (with 2^n-1 prices) for small bundles and using coarse grids for search. We then explore a practical specific form of mixed bundling: selling with a bundle price alongside a uniform à la carte price.

A Simple Implementation of General Mixed Bundling

[16] In principle a nonlinear tariff can reproduce the two part tariff as a special case, so nonlinear pricing should produce weakly better results than the two part tariff. Our restrictive parameterization of the tariff is likely responsible for the nonlinear tariff producing less revenue than the two part tariff.

Because of the large number of bundles, mixed bundling is difficult to implement with 50 products. For this reason, we examine the simpler problem of finding optimal mixed bundling pricing schemes for a much smaller set of products, 3. Even with 3 products, there are 7 bundle prices. To keep the problem manageable, we search over the following grids among prices between 0 and 3 in increments of 0.25 for one song bundles, between 0 and 6 for two song bundles, and between 0 and 9 for three song bundles. We employ a 10% sample of the 2009 parametric dataset to speed computation.

In the course of searching for the best mixed bundling schemes, we also find the best UP, CP, PB, 2-part, and BSP schemes on this coarse grid. This allows us to compare the performance of MB with all other schemes. Because each other scheme is special case of MB, we expect MB to provide the most revenue.

We want systematic insight into the relative performance of each of the pricing schemes, so we need to perform the 3-song analysis repeatedly, on different randomly selected groups of three songs. We repeat the analysis 5 times, and Table 4 reports average results. Results on schemes apart from mixed bundling largely reflect what we have found so far. Component pricing generates a bit – 3 percent - more revenue than uniform pricing. Bundling, two-part tariffs, and nonlinear pricing generate substantially more, at 10, 12, and 13 percent more revenue than uniform pricing, respectively. Mixed bundling is better still. On average, MB generates 15 percent more revenue than uniform pricing.

The glass is both half empty and half full. MB achieves more relative to UP than the next-best schemes. MB raises revenue 20 more than its next best alternative. But even with mixed bundling, revenue’s share of surplus reaches an average of only 34 percent of surplus across the 5 song groups in Table 4 (compared with 29 percent for profit-maximizing uniform pricing).

Bundling alongside à la Carte Sales

We have analyzed pure bundling as though the à la carte option ceased to exist. It seems entirely possible that bundling, if pursued, would be added alongside existing à la carte sales. This possibility suggests a number of alternative pricing schemes to explore that constitute simple versions of mixed bundling. Consider a family of two parameter pricing schemes where p_A is the uniform à la carte price and p_B is the bundle price.

One question is, what unconstrained combination of $\{p_A,\ p_B\}$ maximizes revenue? A second question is, what p_B maximizes revenue, given that p_A is constrained to some value, such as its profit-maximizing uniform price.

The question of the best bundle price to complement an existing à la carte scheme is interesting in contrast to the question surrounding current policy debates concerning cable television.[17] Bundling is the default in that circumstance, and various constituencies are advocating à la carte pricing as a complement. Of course, the optimal values of p_A and p_B are related, so it may be difficult to change one without affecting the other.

As Table 5 indicates, when p_A and p_B are unconstrained, the revenue-maximizing combination using the 2009 data is a bundle price of \$37.30 with a single song price of \$7.50. With prices chosen to maximize revenue, the introduction of an à la carte option does not bring about a low à la carte price. Bundle sales generate the vast majority of revenue (99 percent), and few songs are sold à la carte. Revenue is slightly higher than it is under pure bundling without à la carte sales: Revenue in this case is 29.1 percent above revenue under uniform pricing (compared with 28.8 percent for pure bundling).

Our second exercise is the determination of an optimal bundle price when the à la carte price is constrained to its current level of $1.46. Given this constraint, the best bundle price is $31.05, and bundle sales generate just over two thirds of revenue. This scheme generates 14.5 percent more revenue than uniform pricing, only half the improvement over uniform pricing produced with either pure bundling or unconstrained simple mixed bundling. That is, maintaining the à la carte price at $1.46 significantly handicaps the ability for bundling to raise revenue.

7. Third Degree Price Discrimination

One class of pricing schemes we have not yet explored is schemes that treat people differently according to exogenous characteristics (as opposed to endogenous behavior of how many songs to buy, given the schedule). Examples of these could, in principle, include price discrimination by race, gender, geography, or income. It should be noted that many such forms of price discrimination are both illegal and, at times, morally questionable. Our exploration of this class of pricing schemes is merely aimed at determining what classes of pricing schemes could, in principle, fulfill non-uniform pricing's promise of making surplus appropriable.

Before proceeding further it makes sense to note that the vast majority of the variation in valuations occurs between individuals as opposed to between products (within individuals). This suggests that schemes that can divide consumers according to their valuations will be able to extract more of their valuations as producer surplus.

A conceptually simple scheme is person-specific pricing. Third-degree pricing schemes – price discriminating by type of person – are special cases, so this will give an upper bound on the effect of such schemes on the distribution of surplus. To calculate the person-specific profit-maximizing price, we create person-specific demand curves, ordering their valuations from highest to lowest. We then find the maximum revenue. The valuation associated with maximal revenue for each person is the person-specific revenue-maximizing price. Because our parametric estimates do not include person characteristics, we can only implement third degree price discrimination based on observable characteristics using the raw data. Table 6 reports results for various kinds of third degree price discrimination.

Revenue-maximizing person-specific pricing gives rise to substantial variation in prices. Using the 2008 data, the inter-quartile price range runs from \$0.56 to \$2.78. Using 2009 data, it runs from \$0.61 to \$2.15.

Person-specific pricing gives substantially more benefit to producers than do the non-discriminatory schemes explored above. They raise revenue by two thirds using the 2008 sample and by half using the 2009 sample. This benefit to producers comes at substantial cost to consumers: CS falls by a third relative to its value under uniform pricing in the 2008 data and by a quarter in the 2009 data. The gains to producers under person-specific pricing outweigh the losses to consumers, so deadweight loss falls using both samples.

Person-specific pricing may be difficult to implement if it's hard to know each individual's demand curve *a priori*. This raises the question of what revenue improvement third degree price discrimination schemes based on observable characteristics can achieve. To this end we use the raw data to explore schemes based on the scant observables in our data, using the raw data: gender, ethnicity, whether a respondent is a resident alien, and age (whether under 20). As with the parametric approach, the raw data allow large increases in producer surplus with person-specific prices: 81 percent for 2008 and 68 percent for 2009. But the effects of price discrimination based on observables are small: none raise revenue more than 6 percent, and most accomplish less. Despite the large revenue enhancing effects of individually customized uniform prices, forms of third degree price discrimination that might more feasibly be implemented produce only negligible revenue improvements.

8. Precision

The results that are the main focus of this paper – the shares of surplus going to producers and the portion foregone as deadweight loss, along with the changes in these shares with alternatives to uniform pricing – are each statistics subject to sampling error. This section explores the precision of our estimates.

As mentioned above, our calculated optimal prices (and other ensuing statistics, such as the PS shares under various pricing schemes) are each functions of the estimated parameters of our basic statistical model (the zero-inflated multinomial lognormal model). To get estimates of the precision of our estimates, we have re-estimated the basic model 500 times. For the pricing schemes that allow quick calculation of optimal prices (uniform pricing, component pricing, bundling, and person-specific pricing), we can then calculate revenue-maximizing pricing schemes for each set of parameters.

Table 7 reports estimated standard errors for both the surplus shares (CS, PS, and DWL) and the percent change in these shares relative to uniform pricing based on the 500 bootstrapped parameter estimates. The top and bottom panels report 2008 (2009) estimates, and the surplus shares are estimated precisely. For example, the standard error of the CS share under uniform pricing (roughly 44 percent of surplus) is 1.41 percentage points, while the standard error of the PS share (just over a quarter) is 0.33 percentage points.

The latter columns report the standard errors for the percent changes in surplus components relative to uniform pricing. As Table 2 indicated, component pricing raises 2008 PS by 2.65 percent. Table 7 indicates that the standard error for this PS improvement is 0.33 percent. Bundling raises 2008 revenue by nearly 17 percent, and the standard error for this improvement is 4.87 percent. Similar patterns hold for 2009. Improvements in PS are all large relative to their standard errors. The same cannot be said of the changes in CS and DWL.

Table 8 focuses specifically on the improvement in PS under each scheme, reporting the 5th, 25th, 50th, 75th, and 95th percentiles of the distributions of percentage PS improvements relative to uniform pricing. For example, the table shows that in the 2008 sample, the 5th percentile increase in PS under component pricing is 2.47 percent, while the 95th percentile improvement is 3.59 percent. The table shows that the PS increases are rather precisely estimated.

Although we are unable to calculate standard errors for measures surplus shares under nonlinear pricing and two-part tariffs, we note that these schemes yield surplus shares – and improvements over uniform pricing – that are very similar to those achieved by bundling, for which we can calculate measures of precision.

V. Pareto Improving Pricing Schemes

Within each family of prices (e.g. single price uniform pricing or two part tariffs), each particular scheme gives rise to particular values of consumer and producer surplus. For example, when p=0 under uniform pricing, then the entire surplus is distributed to consumers. By contrast, as p rises, consumer surplus falls as producer surplus rises, at least to a point. The CS and PS resulting from various uniform prices define a frontier in PS-CS space.

For schemes that raise PS relative to uniform pricing – bundling and related schemes – we can look for ranges that raise both consumer and producer surplus, relative to their values under uniform pricing. Pure bundling, which produces results nearly identical to two part tariffs and nonlinear pricing, provides a ready example. Using the 2009 sample, we have shown that we can raise PS by 29 percent while reducing CS by 5 percent. Figure 10a depicts a locus of possible changes in PS and CS, relative to their values under profit maximizing uniform pricing. For example, it is possible to raise PS nearly 30 percent while holding consumers – as a group – harmless. By contrast, with a lower bundle price, aggregate consumer surplus rises. Indeed, it is possible to raise consumer surplus by over 50 percent while maintaining revenue at its maximized level under uniform pricing (with a bundle price of \$16.95, compared with its revenue-maximizing value of \$36.84).

While it is interesting that there exist bundle pricing schemes that could raise PS while holding overall CS constant at its level under revenue-maximizing uniform pricing, this is arguably the wrong Pareto-improvement calculation. Apple Corp. derives revenue from selling hardware to individuals whose surplus from music buying is sufficiently large to make the complementary hardware attractive. What matters is therefore not aggregate CS but rather the number of persons who find their CS sufficiently large to induce hardware purchase.

For each potential bundle price, we can determine the share of consumers experiencing higher CS than they would under revenue maximizing uniform pricing. Figure 10b depicts the relationship between the percent improvement in PS (relative to uniform pricing) and the share of consumers better off. Each point along the curve represents a different bundle price. As the bundle price falls, producer surplus initially rises, until we reach the profit-maximizing bundle price. As the bundle price continues to fall below the profit-maximizing bundle price, more consumers obtain higher surplus with the bundle than they would under uniform pricing. The revenue maximizing pure bundling price (which lowers CS by 5 percent) makes roughly two thirds of consumers worse off relative to their well-being under uniform pricing. But lower bundle prices, which generate less PS, make higher shares of consumers better off. At the extreme, nearly 80 percent of consumers could be made better off than under uniform pricing while holding producers harmless at their uniform pricing PS. With a bundle price of $28.55, half of consumers would be better off than under profit-maximizing uniform pricing, and revenue would be 24.6 percent above its uniform pricing maximum. Thus, alternatives to uniform pricing can make both sellers and most consumers better off.

VI. Conclusion

Using survey data on individuals' valuations of popular songs in 2008 and 2009, we are able to directly calculate the revenue – and overall division of surplus – from various pricing schemes. We have two major results, one positive and one negative.

First, various alternative schemes – including simple schemes such as pure bundling and two-part tariffs –raises producer consumer surplus by about 17 percent in one sample and by about 30 percent in another.

Uniform pricing delivers about a third of surplus as revenue, and while various alternative pricing schemes can raise revenue substantially, none of the self-selecting schemes raise revenue's share of surplus above 37 percent. Although individual-specific pricing – an extreme form of third degree price discrimination – raises revenue by 50 percent or more, third degree price discrimination based on available observable criteria raises revenue very little. Hence, results based on these data indicate that even with sophisticated pricing, much of the area under the demand curve for this product is beyond the reach of appropriation by sellers. For products with substantial fixed costs, this leaves open the possibility of inefficient under-provision.

We are aware that this analysis covers a particular product and a small and potentially unrepresentative sample of consumers. Further study with other samples and other products can help clarify our understanding of both pricing strategy and surplus appropriation in product markets.

Table 1: Survey Songs and their Valuations, 2008 Sample

Song name	mean	25th pctile	median	75th pctile
Apologize (feat. OneRepublic) - Timbaland	$2.37	$0.59	$1.39	$2.67
Big Girls Don't Cry (Personal) - Fergie	$1.16	$0.08	$0.53	$1.22
Bubbly - Colbie Caillat	$1.47	$0.08	$0.68	$1.73
Clumsy - Fergie	$0.78	$0.04	$0.29	$1.01
Crank That (Soulja Boy) - Soulja Boy Tell 'Em	$2.00	$0.28	$1.01	$2.10
Crushcrushcrush - Paramore	$0.58	$0.01	$0.13	$0.71
Cyclone (feat. T-Pain) - Baby Bash	$1.29	$0.08	$0.56	$1.45
Don't Stop the Music - Rihanna	$1.40	$0.11	$0.63	$1.44
Feedback - Janet	$0.63	$0.01	$0.11	$0.57
Hate That I Love You (feat. Ne-Yo) - Rihanna	$1.30	$0.10	$0.55	$1.47
Hero/Heroine (Tom Lord-Alge Mix) - Boys Like Girls	$0.77	$0.02	$0.26	$1.00
Hey There Delilah - Plain White T's	$2.02	$0.15	$0.94	$2.02
How Far We've Come - Matchbox Twenty	$1.41	$0.10	$0.69	$1.47
Hypnotized (feat. Akon) - Plies	$1.15	$0.06	$0.48	$1.12
I Don't Wanna Be In Love (Dance Floor Anthem) - Good Charlotte	$1.06	$0.06	$0.47	$1.20
Into the Night (feat. Chad Kroeger) - Santana	$1.49	$0.09	$0.71	$1.53
Kiss Kiss (feat. T-Pain) - Chris Brown	$1.45	$0.12	$0.85	$1.70
Love Like This - Natasha Bedingfield	$1.04	$0.06	$0.43	$1.06
Love Song - Sara Bareilles	$1.02	$0.05	$0.37	$1.07
Low (feat. T-Pain) - Flo Rida	$1.60	$0.11	$0.88	$1.93
Misery Business - Paramore	$0.69	$0.01	$0.17	$0.90
No One - Alicia Keys	$1.59	$0.13	$0.83	$1.86
Our Song - Taylor Swift	$0.81	$0.01	$0.12	$0.80
Over You - Daughtry	$1.22	$0.05	$0.47	$1.12
Paralyzer - Finger Eleven	$1.11	$0.03	$0.34	$1.17
Piece of Me - Britney Spears	$0.77	$0.01	$0.11	$0.85
Ready, Set, Don't Go - Billy Ray Cyrus feat. Miley Cyrus	$0.59	$0.00	$0.09	$0.58
Rockstar - Nickelback	$1.39	$0.06	$0.50	$1.47
S.O.S. - Jonas Brothers	$0.68	$0.01	$0.15	$0.76
See You Again - Miley Cyrus	$0.68	$0.00	$0.09	$0.59
Sensual Seduction (Edited) - Snoop Dogg	$1.18	$0.04	$0.29	$1.07
Shadow of the Day - Linkin Park	$1.24	$0.07	$0.52	$1.23
Sorry - Buckcherry	$0.64	$0.00	$0.13	$0.76
Start All Over - Miley Cyrus	$0.47	$0.00	$0.08	$0.32
Stay - Sugarland	$0.64	$0.00	$0.10	$0.59
Stop and Stare - OneRepublic	$1.05	$0.07	$0.44	$1.10
Stronger - Kanye West	$2.79	$0.87	$1.74	$3.04
Sweetest Girl (Dollar Bill) [feat. Akon, Lil Wayne & Niia] - Wyclef Jean	$1.79	$0.14	$0.88	$1.98
Take You There - Sean Kingston	$1.37	$0.13	$0.78	$1.58
Tattoo - Jordin Sparks	$0.94	$0.04	$0.39	$1.00
Teardrops On My Guitar - Taylor Swift	$0.92	$0.01	$0.17	$0.93
The Great Escape - Boys Like Girls	$1.11	$0.05	$0.44	$1.25
The Way I Am - Ingrid Michaelson	$0.91	$0.02	$0.26	$0.97
The Way I Are (feat. Keri Hilson & D.O.E.) - Timbaland	$2.24	$0.42	$1.13	$2.61
Through the Fire and Flames - Dragonforce	$0.73	$0.00	$0.11	$0.90
Wake Up Call - Maroon 5	$1.55	$0.17	$0.87	$1.92
When You Were Young - The Killers	$1.61	$0.17	$0.90	$1.98

Witch Doctor - Alvin and the Chipmunks	\$0.69	\$0.00	\$0.08	\$0.43
With You - Chris Brown	\$1.34	\$0.08	\$0.49	\$1.14
Won't Go Home Without You - Maroon 5	\$1.43	\$0.17	\$0.86	\$1.57

Notes: The list is the top 50 songs on iTunes January 11, 2008. Respondents indicated their maximum willingness to pay for each song from its hypothetical sole authorized source.

Table 2: Alternative Pricing Approaches using Parametric Estimates

Pricing Method	PS share	CS share	DWL share	% ΔPS	% ΔCS	% Δ DWL	sample
Uniform	27.0%	44.2%	28.8%				2008
Component	27.7%	41.7%	30.6%	2.65%	-5.70%	6.27%	2008
Pure Bundling	31.5%	37.6%	30.9%	16.61%	-14.87%	7.27%	2008
Two-Part Tariff	31.8%	39.2%	28.9%	17.97%	-11.27%	0.45%	2008
Nonlinear	31.8%	38.3%	29.9%	17.72%	-13.28%	3.79%	2008
Uniform	28.4%	42.4%	29.2%				2009
Component	29.2%	44.7%	26.1%	2.90%	5.42%	-10.66%	2009
Pure Bundling	36.5%	40.3%	23.2%	28.81%	-5.00%	-20.68%	2009
Two-Part Tariff	36.9%	43.1%	20.0%	30.30%	1.55%	-31.63%	2009
Nonlinear	36.6%	39.5%	23.9%	29.04%	-6.83%	-18.26%	2009

Note: We calculate the profit-maximizing scheme in each pricing family (uniform, component, etc.). The first three columns show how the total area under the demand curve is divided among producer surplus (PS), consumer surplus (CS), and deadweight loss (DWL). The next three columns show how the components fare under each pricing scheme, relative to profit-maximizing uniform pricing. All exercises are performed on simulated data based on the parametric zero-inflated multivariate lognormal model. See text for details.

Table 3: Bundle Size and Effect on Revenue Relative to Uniform Pricing

Bundle Size	*2008, parametric*	*2009, parametric*
2	4.9%	7.0%
3	7.6%	11.1%
4	9.3%	14.3%
5	10.6%	16.4%
10	13.6%	21.6%
25	15.9%	26.6%
50	16.6%	28.8%

Note: table reports average percent change in revenue by selling via pure bundling with bundles of different sizes. Each entry for bundle sizes below 50 is based on 500 random bundles drawn from all possible bundles.

Table 4: Performance of Mixed Bundling and other Schemes Relative to Uniform Pricing

	PS	*CS*	*DWL*	*Q*
Uniform	1.00	1.00	1.00	1.00
Component	1.03	0.90	1.16	0.92
Pure Bundling	1.10	0.91	1.06	0.49
Two Part	1.12	0.99	0.87	1.22
Nonlinear	1.13	0.97	0.90	1.23
Mixed Bundling	1.15	0.87	1.08	1.17

Note: Based on 10 groups of 3 songs. For each group of songs and each pricing scheme, we calculate profit-maximizing pricing on a coarse grid of possible prices that allows us to find the best schedule for mixed bundling. We then average the results across the 10 groups of 3 songs. See text for details.

Table 5: Simple Mixed Bundling

			Revenue share		% above uniform		
Scheme	P_A	P_B	á la carte	Bundle	PS	CS	DWL
Uniform Only	1.46		100.0%	0.0%			
Pure Bundling		36.84	0.0%	100.0%	28.8%	-5.0%	-20.7%
Bundling with Fixed á la Carte Price	1.46	31.05	27.6%	72.4%	14.5%	20.9%	-44.5%
unconstrained	7.50	37.30	1.4%	98.6%	29.1%	-5.7%	-19.9%

Note: The first two rows repeat the profit-maximizing schemes under uniform pricing and pure bundling, respectively. The latter two rows explore simple mixed bundling. The third finds an optimal bundle price, given an á la carte pricing fixed at the profit-maximizing uniform pricing value. The last row finds unconstrained optimal á la carte and bundle prices.

Table 6: Third Degree Price Discrimination

type	PS rel to uniform	CS rel to Uniform	DWL rel to uniform	approach	sample
person	65.7%	-32.8%	-11.3%	parametric	2008
person	48.0%	-26.2%	-8.6%	parametric	2009
person	81.3%	-50.3%	-13.8%	raw	2008
gender	0.4%	27.0%	-41.4%	raw	2008
ethnicity	5.9%	5.1%	-14.3%	raw	2008
Resident alien	1.7%	33.6%	-52.9%	raw	2008
age	1.0%	-2.6%	2.8%	raw	2008
person	67.9%	-55.8%	43.8%	raw	2009
gender	0.0%	0.0%	0.0%	raw	2009
ethnicity	1.6%	-5.2%	15.1%	raw	2009
Resident alien	0.0%	0.0%	0.0%	raw	2009
age	0.4%	0.6%	-3.1%	raw	2009

Note: The first two rows are based on exercises in which each consumer is charged his or her profit-maximizing price, based on the data simulated from the parametric estimates for 2008 and 2009, respectively. The remainder of the table explores third-degree price discrimination using the raw data to allow differential pricing based on observable characteristics of consumers.

Table 7: Precision of Estimates of Surplus Shares and Percent Increase under Alternative Pricing

	Standard error of Percent of Surplus			Standard Error of Percent Increase relative to UP			
Pricing scheme	**PS**	**CS**	**DWL**	**PS**	**CS**	**DWL**	sample
uniform	0.33	1.41	1.53				2008
component	0.34	1.13	1.29	0.33	2.7	3.79	2008
bundling	1.54	2.86	3.19	4.87	7.28	10.26	2008
person-specific	1.57	0.98	2.08	7.06	2.99	8.5	2008
uniform	0.47	0.96	1.07				2009
component	0.49	0.54	0.7	0.41	2.23	3.76	2009
bundling	1.37	2.21	2.26	4.27	5.43	9.14	2009
person-specific	0.67	0.76	1.25	3.14	2.17	5.18	2009

Notes: Standard errors are calculated from 500 optimal pricing schemes resulting from 500 bootstrap estimates of the parameters of the zero-inflated multivariate lognormal model.

Table 8: Bootstrap Distribution of Improvement in Producer Surplus Relative to Uniform Pricing

	Distribution of PS Percent Improvement					
	5th	25th	median	75th	95th	
component	2.47	2.73	2.95	3.19	3.59	2008
bundling	9.65	12.93	15.66	19.04	25.24	2008
person-specific	57.32	63.4	68.48	73.58	80.62	2008
component	2.37	2.7	2.94	3.27	3.75	2009
bundling	19.36	23.52	26.4	29.37	33.59	2009
person-specific	44.36	47.06	49.34	51.29	54.74	2009

Notes: Standard errors are calculated from 500 optimal pricing schemes resulting from 500 bootstrap estimates of the parameters of the zero-inflated multivariate lognormal model.

References

Adams, W.J. and J.L. Yellen (1976): "Commodity Bundling and the Burden of Monopoly," Quarterly Journal of Economics, 90, 475-98.

Armstrong, M. (1999): "Price Discrimination by a Many-Product Firm," Review of Economic Studies, 66(1), Special Issue: Contracts, 151-68.

Bakos, Y. and E. Brynjolfsson (1999): "Bundling Information Goods: Pricing, Profits and Efficiency," Management Science, 45(12), 1613-30.

Christman, Ed. "Dollars and Cents: iTunes Store." Billboard Magazine. March 15, 2008. (accessed at http://www.billboard.biz/bbbiz/content_display/magazine/upfront, March 14, 2008).

Crawford, G.S. "The Discriminatory Incentives to Bundle in the Cable Television Industry" "*Quantitative Marketing and Economics*, v6n1 (March 2008), 41-78.

Crawford, Gregory S. and Joseph Cullen. "Bundling, Product Choice, and Efficiency: Should Cable Television Networks Be Offered à la Carte?" *Information Economics and Policy*, vol 19, issues 3-4, (October 2007), 379-404.

Chu, Chenghuan Sean, Phillip Leslie, and Alan Sorenson. "Nearly Optimal Pricing for Multiproduct Firms." Working paper, 2007. http://www.stanford.edu/~pleslie/bundling.pdf .

Diamond, Peter A and Jerry A. Hausman. "Contingent Valuation: Is Some Number Better than No Number?" Journal of Economic Perspectives, vol 8, number 4, fall 1994, pp 45-64.

Edlin, Aaron S.; Mario Epelbaum; Walter P. Heller, "Is Perfect Price Discrimination Really Efficient?: Welfare and Existence in General Equilibrium." *Econometrica*, Vol. 66, No. 4 (Jul., 1998), pp. 897-922.

Fang, H. and P. Norman (2006): "To Bundle or Not to Bundle," RAND Journal of Economics, Winter 2006.

Geng, X., M.B. Stinchcombe, and A.B. Whinston (2005): "Bundling Information Goods of Decreasing Value," Management Science, 51(4), 662-67.

Hanson, Ward and R. Kipp Martin. "Optimal Bundle Pricing." Management Science. Vol 36, no. 2. February 1990, 155-174.

Heal, Geoffrey. “Spatial Structure in the Retail Trade: A Study in Product Differentiation with Increasing Returns.” Bell Journal of Economics, Autumn 1980. Vol. 11, Iss. 2; pg. 565.

Jedidi, K., S. Jagpal and P. Manchanda (2003): “Measuring Heterogeneous Reservation Prices for Product Bundles,” Marketing Science, 22(1), 107-30.

Kalish, Shlomo and Paul Nelson. “A Comparison of Ranking, Rating, and Reservation Price Measurement in Conjoint Analysis.” Marketing Letters 2:4 (1991): 327-335.

Knetsch, Jack L & Sinden, J A, 1984. “Willingness to Pay and Compensation Demanded: Experimental Evidence of an Unexpected Disparity in Measures of Value,” The Quarterly Journal of Economics, MIT Press, vol. 99(3), pages 507-21, August.

Leslie, P. (2004): “Price Discrimination in Broadway Theater,” RAND Journal of Economics, 35(3), 520-41.

Long, J.B. (1984): “Comments on Gaussian Demand and Commodity Bundling,” The Journal of Business, 57(1), S235-S246.

McAfee, R.P., J. McMillan and M. Whinston (1989): “Multiproduct Monopoly, Commodity Bundling, and Correlation of Values,” Quarterly Journal of Economics, 114, 371-84.

McMillan, R. (2007): “Different Flavor, Same Price: The Puzzle of Uniform Pricing for Differentiated Products,” Mimeo, Federal Trade Commission.

Musgrove, Mike. “Itunes to Sell Songs Without Restrictions; Four Record Companies Back Tiered Pricing Plan,” The Washington Post, Financial, P. D01, January 7, 2009.

Orbach, Barak Y. and Liran Einav. “Uniform Pricing for Differentiated Goods: The Case of the Movie Theater Industry.” International Review of Law and Economics 27 (2007): 129-153.

Pindyck, Robert S. and Daniel L. Rubinfeld, Microeconomics, sixth ed. Pearson Prentice Hall: Upper Saddle River, NJ. 2006.

Rob, Rafael and Joel Waldfogel. “Piracy on the High C’s: Music Downloading, Sales Displacement, and Social Welfare.” *Journal of Law & Economics* (2006):

Schmalensee, R. (1984): “Gaussian Demand and Commodity Bundling,” Journal of Business, 57(1), 211-231.

Stigler, G.J. (1963): “United States v. Loew's Inc.: A Note on Block Booking,” Supreme Court Review, 152-7.

Sundararajan, Arun. (2004) “Nonlinear Pricing of Information Goods.” *Management Science* 50 (12), 1660-1673.

Venkatesh, R. and V. Mahajan (1993): “A Probabilistic Approach to Pricing a Bundle of Products or Services," Journal of Marketing Research, 30(4), 494-508.

Wilson, Robert B. *Nonlinear Pricing*. New York: Oxford University Press. 1993.

Figure 1: Distribution of Valuations on [0,2]

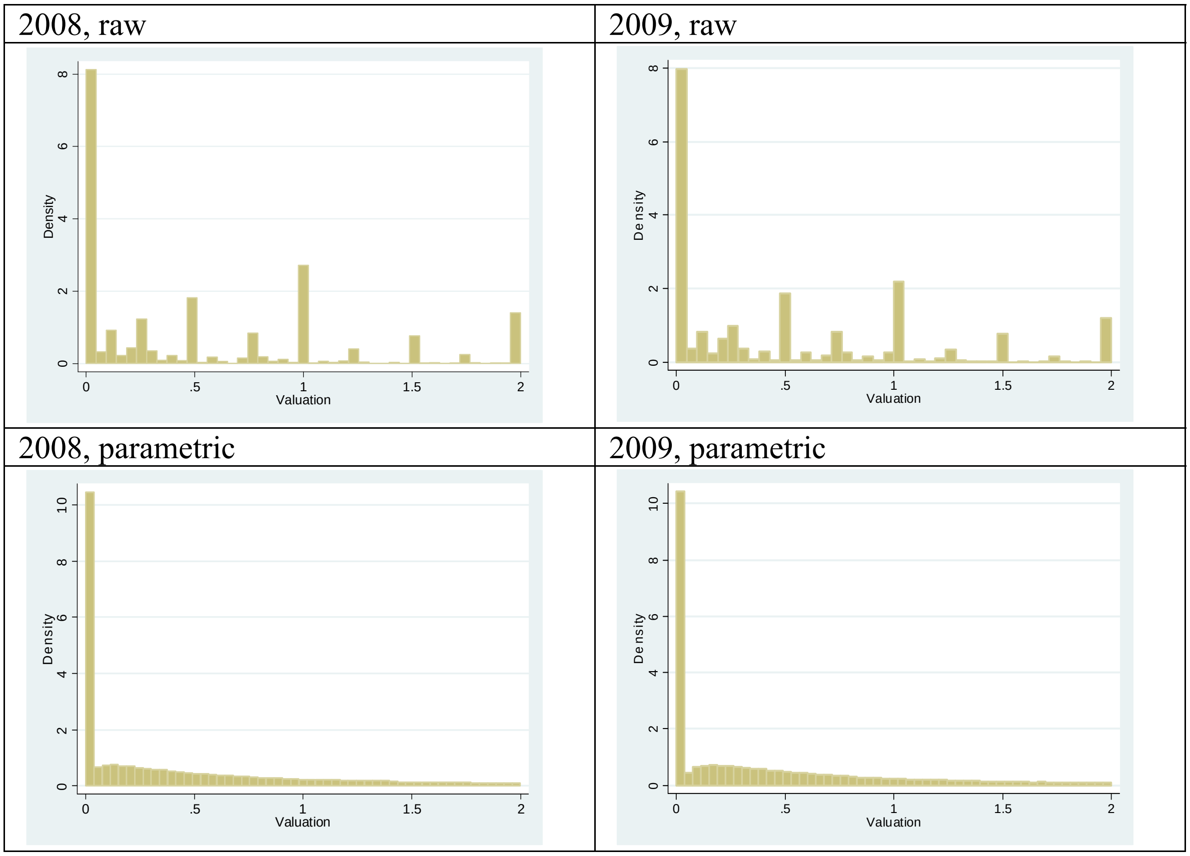

Notes: the distribution of valuations between $0 and $2. The top panels report raw data. The bottom panels employ the simulated data derived from our zero-inflated multivariate lognormal estimates.

Figure 2: Song Valuation and Ownership, 2009 Sample

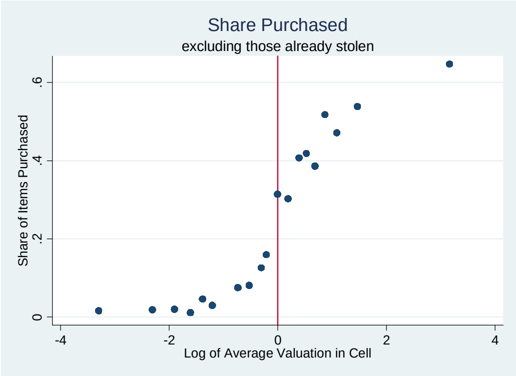

Note: Figure excludes songs obtained via file-sharing.

Figure 3: Valuations across Individuals and Quantities of Music

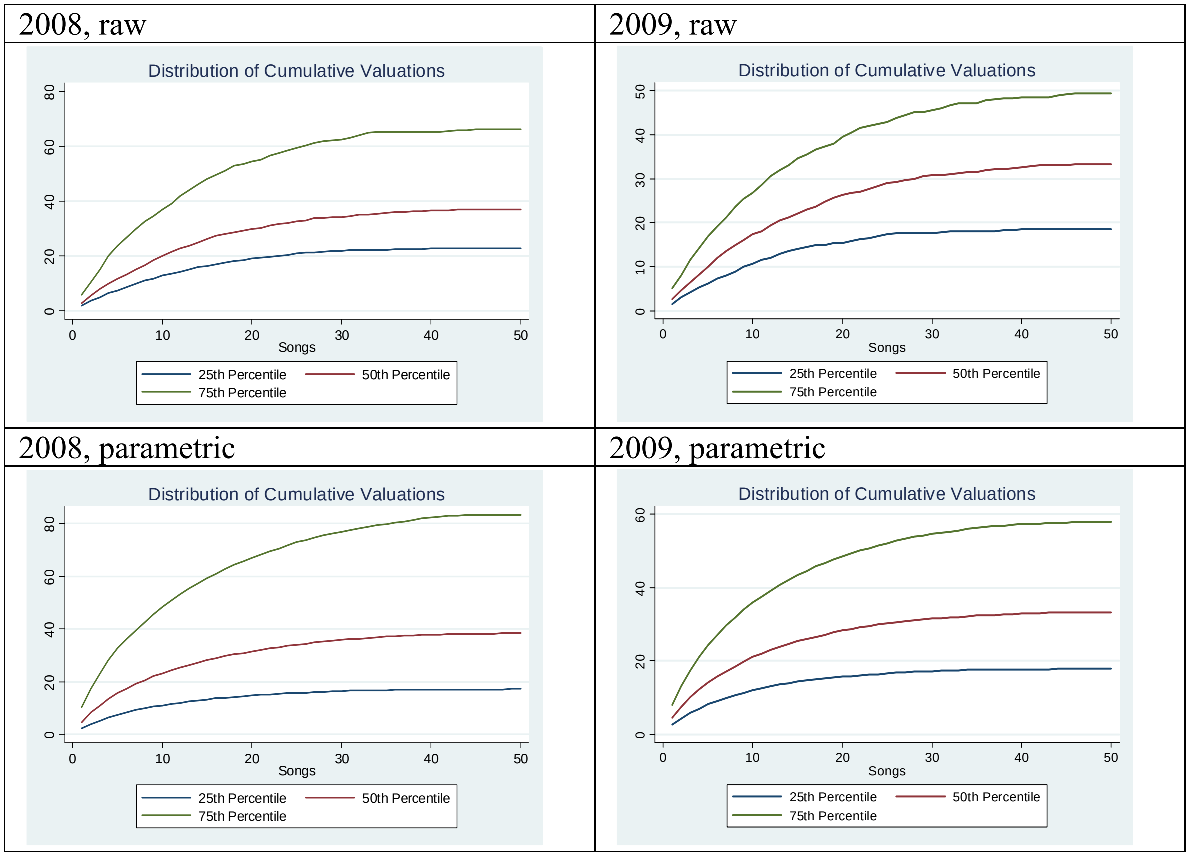

Notes: For any number of songs *x*, the figures depict the 25th, 50th, and 75th percentile of the distribution of the valuations of the *x* songs that the individual respondents value most highly. The top panels are based on raw data. The bottom panels are based on the simulated data derived from our zero-inflated multivariate lognormal estimates.

Figure 4: Pairwise Correlation of Song Valuations

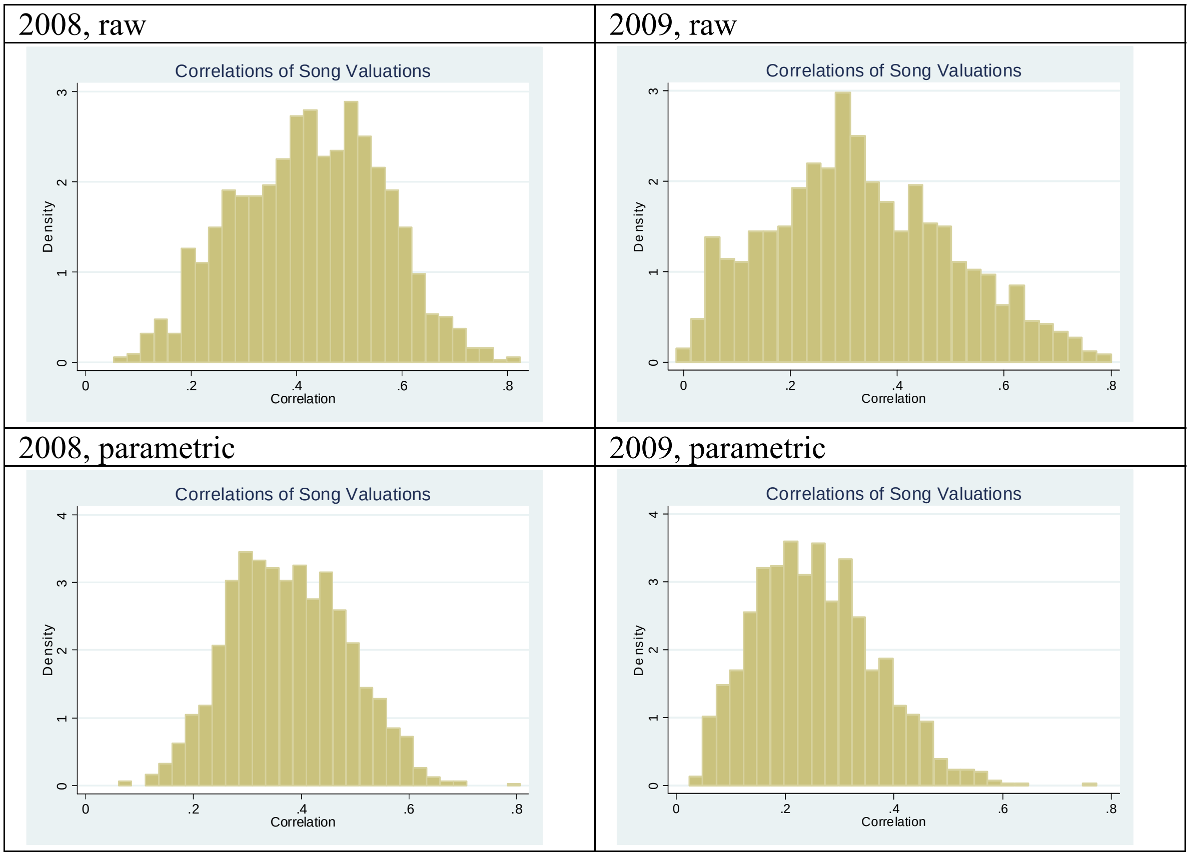

Notes: Figures report the distribution of pairwise correlations in valuation for the 50 songs. Each histogram contains 1225 (50*49/2) correlations. The top panels are based on raw data. The bottom panels are based on the simulated data derived from our zero-inflated multivariate lognormal estimates.

Figure 5: Overall Demand Curve

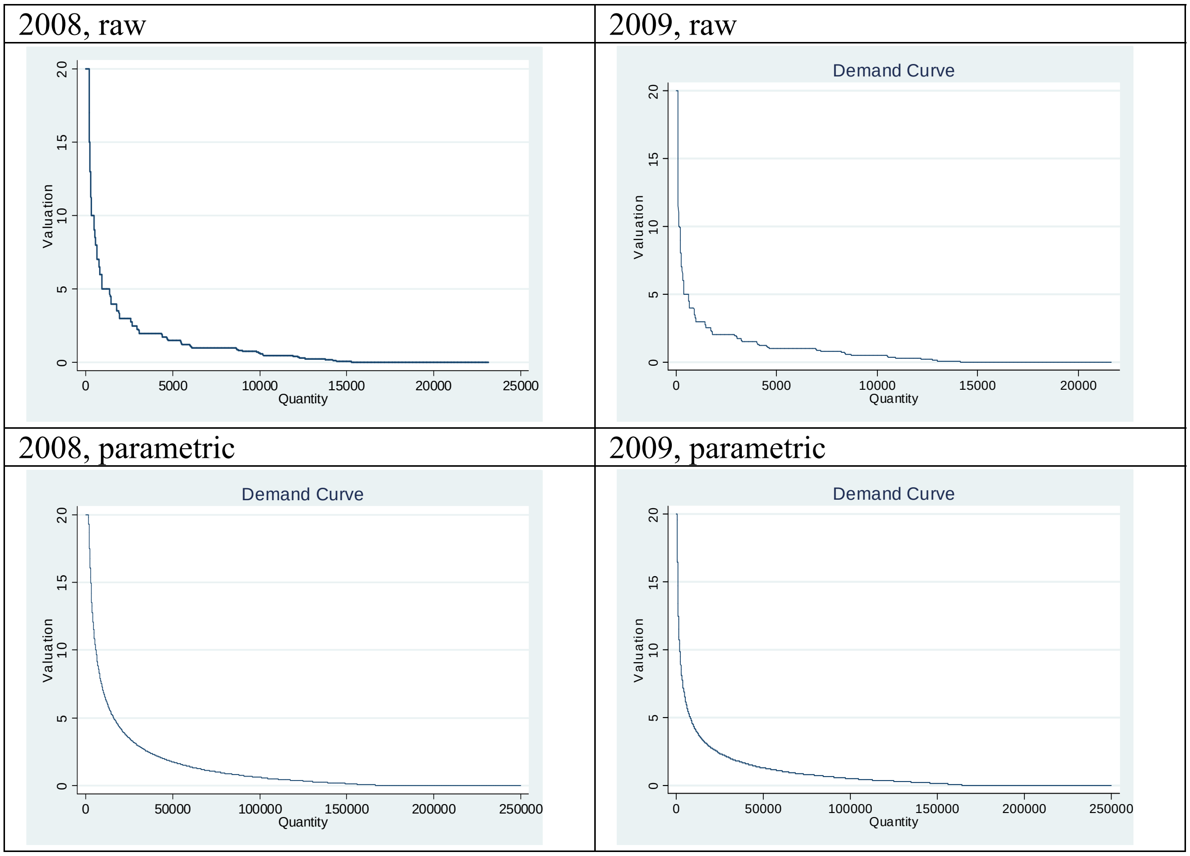

Notes: Figures report the demand curves constructed by ordering individual song valuation observations from highest to lowest. The top panels are based on raw data. The bottom panels are based on the simulated data derived from our zero-inflated multivariate lognormal estimates.

Figure 6: Single-Price Revenue Function

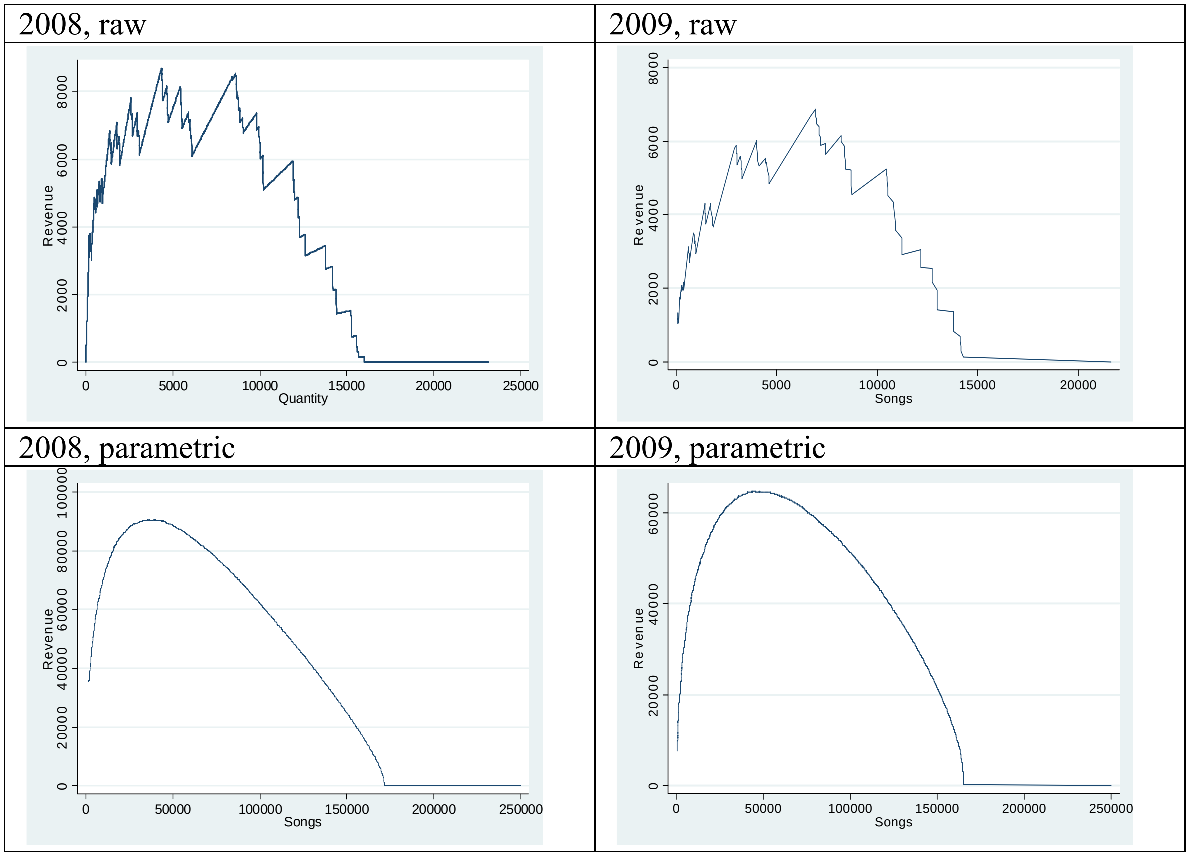

Notes: Figures report the revenue functions curves constructed by ordering individual song valuation observations from highest to lowest, defining $V(n)$. The maximum revenue available for selling any quantity n is then $nV(n)$. The top panels are based on raw data. The bottom panels are based on the simulated data derived from our zero-inflated multivariate lognormal estimates.

Figure 7: Distribution of Song-Specific Revenue-Maximizing Prices

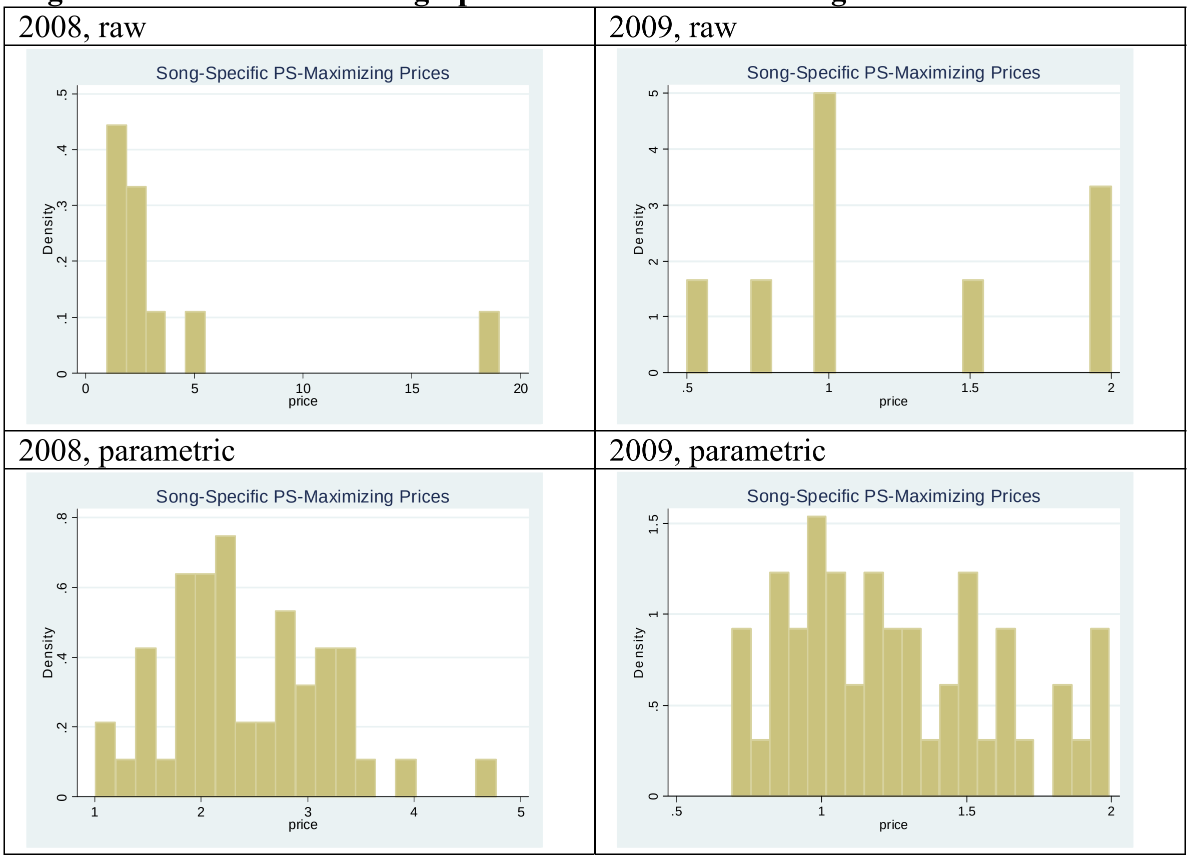

Notes: Figures report the distribution of song-specific revenue maximizing prices. The top panels are based on raw data. The bottom panels are based on the simulated data derived from our zero-inflated multivariate lognormal estimates.

Figure 8: Top Two-Part Tariffs

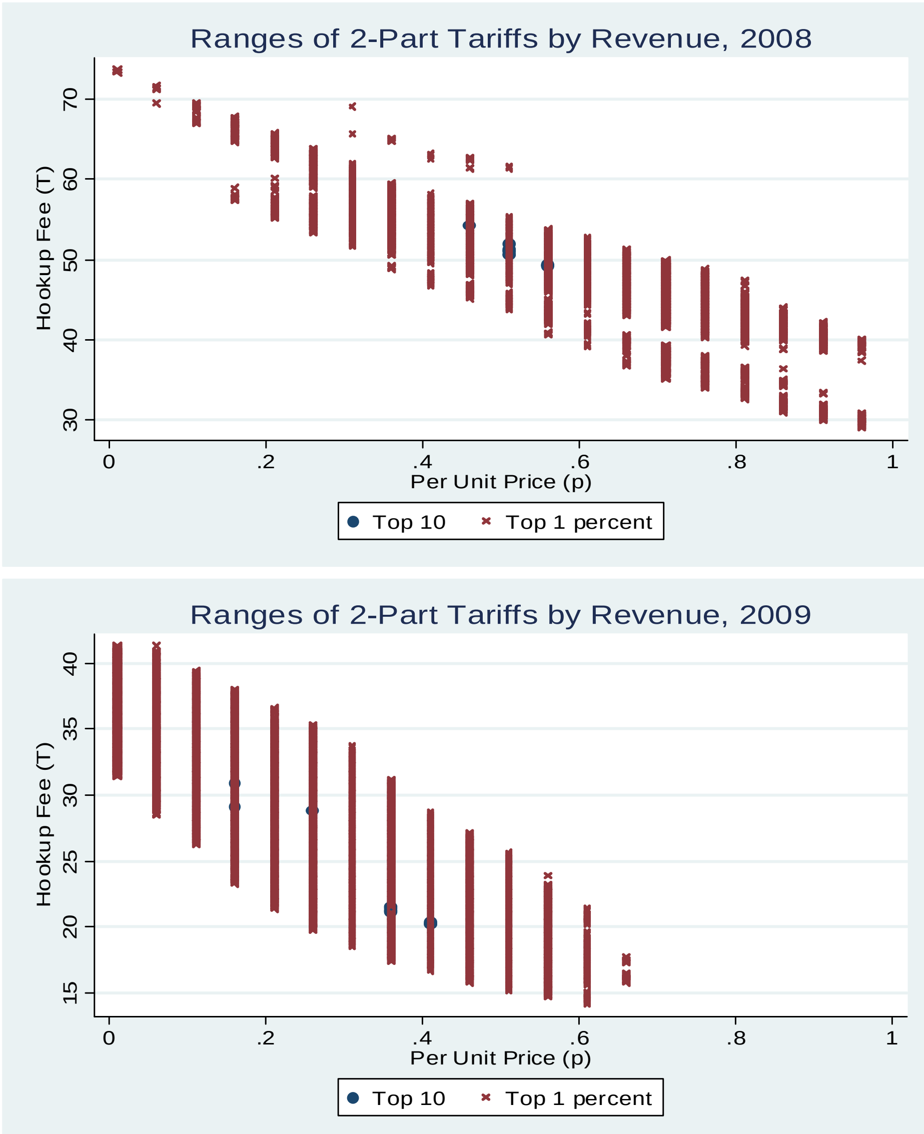

Notes: The figures report the best 10, and the best 1 percent of two-part tariffs identified by grid search.

Figure 9: Budget Constraint with Nonlinear Price Schedule (2009 sample)

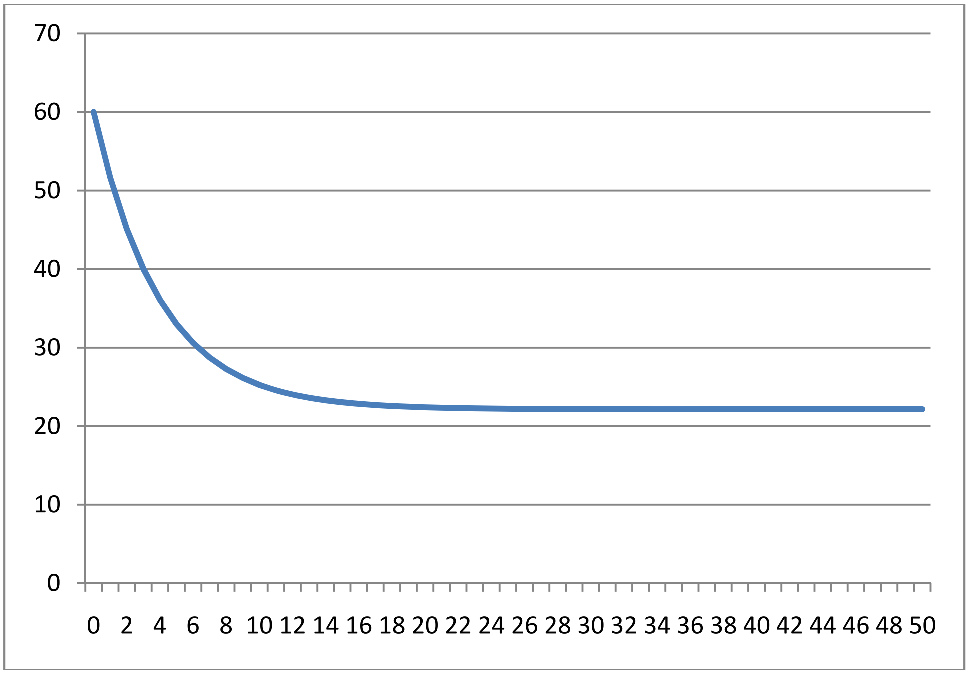

Note: The figure depicts a budget constraint with an endowment of $60. The vertical axis shows dollars, while the horizontal axis shows the number of songs the individual can purchase.

Figure 10: The Tradeoff between CS and PS Improvement in Bundle Pricing (2009 Sample)

10a

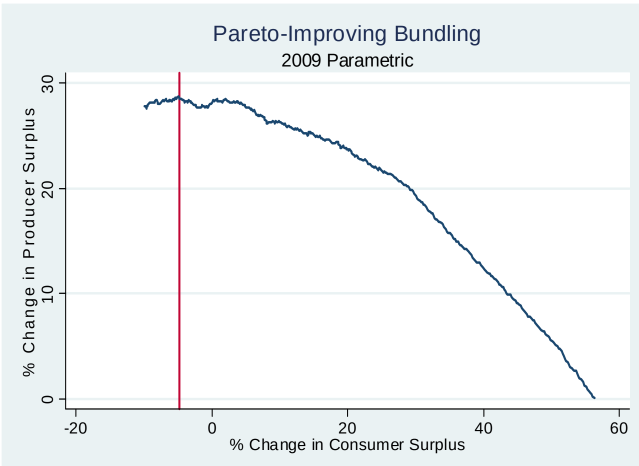

Note: Figure shows the frontier relating producer and consumer surplus across pure bundling schemes differing in their price. The vertical line identifies the profit-maximizing scheme. As the price declines, the PS improvement declines, and CS improvement increases.

10b

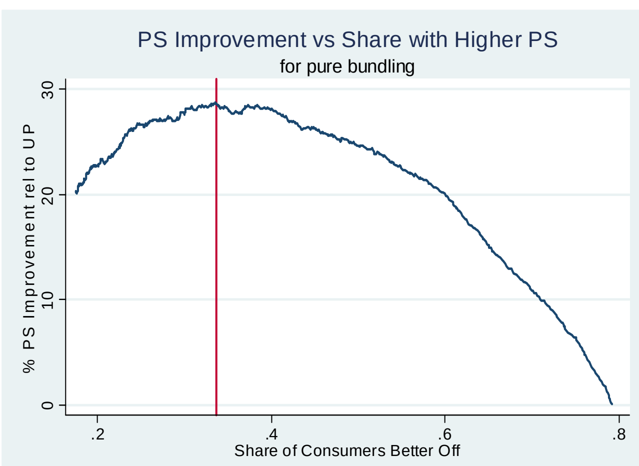

Note: Figure shows the frontier relating total producer surplus with the share of consumers better off under a bundling scheme (relative to uniform pricing) across pure bundling schemes differing in their price. The vertical line identifies the profit-maximizing scheme. As the price declines, the PS improvement declines, and the share of consumers better off under the bundling scheme increases.

Appendix

Table A1: Alternative Pricing Approaches using Raw Data

Pricing Method	**PS share**	**CS share**	**DWL share**	**% ΔPS**	**% ΔCS**	**% Δ DWL**	
Uniform	30.6%	41.8%	27.6%				2008
Component	32.0%	49.0%	19.0%	4.60%	17.19%	-31.19%	2008
Pure Bundling	30.8%	53.5%	15.7%	0.62%	27.82%	-42.91%	2008
Two-Part Tariff	32.3%	53.0%	14.7%	5.50%	26.80%	-46.79%	2008
Nonlinear	30.9%	53.5%	15.6%	1.06%	27.83%	-43.42%	2008
Uniform	36.4%	47.9%	15.7%				2009
Component	37.2%	43.5%	19.4%	2.13%	-9.30%	23.46%	2009
Pure Bundling	38.3%	41.3%	20.4%	5.26%	-13.87%	30.15%	2009
Two-Part Tariff	38.9%	39.9%	21.2%	6.90%	-16.77%	35.22%	2009
Nonlinear	38.3%	41.2%	20.5%	5.25%	-13.99%	30.55%	2009

Table A3: Bundle Size and Change in PS, Raw Data

Bundle Size	*2008, raw*	*2009, raw*
2	-3.3%	-1.9%
3	-3.9%	-1.9%
4	-3.6%	-1.6%
5	-3.1%	-1.3%
10	-1.1%	0.4%
25	0.7%	3.1%
50	0.6%	5.3%